



MIT Sloan School of Management

MIT Sloan School Working Paper 4710-08
2/13/2008

Preventing Accidents and Building a Culture of Safety: Insights from a Simulation Model

John Lyneis, Stuart Madnick

© 2008 John Lyneis, Stuart Madnick

All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission, provided that full credit including © notice is given to the source.

This paper also can be downloaded without charge from the
Social Science Research Network Electronic Paper Collection:
<http://ssrn.com/abstract=1256143>

**Preventing Accidents and Building a Culture of Safety:
Insights from a Simulation Model**

**John Lyneis
Stuart Madnick**

Working Paper CISL# 2008-03

February 2008

Composite Information Systems Laboratory (CISL)
Sloan School of Management, Room E53-320
Massachusetts Institute of Technology
Cambridge, MA 02142

Preventing Accidents and Building a Culture of Safety: Insights from a Simulation Model

John Lyneis, Stuart Madnick
Sloan School of Management
Massachusetts Institute of Technology
January 2008

Abstract

Research has approached the topic of safety in organizations from a number of different perspectives. On the one hand, psychological research on safety climate gives evidence for a range of organizational factors that predict safety across organizations. On the other hand, organizational learning theorists view safety as a dynamic problem in which organizations must learn from mistakes. Here, we synthesize these two streams of research by incorporating key organizational factors from the safety climate literature into a dynamic simulation model that also includes the possibility for learning. Analysis of simulation results sheds insight into the nature of reliability and confirms the dangers of over-reliance on ‘single loop learning’ as a mechanism for controlling safety behaviors. Special emphasis is placed on strategies that managers might use to encourage learning and prevent erosion in safety behaviors over time.

Introduction

Safety and the prevention of accidents is a topic that has interested both managers and organizational theorists for some time. Accidents or violations of safety regulations are often reported in the news and can bring disastrous consequences to individuals, organizations and to society as a whole, from the loss of human life to a loss of public confidence in the services that organizations provide. Recent notable examples include the deaths of several mine workers in Utah due to the alleged use of unsafe mining practices (Borenstein, 2007), and the death a subway worker in New York City who failed to follow regulations regarding the crossing of live tracks while performing maintenance (Neuman, 2007). Similarly, in the United States conflict in Iraq, poor adherence to safety regulations has contributed to a large number of vehicle accidents, many of which have claimed soldiers’ lives (Minami & Madnick, 2007). Less dramatic examples are even more common: for example, the Boston Globe reported that many high-end Boston restaurants consistently fail to observe safety regulations related to the preparation of food, despite the risk of infecting customers with food poisoning (Nelson and Hankinson, 2007). Given the prevalence of news reports like these, the tremendous losses that organizations risk in the event of an accident, and the emphasis often placed on safety in public statements, why is it that so many organizations do such a poor job of adhering to safety regulations and preventing accidents?

Safety and Organizational Theory

Organizational theorists have debated the topic of safety and accident prevention in organizations for some time. On the one hand, Perrow (1984) argues that in certain ‘high hazard’ industries accident prevention may be impossible, despite the presence of regulations designed to prevent them. Because in such “complex” organizations interactions between parts are numerous, slack is limited and processes are highly time dependent, individuals face an especially difficult time interpreting cues and responding adequately to potential dangers (Perrow, 1984). In

response to this argument, others have argued instead that organizations, including those in high hazard industries, can be designed and can learn to undertake difficult tasks with reliability (Roberts, 1990; Weick, 1987). Roberts (1990) defines “high reliability organizations” and suggests several strategies that these organizations use to overcome the deficiencies cited by Perrow, including culture, redundancy, continuous training, and organizational design. Weick (1987) further argues that reliability can be developed through a culture of “storytelling” that encourages the exchange of rich and varied information through face to face communication.

Social Psychology

Consistent with this second view, there is also a rich literature in social psychology that identifies characteristics of organizations that promote safety. Although the earliest psychological research on occupational safety emphasizes either human error or technical failure as the source of accidents, organizational factors are now widely recognized as having a high importance (Clarke, 2000). A substantial portion of this work, beginning with Zohar’s (1980) influential study, draws on the notion of “safety climate” to explain the behavior of individuals with regard to safety. Zohar defines climate as “perceptions that employees share about their work environment... that serve as a frame of reference for guiding appropriate and adaptive task behaviors (pg. 96).” Climate is hypothesized to mediate the relationship between a number of organizational characteristics, including management commitment to safety, the openness of communication links, and the stability of the workforce, and safety behavior. Thus, beyond the propensity of individuals to engage in safe acts independently, this research introduces the notion of a “shared cognition regarding safety” that can be established within organizations (pg. 101).

Building on Zohar’s original work, a number of studies have developed measures for safety climate and verified the relationship between this construct and the safety behaviors of individuals within an organization. Examples include Griffin & Neal (2000), Brown & Holmes (1986) in manufacturing, Hofmann & Stetzer (1996) in a chemical processing plant, and Mearns et al (1998) in a study of oil and gas production. In addition, Zohar (2000) extends this construct by showing that it applies to groups, in addition to organizations as a whole, through the actions of supervisors. Defining the organizational factors that predict safety climate has also received a lot of attention, with a focus especially on management commitment (Brown & Holmes, 1986; Dedobbeleer & Beland, 1991; Huang et al, 2004), managers safety practices (Naveh et al, 2005), leadership style (Zohar, 2002; Barling 2002), work pressure or the conflict between production and safety (Diaz & Cabrera, 1997; Clarke, 2006), overall job satisfaction (Barling & Kelloway, 2003), and the quality of communication or exchange relationships between managers and employees (Hoffman & Morgeson, 1999).

In addition to the construct of safety climate, which emphasizes shared perceptions, a related concept of safety *culture* has also gained popularity in the literature. Drawing on established theories of organizational culture (Schein, 1985), this research looks beyond specific perceptions regarding safety to consider also more general “basic assumptions” that can shape safety behavior (Guldenmund, 2000). Along these lines, Pidgeon (1991) argues that a “good safety culture” will include not only positive norms and attitudes regarding safety, but also a basic reflexivity that allows organizations to discover and learn about potential new hazards (Pidgeon, 1991, cited in Clarke, 2000). Weick’s (1987) notion of “culture as a source of high reliability,” noted above, is very similar. Despite this broader theoretical foundation for culture, however, much safety research continues to operationalize ‘culture’ as “shared attitudes towards safety (Clarke, 2000, pg. 68),” a notion that is not inconsistent with safety climate. For example, Cheyne et al (2002) include a measure of ‘attitude’ in their study of safety climate in two manufacturing firms, and conclude that “general attitudes to safety” do influence safety behavior and may also be seen to “facilitate climate change.”

Organizational Learning

While the literature on safety climate and safety culture goes a long way towards explaining variation in safety practices across organizations, scholars have also had a lot to say about how individual organizations can develop cultures of safety, and why such efforts often fail. In most of the safety climate literature, the question of implementation is not considered, under the implicit assumption that implementation requires little more than the right management actions. For example, this research suggests that to improve safety, managers should place a high emphasis on it, model safe practices, foster effective communication, and limit work pressure to within manageable levels. Reason (1998) goes so far as to suggest that a safety culture may be “socially engineered” (pg. 302) by way of reporting systems and management practices that encourage documentation of “near misses.” According to Reason’s model of accidents (Reason, 1990), organizational factors that weaken organizational defenses against accidents (what he terms the “latent” pathway of accident causation) can be proactively identified and corrected (Clarke, 2000).

In opposition to this view, others have argued that in fact the implementation of a culture of safety is not such an easy task. Here, the focus shifts from prescribing elements of an effective safety culture to managers to an examination of why it is that organizations so often fail to learn from mistakes. In Weick’s words, organizations should be viewed not as “decision makers” but instead as “interpretation systems that generate meaning (Weick, 1987, pg. 123).” While norms of open communication can enable effective interpretation, individuals’ ability to interpret events can also be quite limited, particularly in the context of learning from mistakes. Crucially, individuals and organizations have a tendency to overemphasize “human error” when placing blame, rather than looking beyond to the social and cultural systems that most often also play a role – an example of the fundamental attribution error (Carroll, 1993). This is especially significant given the preponderance of organizational factors shown to influence safety climate and safety culture. If organizations routinely fail to recognize these larger sources of error, it is not surprising that mistakes can go uncorrected.

The distinction between ‘single loop’ and ‘double loop’ learning (Argyris, Putnam & Smith, 1985) is another useful characterization of the trap that organizations sometimes fall into with regard to learning. Single loop learning represents the immediate and local actions that individuals and organizations take in response to a perceived problem. For example, following an accident, such actions might include increasing enforcement of existing rules and procedures to prevent a reoccurrence of the same accident (Carroll, Rudolph & Hatakenaka, 2002). Double loop learning, on the other hand, represents the kind of learning in which assumptions are challenged and mental models are adjusted. Here, instead of focusing on enforcement, individuals might question why rules were not originally followed, whether the existing rules are effective, and whether other underlying causes of an accident might exist (Carroll, Rudolph & Hatakenaka, 2002). Clearly, this second type of learning presents the greatest potential benefit in terms of increased organizational safety.

Despite this potential, double loop learning presents a difficult challenge to organizations. For one, this type of learning threatens existing bureaucratic structures and the control system of the organization (Carroll, Rudolph & Hatakenaka, 2002). More importantly, in a dynamic sense the immediate *success* of single loop learning can undermine both the motivation and the perceived need to follow through on more substantial improvement efforts (Repenning & Sterman, 2002; Tucker & Edmondson, 2003). For example, Tucker & Edmondson (2003) argue that because of the gratification that hospital nurses receive from solving problems quickly and efficiently using ‘first order’ means these nurses are less motivated to raise underlying issues with superiors. Similarly, increasing utilization of a plant will yield immediate improvements in output, reducing the perceived need to take poorly performing equipment down to fix underlying maintenance problems (Repenning & Sterman, 2000). In the context of safety, focusing on compliance while blaming individuals for accidents reduces managers’ anxiety (Carroll, 1993).

Thus, two streams of literature provide complementary insights into the problem of safety in organizations. On the one hand, psychological research on safety climate and safety culture illustrates a range of organizational variables that can impact safety behavior. This research does not, however, answer the dynamic question of how safety in an organization can be improved, or why in some instances safety culture may erode over time. Research in organizational learning addresses this gap, highlighting the challenges that organizations face in learning from mistakes in fundamental ways. Too often, though, research in organizational learning focuses on the learning process rather than seeking to draw general lessons regarding the organizational factors that promote safety. While certainly the challenges that organizations face and the causes for particular accidents are likely to be unique and highly context dependent, research on safety culture suggests that general lessons may also exist. For example, if it is known that management commitment and production pressure are strong predictors of safety, how do these variables relate to an organization's ability to learn and sustain improvements in the long run?

Dynamic Simulation

To integrate these two streams of research, we make use of a dynamic simulation model, using the System Dynamics methodology (Sterman, 2000). System dynamics has been used successfully on numerous occasions to model the experience of particular organizations with regard to safety and accidents (Leveson et al, 2005; Cooke, 2003; Minami & Madnick, 2007), and to generate more general theory concerning both the causes of disaster (Rudolph & Repenning, 2002), and accident prevention (Cooke & Rohleder, 2006). In addition, causal loop diagramming, a subset of System Dynamics that entails the generation of qualitative models of feedback relationships, has been used by a number of authors to explain the erosion of safety behaviors (Marais et al, 2006) and the difficulty that organizations have in learning from mistakes (Senge, 1990; Sterman, 1994; Tucker & Edmondson, 2003).

Studies using System Dynamics to study safety have made a number of important theoretical contributions. In the most general sense, Rudolph & Repenning (2002) show that the stock (or accumulation) of interruptions facing organizational actors, and not the severity of any one interruption, most contributes to the occurrence of a disaster or adverse event. This emphasis on stocks and flows is central to the System Dynamics method. In addition, System Dynamics highlights the role of feedback in leading complex organizations towards often unanticipated outcomes. For example, Minami & Madnick (2007) argue that short term efforts to enforce safety behaviors among combat troops may also encourage fatigue and complacency in the long run, thereby undermining the original policy. Similarly, Leveson et. al (2005) use a System Dynamics model to describe and analyze safety culture at NASA during the time leading up to the Columbia disaster, and find that periods of apparent success can encourage heightened risk taking and complacency towards new safety investments. Cooke (2003) investigates similar dynamics in a comprehensive model of a production setting (the Westray mine), and shows how the pressure to produce contributes to feedback that undermines safety.

Unlike these studies that model the experience of a particular organization, our approach instead is to capture general relationships that are known to exist across a range of organizations, drawing particularly from the safety climate literature and qualitative models of organizational learning. We go beyond most qualitative models, however, by placing them within a mathematical simulation of the "physical structure" of task accomplishment within an organization (Repenning & Sterman, 2002). In particular, we show that this physical structure, and the related notion of 'production pressure,' can have a significant influence on the ability of an organization to learn, regardless of the strength of communication norms that may exist, and that the role of decision makers with regard to production pressure is crucial to fostering a culture of learning. Along these lines, we also develop an improved understanding of "management safety actions," "management commitment," and "safety priority," concepts from the safety culture literature that are not usually thought of in a dynamic sense. A dynamic model shows

why commitment to safety *must* involve more than simply enforcing rules and should also recognize the time that is required to complete safety tasks.

In the next section, we describe the assumptions and structure of the proposed model of organizational safety. Then, we provide an analysis of the model behavior, with particular attention towards relevant insights and the range of outcomes that are possible.

A Dynamic Model of Safety

At the heart of most organizations' attempts to prevent accidents and build a culture of safety are rules and procedures that individuals and organizations are expected to follow. Adherence to rules and procedures is an important component of safety climate; for example, Zohar uses a questionnaire to measure safety climate that includes several questions explicitly related to following rules (Zohar, 2000, pg. 591). High hazard organizations are especially compliance focused, with the need to satisfy regulators and prevent accidents - leading to detailed analysis and structuring of work (Carroll, Rudolph & Hatakenaka, 2002). As a result, a central construct in our model is 'adherence to rules and procedures', and this is assumed to have a positive effect on safety behavior, which in turn has a negative influence on the incident rate. It should be noted that we intentionally use the term "incident rate" rather than "accident rate"; this broader definition includes any abnormal event that could have the potential to lead to an accident, and that could also serve as the basis for either management action or learning (Cooke & Rohleder, 2006).

The evidence is also clear that adherence to rules and procedures alone will not prevent incidents. Certainly, Weick's notion of reliability is far broader; for example, reliability must also include the ability to handle the "unanticipated effects of emotional, social, and interpretive processes (1987, pg. 114)," the strength of communication norms, and the ability to learn and use learning to guide new routines (Carroll, Rudolph & Hatakenaka, 2002). Similarly, in the safety climate literature Katz-Navon, Naveh & Stern (2005) argue that the relationship between rules and procedures is curvi-linear, due to the costs of added complexity that come with too many rules. For all of these reasons, we introduce a second causal driver of safety behavior, termed the "effectiveness of rules and procedures." Broadly, this variable represents the outcome of learning and the general effectiveness of organizational safety routines and procedures, assuming full "adherence." Thus, even if adherence is 100%, this formulation allows for the possibility that the rules themselves are flawed, by way of a low value for "effectiveness of rules and procedures." Figure 1 shows the base causal structure of incidents that is used in the model.

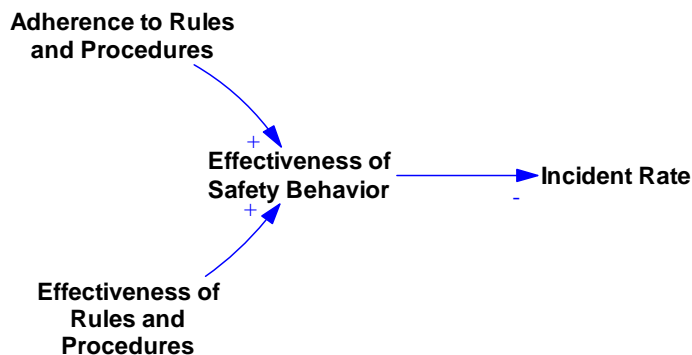


Figure 1: Basic causal structure for the Incident Rate

Individual Risk and Self Preservation

System Dynamics models are based on the concept of feedback. Negative feedback loops are “homeostatic processes” that bring system behavior towards a desired state, while positive feedback loops “amplify deviations” and cause exponential growth (Sterman, 2000). Figure 2 shows one simple balancing feedback loop regarding organizational safety, labeled ‘self preservation.’ The logic of this loop is as follows: an increase in the incident rate causes an increase in perceived personal threat, causing a greater adherence to rules and procedures, more effective safety behavior, and a *lower* incident rate. Thus, the incident rate is partly controlled through this balancing process of individuals recognizing the risk and adjusting their behavior.




Figure 2: The ‘Self-Preservation’ Balancing Loop

In the extreme case, the evidence for this ‘self-preservation’ loop is intuitive: individuals will act to preserve their own safety. As a mechanism for maintaining an incident rate that is acceptable to organizations, however, the evidence suggests that self-preservation is not effective. Zohar and Erev (2007) argue that a number of individual biases work to discourage safe behavior during routine work, including melioration bias (delayed outcomes), rare or uncertain outcomes, and social externalities. As a result, safety behavior is largely a function of management pressures and rewards.

Still, there is some evidence for a weak effect of self preservation. If individuals hold previous encounters or training related to specific threats, they will be more likely to participate in safety programs (Goldberg, et. al, 1991). In addition, employees will practice safe behaviors if they believe they have control over safe outcomes (Huang, et. al, 2004). As a result, we include ‘self preservation’ as a weak balancing loop that is most effective at high incident rates.

Management Safety Actions

Much of the safety literature acknowledges a large role for management actions in preventing accidents. Dimensions include perceived management attitudes to safety (Zohar, 1980), management concern (Brown & Holmes, 1986), management commitment (Dedobbeleer & Beland, 1991), management actions (Cox, et. al, 1998), and senior management support (Gershon, 2000). To take one example, Huang et. al (2004) define supervisor support as “the extent to which supervisors encourage safe working practices among their subordinates (pg. 485),” and construct a measure using two questions, one directed at positive reinforcement (“My supervisor acknowledges when I work safety”) and another at negative reinforcement (“My supervisor tells me when I am not working safely”). Research has also examined the impact of leadership style on safety, with evidence favoring “transformational” leadership over “corrective”

leadership (Zohar, 2002; Barling, 2002). In most cases, leadership and management actions have been found to be highly significant predictors of safety climate and the accident rate (eg. Barling, 2002). Zohar (2007) goes as far as to argue that the “key to success” in ensuring safe behavior “lies in providing frequent, personally meaningful, and immediate rewards for safe conduct, overriding the costs associated with that behavior and exceeding the benefits of unsafe behavior (pg. 122).”

To incorporate management action into a dynamic model of safety, we distinguish between two balancing feedback mechanisms, designed to correspond again with Zohar’s (2000) two factor measure of safety climate. The first of these is what Zohar terms “supervisory action,” and refers to “overt supervisor reaction to subordinate conduct (pg. 591).” As above, this factor includes both positive and negative reinforcement. Translated into feedback terms, Figure 3 shows an additional balancing loop labeled ‘Management Safety Actions’: when the incident rate rises above acceptable levels, managers respond with direct action so as to control safety behavior and keep the incident rate within an acceptable range.

The balancing loop ‘Management Safety Action’ is also consistent with theories of organizational learning. Specifically, this loop represents a prominent example of ‘single loop’ learning (Argyris, Putnam & Smith, 1985) or of ‘first order problem solving’ (Tucker & Edmondson, 2003). When faced with an incident rate that is too high, the natural and most immediately effective response for managers is to focus the blame on individual compliance with rules. Carroll, Rudolph & Hatakenaka (2002) provide an excellent example of this type of response: following an incident in which a maintenance worker falls off of a ladder, the accident report found that “the cause of the accident was a failure of the employee, the employee in charge, and the supervisor to properly follow the Accident Prevention Manual requirements for working in elevated positions.” Next steps included “appointing a full time safety person” so as to better “communicate company expectations (pg. 19).” This example suggests that enforcing adherence to rules and procedures by overt means is a method often used by managers to develop a positive safety climate.

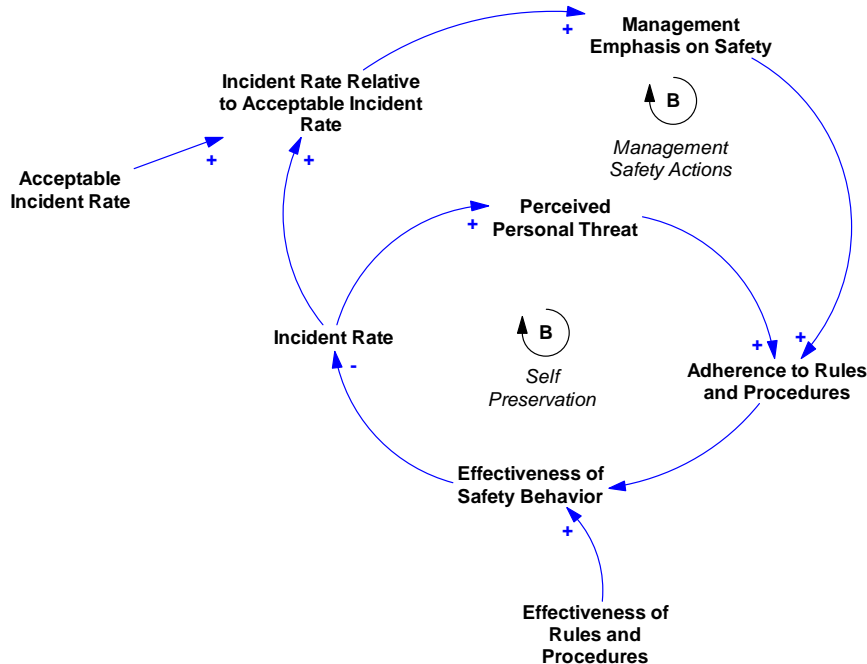


Figure 3: The ‘Management Safety Actions’ Balancing Loop

The second factor in Zohar's (2000) measure of safety climate is "expectation" regarding supervisory action. This is closely linked to the "safety vs. productivity" tradeoff (pg. 591), and is measured via survey items such as "whenever pressure builds up, my supervisor wants us to work faster, rather than by the rules" and "as long as there is no accident, my supervisor doesn't care how the work is done." In contrast to management safety *actions*, expectations are expected to convey a more general sense of the "overriding priority" of safety within an organization (pg. 595). Thus, "priority" has come to symbolize the relative importance of safety tasks in relation to other demands that workers face, with a high priority indicating that safety tasks will be given precedence even when work pressures are high (Katz-Navon, et al, 2002). Numerous studies have documented the importance of performance pressure as a predictor of safety behavior (Wright, 1986; Embrey, 2002; Katz-Navon, et. al, 2002; Clarke, 2006).

Two important feedback loops are introduced by this discussion (Figure 4). The first, labeled 'Cutting Back on Safety' represents the natural response that employees have to work pressure when safety is not a priority. If work requirements increase, production pressure will also increase, causing adherence to safety rules and procedures to fall until work demands are in line with what individuals are capable of completing. Evidence for the existence of this loop is strong, and includes both quantitative studies that show a relationship between production pressure and safety, and qualitative studies that document how individuals may cut corners when under time pressure (Wright, 1986; Oliva & Sterman, 2001).

The second feedback loop introduced in Figure 4, labeled 'Safety Priority,' represents the response that managers might take when expectations regarding safety are high. Here, work requirements are no longer exogenous, but instead reflect the safety needs of the organization. Thus, in a high priority environment, work requirements must necessarily fall when the safety needs of the organization rise, so as to contain the incident rate within acceptable levels. The positive correlation between safety priority, defined as the adherence to procedures even when under pressure, and safety climate (Zohar, 2000) suggests that in fact this loop is active in high performing organizations.

As with the 'Management Safety Actions' loop, 'Safety Priority' is also consistent with research in organizational learning. If 'Management Safety Actions' are examples of 'first order' problem solving or single loop learning, safety priority may represent one additional level of learning in response to incidents. (Although it is important to note that this is not yet 'double loop' learning in the sense defined above). For example, in the hospital nurse example cited earlier (Tucker & Edmondson, 2003), nurses often failed to report incidents to superiors in part due to a culture of independence and efficiency that had developed. The notion of safety priority is directly opposed to this concern for efficiency: if safety priority were high, nurse task requirements would instead reflect the need to follow through on problems that emerge, even if doing so required taking time away from regular work. Similarly, Edmondson (2002) argues that "reflective sessions where task and time pressure are temporarily removed" are important to psychological safety and the process of team learning. Thus, we hypothesize that the 'Safety Priority' loop is one important mechanism that organizations might use to prevent accidents and develop a culture of safety.

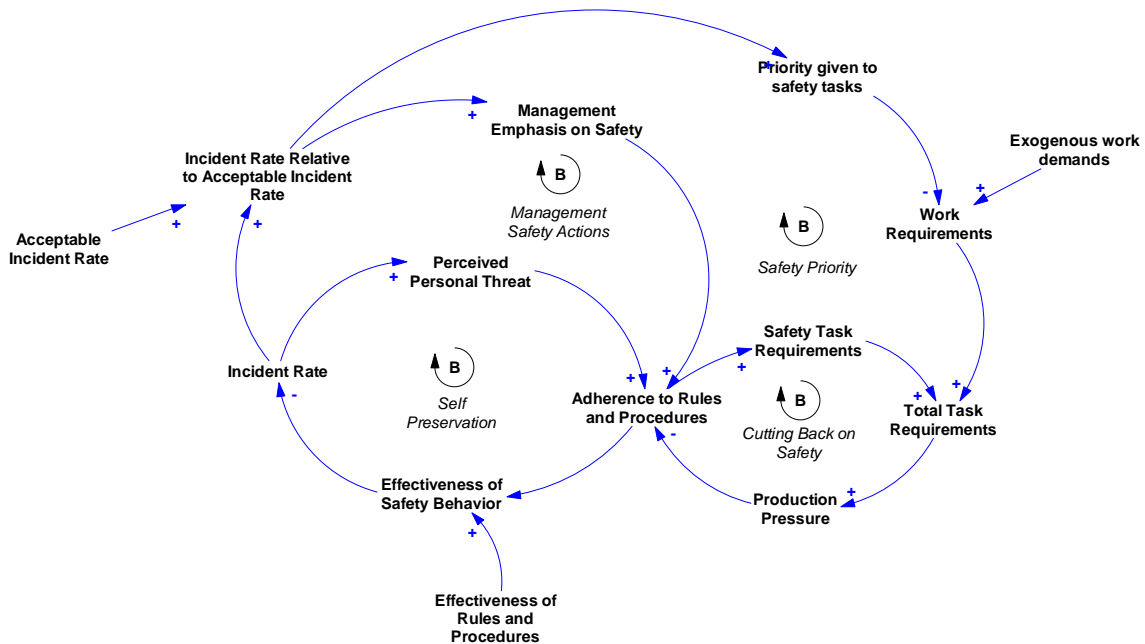


Figure 4: The 'Safety Priority' and 'Cutting Back on Safety' Balancing Loops

Organizational Learning

Of course, team learning and organizational learning involve far more than reducing time pressures so that existing rules and procedures can be followed. In addition, a large part of learning must involve generating *new* knowledge and routines in response to new challenges that will inevitably arise. Again, Weick's (1987) notion of reliability rests on rich communication between individuals, so as to generate "requisite variety" that enhances each individual's ability to perceive and act on information cues. In addition, Carroll, Rudolph & Hatakenaka (2002) argue that a focus solely on compliance is overly controlling, and as such may alienate workers.

As a result, reliable organizations must have additional feedback mechanisms in place that extend beyond simple adherence to rules and procedures. We suggest above that the term 'effectiveness of rules and procedures,' defined in the broadest sense possible, may capture some of the results of learning in high reliability organizations. To use Weick's model, we assume that in an organization that possesses a strong safety culture, incidents lead to the generation of a shared stock of knowledge that can aid in the reliable performance of daily tasks, thus making rules and procedures more "effective" and leading to fewer incidents. This specification also depends upon the broad definition for "incidents" mentioned above and an even broader definition for learning, such that "learning from incidents" goes beyond simple trial and error. Admittedly, this loop lacks a great deal of the richness characteristic of most theories of organizational learning, but nevertheless we believe it to be a good starting point for the investigation of model behavior.

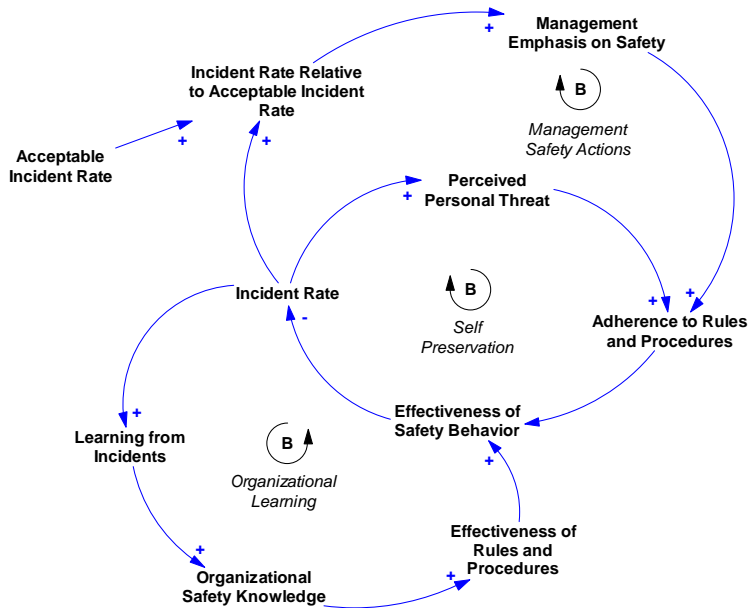


Figure 5: The 'Organizational Learning' Balancing Loop

Model Specification

We next incorporate each of the feedback loops described above into a mathematical simulation model of a generic organization, using the System Dynamics methodology (Sterman, 2000). Figure 6 shows the main stock and flow structure of task completion; detailed equation listings are provided in the appendix. Here, tasks are defined in a deliberately broad sense, to mean any work activities that must be completed. In System Dynamics models, 'stocks' are shown as rectangles and represent the memory of the system, flows represent the rate of change of stocks, and auxiliary variables are used to compute flows for each step that the simulation is run. In this example, new tasks flow into the "Task Backlog" stock, and flow out according to task completion. Task completion is in turn a function of capacity, capacity utilization, and the time required per task.

This structure is linked to the feedback loops shown above in two ways. First, adherence to rules and procedures determines the time required to complete tasks and thus the desired work rate (measured in hours per week). Second, the desired work rate (which is a function of backlog, desired completion time, and time required per task) together with capacity determine schedule or production pressure, one of the key drivers of adherence to rules and procedures. In sum, this simple model allows us to examine organizational responses to accidents in the context of the physical structure of task completion, an endeavor that is much more difficult using only qualitative theories. The generic structure used is very similar to that used in other System Dynamics models of organizations (see for example Sterman, 2000; Sterman, Henderson et. al, 2007; Paich & Sterman, 1993).

Two additional feedback loops are introduced by the task completion structure. The first is straightforward: as task backlog increases, schedule pressure increases causing capacity utilization to rise accordingly. The second loop, labeled 'hire more people,' however, deserves some mention. Operationally, this loop represents any hiring, firing, or transferring of staff such that over time, capacity matches desired capacity. Adherence to rules and procedures determines the amount of work that must be completed, which determines the desired staff level and - after a delay - the actual capacity of the organization. Given that in many organizations, this adjustment can in effect be quite slow, due to the strains involved with training new employees or the

pressures against layoffs, we choose a time constant for this adjustment that is reasonably long (6 months).

Beyond explicit hiring and firing, though, there is a broader meaning to this feedback process. Specifically, this loop represents the chief means by which safety practices become institutionalized into the memory of the system. Whether through hiring, firing, or simply the way that time is used, over time the capacity of the organization to complete tasks will adjust to established norms regarding how long tasks take. Crucially, the amount of time that tasks take necessarily reflects the amount of effort that is put towards safety. This result corresponds to the observation that “unsafe behaviors and routines can become ‘normal’ or habitual, in the sense that everyone does them (Hofmann & Stetzer, 1996, pg. 310, citing Wright, 1986).” A very similar result is established empirically by Oliva & Sterman (2001), who develop a System Dynamics model of backlog and order fulfillment in a call center. Over time, service quality (measured as the average time per call) declines, resulting in both lower standards and lower capacity to handle future calls. In a sense, we hypothesize the existence of a similar dynamic for the case of organizational safety, with safety standards as analogous to Oliva & Sterman’s notion of service quality.

This feedback loop is also crucial to the model’s treatment of complacency, a concept that is central to many accounts of safety behavior in organizations (Minami & Madnick, 2007; Leveson, 2005). If the incident rate is low for an extended period of time and safety is not deliberately given a high priority by managers, adherence will gradually fall as individuals respond to production pressure and the demands of daily tasks. An organization becomes complacent, however, as these changes are institutionalized in desired staff and, after a delay, in capacity. As soon as capacity reflects new assumptions about time per task, production pressures reflect this new state, making re-adherence to safety more and more difficult. Thus, as it is modeled here, a falling incident rate does contribute to increased complacency over time in the absence of management attention to safety.




Figure 6: The Stock and Flow structure of Task Completion

The stock and flow structure for adherence to rules and procedures is shown in Figure 7. Adherence is also modeled as a stock, reflecting the fact that individual actions form habits that change only gradually, especially in the context of large, bureaucratic organizations. As the preceding analysis suggests, three forces influence adherence: personal threat, management emphasis, and production pressure. The sum of these threats determines the fractional change in adherence, and as long as this sum is positive, adherence will gradually increase until the pressure is relieved, according to a standard hill climbing heuristic (Sterman, 2000, pg. 537). Similarly, when the sum is negative, adherence will fall. For example, if production pressure is positive, in the absence of other pressures adherence will gradually decrease, freeing resources until pressure is at a more manageable level. The variables “Maximum Increase in Adherence” and “Minimum time to increase adherence” are included to ensure that adherence remains between 0 and 100%. A complete listing of model equations is provided in the appendix.

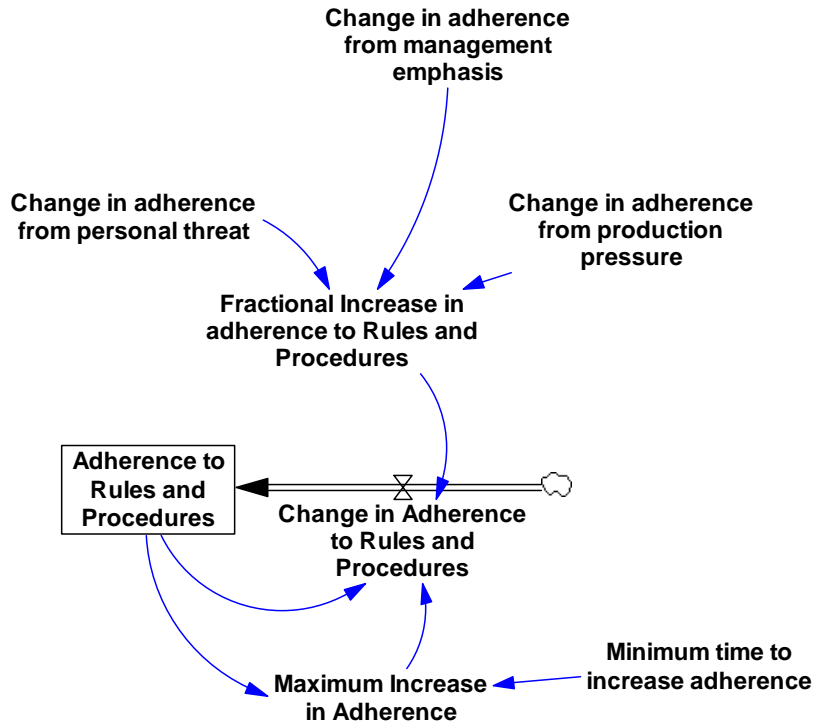


Figure 7: The Stock and Flow Structure of Adherence

Model Analysis and Results

Next, we present an analysis of select model results, designed to show the relative influence of each of the balancing mechanisms, detailed above, that organizations might use to prevent accidents and build a culture of safety. We start with the simplest mechanisms, self preservation and overt management action, and move through those mechanisms that we might expect to find in organizations with more sophisticated approaches towards safety.

1) Self-Preservation Alone

Figure 8 shows model results for a run in which only self preservation is active. Beginning in month five, there is a 5% step increase in the rate of new tasks, causing a rise in backlog and therefore a rise in production pressure. The rise in production pressure causes adherence to rules and procedures to fall to meet completion goals. In this example, the incident rate rises above the self preservation threshold, which is assumed to be four incidents per month. As a result, there is a slight rise in adherence as individuals react to the unsafe environment. However, self preservation alone is not sufficient to bring the incident rate back to within levels that are acceptable to the organization. Thus, this behavior is consistent with the expectation that self preservation has a weak but insufficient influence on accident prevention.




Figure 8: Model results for a step increase in new tasks. When self-preservation is the only active mechanism, the increased production pressure leads to a permanently higher equilibrium incident rate.

2) Adding Management Safety Actions

In Figure 9, the same step increase in the rate of new tasks is simulated, only now a second balancing loop is added in the form of overt management safety actions (as shown above in Figure 3). The red line represents the previous run with self-preservation alone, and the blue represents the behavior after management actions are added. From the graphs, the added influence of management pressure is clear. Now, the incident rate and adherence return to a point close to their starting values, as management pressure persists until this goal is achieved. The difference is also shown in the graph for capacity: because management pressure forces adherence to rise back to its original level, capacity must increase to meet the new, raised demand. In contrast, when self preservation alone is active, part of the increased demand is satisfied by cutting corners on safety tasks.




Figure 9: Model Results for a step increase in new tasks. When management safety actions are also taken, the incident rate returns to equilibrium, but an increase in capacity is required to complete additional tasks.

The results shown in Figure 9 also illustrate the notion of bounded rationality on the part of managers (Cyert & March, 1963). The theory of bounded rationality states that in complex organizations, rather than solve for the globally optimal response, individuals with limited mental models will instead search for satisfactory outcomes that are consistent with only a small subset of feedbacks, usually those for which time delays are small (Cyert & March, 1963; Sterman, 2000b). The ubiquity of ‘single loop learning’ focused on short term, overt reactions to accidents has been discussed; Figure 9 shows precisely why such a response is intendedly rational on the part of managers. In a simplified system in which demand exhibits only a single step increase and only feedback through self preservation and management action are active, an overt response is entirely effective in bringing the incident rate back to an acceptable level. Thus, the partial model test supports the actions that managers sometimes take along these lines (see Morecroft, 1985 for a full discussion of bounded rationality and partial model testing of System Dynamics models).

Is a policy of overt management safety action similarly effective in a more complex system? Figure 10 begins to explain why the answer to this question is ‘No.’ Here, instead of a step increase in demand, beginning in month 5 demand follows a pattern of stochastic noise, with a standard deviation of 5% and a mean that remains constant at the original value. As the blue line indicates, at first the incident rate remains roughly equivalent to its base value, with only small fluctuations. At around time 20, however, a sudden random increase in backlog causes the system to begin a slow deterioration. Adherence falls and the incident rate rises, as the lower level of adherence becomes institutionalized in the form of lower capacity. The balancing loop ‘Hire more people’ (Figure 6) is crucial to understanding this behavior: during each period that backlog rises, norms concerning safety gradually slip,

resulting in a slow deterioration in safety standards over time. Management action can partly counteract this trend, as evidenced by the sharp periods of increase in adherence, but eroding goals cause production pressure to remain high, making a full recovery almost impossible. While a step increase in the rate of new tasks can send a clear signal that more capacity is needed, here, each small increase in the rate of new tasks causes further erosion in the incident rate before the organization knows to respond. The performance is worse in the “noise” run despite the fact that the mean arrival rate of new tasks remains constant, unlike the “step” run where the mean increases. These results show that a policy of relying solely on overt management pressure to ensure safety, although intendedly rational in a simple system, is not robust even for a small change in the form of the demand.

Notably, this analysis neglects several other added complexities that *further* undermine strategies focused solely on ‘single loop learning,’ several of which are discussed above. Examples include the gradual erosion of trust due to an overemphasis on compliance (Carroll, Rudolph, & Hatakenaka, 2002), the failure to question underlying causes of accidents, or the failure to develop a culture of reliability, to use Weick’s sense of the term. For all of these reasons, a strategy of compliance only, while effective for a simple system, does not withstand the test of a more complex and realistic system.




Figure 10: Model results for a comparison between stochastic noise in the rate of new tasks and a step increase in the rate of new tasks. When demand is noisy, a series of small increases in production pressure cause adherence to safety procedures to gradually erode. In contrast, a step increase eventually sends a clear signal that more resources are needed.

3) Making Safety a High Priority

The safety climate literature provides evidence that in addition to overt action, managers may also improve safety by encouraging individuals to complete safety tasks regardless of other work pressures that may exist. We next simulate this ‘high safety priority’ policy for the same pattern of demand introduced above (stochastic noise with a standard deviation of 5%), to investigate whether a similar erosion in safety occurs. Now, starting at time 5 the

desired capacity shown in Figure 6 is independent of adherence to rules and procedures, and is instead calculating by assuming that adherence will be 100%.

As expected, safety priority has a large influence on the incident rate. By allocating resources under the assumption of full adherence, capacity necessarily grows, but the result is lower production pressure, greater adherence, and a lower incident rate. Notably the volatility in incidents is also far less under this policy: during periods that backlog and schedule pressure rise, erosion is contained by continuing to maintain the same expectation regarding the time required per task. Interestingly, the decreased volatility may be consistent with findings that safety priority *moderates* relationships between other safety variables (eg., Katz-Navon, et al, 2005). For example, higher safety priority and lower volatility implies that over time, the incident rate is less responsive to changes in work pressure and to overt management action. Although this is not the equivalent of comparing across organizations, it is interesting that cross sectional studies do find, for example, a weaker influence of management safety practices when safety priority is high (Katz-Navon, et al, 2005).

Finally, it is important to note that a policy of high safety priority comes at a cost of far lower average productivity, where productivity is defined as tasks completed per month per person. By definition, safety tasks are 'extra' work that require time but that are not necessary to the production of final output. Thus, completing more safety tasks increases the total time per task and lowers productivity. This analysis omits additional feedback that might enhance productivity in the long run: for example, in some cases employees might eventually become more productive (in a broad sense) in an environment that they know is safe.

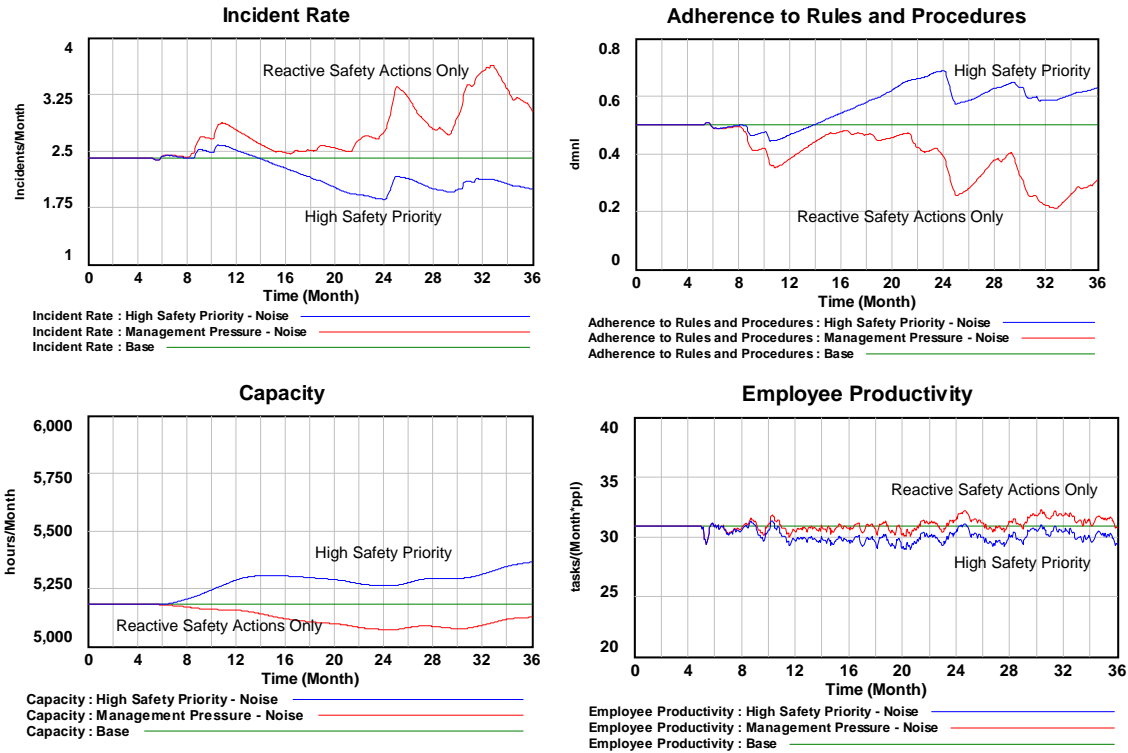


Figure 11: Comparison between a reactive approach to safety (management safety actions), and an approach where safety is always given a high priority. A high priority means increasing capacity in response to task demand, such that safety tasks are still completed. A higher adherence to safety tasks necessarily means a lower productivity (tasks/person/month) on regular tasks.

4) Adding organizational learning

Finally, learning is added to the model. First, we activate feedback through learning without making safety a high priority, to investigate whether norms of communication and learning alone can prevent erosion in safety. Figure 12 shows that although learning increases the effectiveness of rules and procedures and reduces the incident rate, there is still significant erosion in safety over time. Somewhat counter intuitively, increases in effectiveness are accompanied by *lower* adherence, meaning that the overall decrease in the incident rate is far less than it might be. The reason for lower adherence is that higher effectiveness weakens the perceived need for an organizational response through other means, such as overt management emphasis. In addition, without making safety a high priority work pressures remain, causing the same drop in adherence during periods of high backlog. In sum, learning has some benefits, but as long as safety remains a low priority relative to task completion, volatility and erosion in safety standards persist.

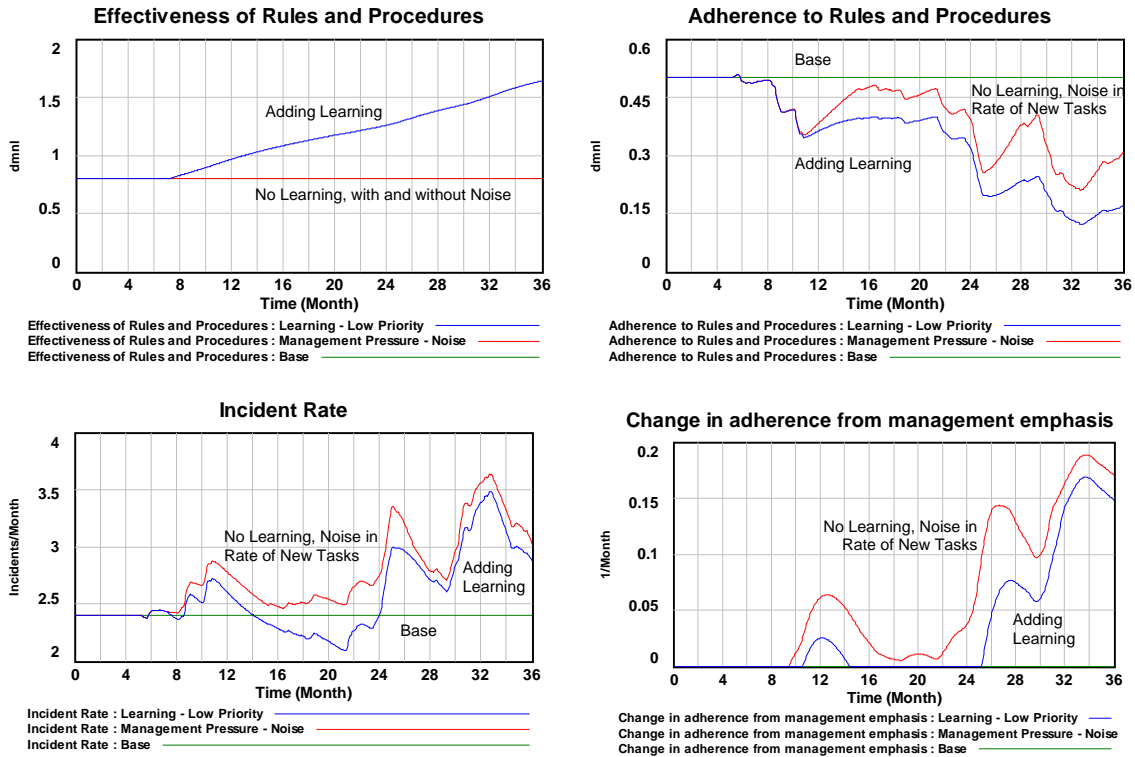


Figure 12: Adding organizational learning, under an assumption of a low safety priority. In this case, learning leads to more effective rules and procedures but only a slight decrease in the incident rate, due to declining adherence.

What if safety is again given a high priority? As expected, Figure 13 shows that under this set of assumptions, safety performance exceeds all other scenarios shown above. The incident rate falls to its lowest level, and also loses the volatility characteristic of runs for which safety priority is low. Both adherence to rules and procedures and the effectiveness of rules and procedures increase. Thus, results support the hypothesis that learning behaviors can contribute to safety most effectively when safety is given a high priority by managers.

Figure 13 also contains a couple of somewhat counter intuitive results. First, the effectiveness of rules and procedures actually increases *more* when safety is given a low priority, due to the fact that more learning is possible when the number of incidents is higher. This is precisely the dilemma faced by high hazard organizations: given that reliability is a “dynamic non-event” (Weick, 1987, pg. 118), individuals must learn from small events, understand “that inertia is a complex state,” and guard against complacency. Here, we see that the most reliable organization, that which exhibits both learning and a high safety priority, avoids complacency by maintaining a high level of adherence, and also learns efficiently from a small and decreasing incident rate. Although the overall stock of knowledge is smaller, the average learning per incident is higher and lessons are more consistently adhered to.




Figure 13: Comparing learning under assumptions of both low and high safety priority. When safety is given a high priority, the resulting improvement in safety is substantial. Adherence also rises and reactive management pressure is no longer necessary.

The situation is even more challenging when incidents are rare *despite* a safety culture that is weak, as was the case at NASA during the time leading up to the Challenger disaster. If incidents are rare and safety priority is low, production pressure will gradually crowd out safety adherence as above, a sign of the organization becoming complacent. In addition, the lack of incidents limits learning, meaning that a pattern of behavior similar to the low priority – low learning trajectory in figure 13 is possible. Thus, in organizations like NASA where incidents are rare, the importance of safety priority is paramount.

A second interesting result from Figure 13 concerns the role of overt management action in an organization for which both learning is active and the priority of safety is high. As the blue line in the bottom right graph of Figure 13 shows, overt management emphasis under this scenario is zero at all times. In other words, due to the positive effects of learning and the decreased influence of production pressure, the incident rate never reaches a point at which managers feel the need to respond reactively. Much has been said about the disadvantages of single loop learning and of a leadership style that overly emphasizes compliance. Should we emphasize eliminating these behaviors by choice, or will single loop learning disappear as a natural result of implementing more effective policies elsewhere?

Discussion and Conclusion

Figure 14 summarizes the results of the above analysis by showing four distinct outcomes in which two policies are varied, the strength of learning and the priority of safety. These outcomes are also presented in Table 1. (Again safety priority refers to the degree to which

individuals are encouraged to place safety above productivity.) A number of conclusions can be drawn.

First, it is clear from the resulting behavior that whatever assumptions are made about the effectiveness of rules and procedures or the strength of learning, making safety a high priority by limiting production pressure is by far the highest leverage policy available to managers seeking to prevent accidents. Both runs in which safety is a high priority reach and sustain lower incident rates than those for which safety is a low priority, regardless of the strength of learning. Furthermore, safety priority can also be a source of reliability: because production pressure is contained, gains in improving the incident rate are sustained and are not vulnerable during the inevitable times when work pressures become large. As a result, the incident rate remains at consistent levels rather than fluctuating in response to each new wave of management emphasis.

This result may provide insight into the nature of reliability in organizations that do hazardous work. While part of reliability certainly involves developing rich norms of communication so as to effectively respond to cues (Weick, 1987), reliability also includes the ability to prevent the type of oscillatory behavior that can result when management pressure or single loop learning is the dominant form of control. An oscillating incident rate is the hallmark of a reactive organization, where successive crises lead to short term fixes that persist only until the next crisis. Partly, eliminating this behavior must involve cultural issues that are unique to each context. However, there is also a general lesson: management decisions must involve giving individuals the time that they need to complete safety tasks without feeling pressure to abandon them.




Figure 14: Comparing the Incident rate for four different scenarios. Giving safety a high priority is by far the most effective policy, with learning providing a substantial incremental benefit only when safety priority is high.

| | Priority of Safety | |
|----------------------|--------------------|-----------------------------|
| | Low | High |
| Strength of Learning | Low | Reliable Moderate Performer |
| | High | Reliable High Performer |

Table 1: 2x2 for combining learning and safety priority

A second important lesson concerns the interaction between learning and safety priority. Figure 14 illustrates clearly that the gains from learning are far more significant when priority is high than when priority is low. In fact, by month 36 many of the gains from learning in the high learning – low priority organization have been entirely lost by gradual erosion in adherence to rules and procedures. In contrast, when safety priority is high, learning creates a large improvement. Figure 15 illustrates this point even more dramatically: even under assumptions of extremely effective learning (in blue), gains eventually erode almost entirely when the priority of safety is low. Thus, a key contribution is that an understanding of the physical structure of task completion and of work pressure is essential to supporting learning in an organization. If adherence to rules and procedures is allowed to erode, it makes little difference how good rules and procedures are in the first place.




Figure 15: The case of extremely effective learning. If priority is low, eventually safety will erode, as above

Interestingly, this observation may be loosely related to what Carroll, Rudolph and Hatakenaka (2002) term ‘structuring’ in their four stage model of organizational learning. To use

their words, “structuring is a collective achievement of consistency and predictability by integrating knowledge into routines and cultural meanings that are institutionalized throughout the organization.” In other words, this last stage of learning involves institutionalizing knowledge gained from learning so that it is used consistently, a concept that in many ways contains our notion of ‘adherence to rules and procedures.’ Indeed, we find that supporting learning by maintaining adherence is crucial, and that if lessons are not institutionalized and instead erode, learning is easily undermined. ‘Safety priority,’ as it is defined in the safety climate literature, may be one essential element of making these lessons endure.

As a final note, it is important to again consider these conclusions in the context of organizational learning. Organizations seeking to learn from accidents and develop a culture of safety may continuously refine a stock of knowledge that increases the effectiveness of rules, as this simulation represents. However, there is also a higher level of learning that transcends what can be captured within a single simulation model. In addition, a simulation may itself serve as a boundary object (Carlile, 2004) that allows managers to experience fundamental learning about the dynamics of their organization as a whole. In this case, managers learn not only about the specifics of individual incidents, but also about the importance of understanding the tradeoff between productivity and safety. The organizational response to accidents is complex, and by adapting the framework proposed in this study to specific organizations we hope to contribute also to this larger form of learning.

Acknowledgements

Work reported herein was supported, in part, by the Singapore Defence Science and Technology Agency (DSTA).

References

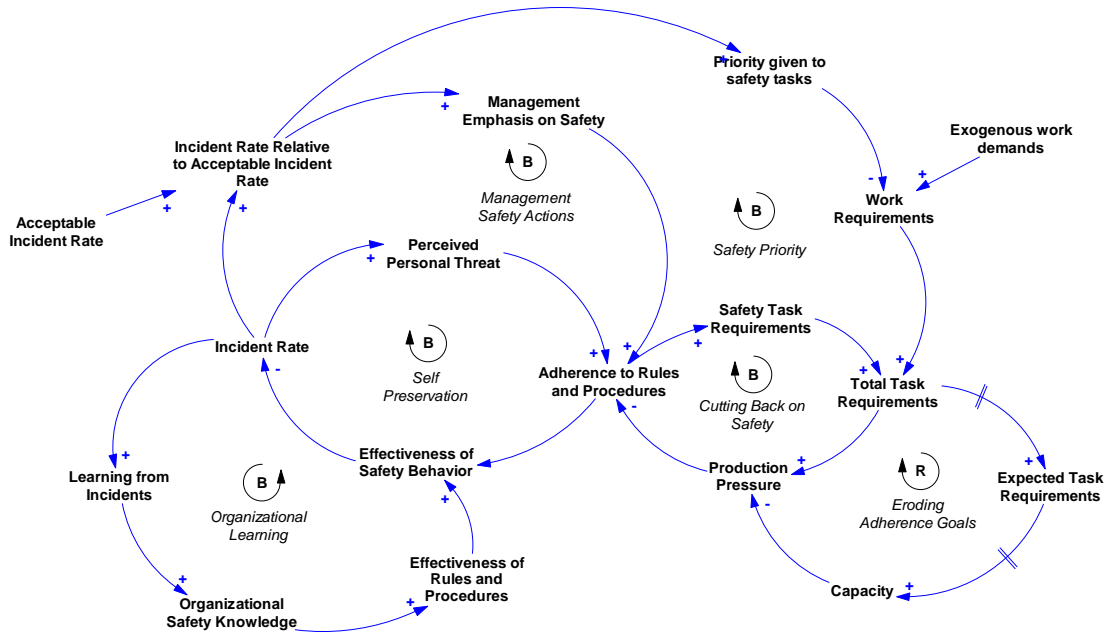
- Argyris, C., R. Putnam, et al. (1985). Action science:[concepts, methods, and skills for research and intervention], Jossey-Bass.
- Barling, J., E. K. Kelloway, et al. (2003). "High-quality work, job satisfaction, and occupational injuries." Journal of Applied Psychology **88**(2): 276–83.
- Barling, J., C. Loughlin, et al. (2002). "Development and test of a model linking safety-specific transformational leadership and occupational safety." J Appl Psychol **87**(3): 488-96.
- Borenstein, Seth. "Deadly Mining Method Used Often," The Associated Press, August 6, 2007
- Brown, R. L. and H. Holmes (1986). "The use of a factor-analytic procedure for assessing the validity of an employee safety climate model." Accid Anal Prev **18**(6): 455-70.
- Carlile, P. R. (2004). "Transferring, Translating, and Transforming: An Integrative Framework for Managing Knowledge Across Boundaries." Organization Science **15**(5): 555-568.
- Carroll, J. S. (1993). "Out of the Lab and Into the Field: Decision Making in Organizations." Social Psychology in Organizations: 38-62.
- Carroll, J. S. (1998). "Organizational learning activities in high-hazard industries: The logics underlying self-analysis." Journal of Management Studies **35**(6): 699.
- Carroll, J. S. (1998). "Safety Culture as an Ongoing Process: Culture Surveys as Opportunities for Enquiry and Change." Work and Stress **12**(3): 272-284.
- Carroll, J. S., J. W. Rudolph, et al. (2002). "Learning from experience in high-hazard organizations." Research in Organizational Behavior **24**: 87-137.
- Cheyne, A., A. Oliver, et al. (2002). "The architecture of employee attitudes to safety in the manufacturing sector A. Cheyne, ." Personnel Review **31**(6): 649-670.
- Clarke, S. (1999). "Perceptions of organizational safety: implications for the development of safety culture." Journal of Organizational Behavior **20**(2): 185-198.
- Clarke, S. (2000). "Safety culture: under-specified and overrated?" International Journal of Management Reviews **2**(1): 65-90.
- Clarke, S. (2006). "Safety climate in an automobile manufacturing plant: The effects of work environment, job communication and safety attitudes on accidents and unsafe behaviour." Personnel Review **35**(4): 413-430.
- Cooke, D. L. and T. R. Rohleder (2006). "Learning from incidents: from normal accidents to high reliability." System Dynamics Review **22**(3): 213-239.
- Cox, S., J. M. Tomas, et al. (1998). "Safety Culture: The Prediction of Commitment to Safety in the Manufacturing Industry." British Journal of Management **9**(s 1): 3-11.
- Cyert, R. M. and J. March (1963). A Behavioral Theory of the Firm.
- Dedobbeleer, N. and F. Beland (1991). "A safety climate measure for construction sites." Journal of Safety Research **22**(2): 97-103.
- Edmondson, A. C. (2002). Managing the Risk of Learning: Psychological Safety in Work Teams, Division of Research, Harvard Business School.
- Embrey, D. E. (1992). "Incorporating management and organisational factors into probabilistic safety assessment." Reliability Engineering & System Safety **38**(1): 199-208.
- Gershon, R. R., C. D. Karkashian, et al. (2000). "Hospital safety climate and its relationship with safe work practices and workplace exposure incidents." Am J Infect Control **28**(3): 211-21.
- Glendon, A. I. and N. A. Stanton (2000). "Perspectives on safety culture." Safety Science **34**(1): 193-214.
- Goldberg, A. I., E. M. Dar-El, et al. (1991). "Threat Perception and the Readiness to Participate in Safety Programs." Journal of Organizational Behavior **12**(2): 109-122.
- Griffin, M. A. and A. Neal (2000). "Perceptions of safety at work: a framework for linking safety climate to safety performance, knowledge, and motivation." J Occup Health Psychol **5**(3): 347-58.

- Guldenmund, F. W. (2000). "The nature of safety culture: a review of theory and research." Safety Science **34**(1): 215-257.
- Hofmann, D. A. and F. P. Morgeson (1999). "Safety-related behavior as a social exchange: The role of perceived organizational support and leader-member exchange." Journal of applied psychology **84**(2): 286-296.
- Hofmann, D. A. and A. Stetzer (1996). "A Cross-Level Investigation of Factors Influencing Unsafe Behaviors and Accidents." Personnel Psychology **49**(2).
- Hofmann, D. A. and A. Stetzer (1998). "The Role of Safety Climate and Communication in Accident Interpretation: Implications for Learning from Negative Events." The Academy of Management Journal **41**(6): 644-657.
- Huang, Y. H., P. Y. Chen, et al. (2004). "Quality of the Execution of Corporate Safety Policies and Employee Safety Outcomes: Assessing the Moderating Role of Supervisor Safety Support and the Mediating Role of Employee Safety Control." Journal of Business and Psychology **18**(4): 483-506.
- Humphrey, S. E., H. Moon, et al. (2004). "Decision making and behavior fluidity: How focus on completion and emphasis on safety changes over the course of projects." Organizational Behavior and Human Decision Processes **93**(1).
- Isla Diaz, R. and D. Diaz Cabrera (1997). "Safety climate and attitude as evaluation measures of organizational safety." Accid Anal Prev **29**(5): 643-50.
- Janssens, M., J. M. Brett, et al. (1995). "Confirmatory Cross-Cultural Research: Testing the Viability of a Corporation-Wide Safety Policy." The Academy of Management Journal **38**(2): 364-382.
- Katz-Navon, T. A. L., E. Naveh, et al. (2005). "Safety Climate in Health Care Organizations: A multidimensional Approach." The Academy of Management Journal **48**(6): 1075-1089.
- Lee, T. (1998). "Assessment of safety culture at a nuclear reprocessing plant." Work and stress **12**(3): 217-237.
- Leveson, N. G., B. Barrett, et al. (2005). Modeling, Analyzing and Engineering NASA's Safety Culture: Phase 1 Final Report. Cambridge, MA, Massachusetts Institute of Technology.
- M Goldenhar, L., L. J. Williams, et al. (2003). "Modelling relationships between job stressors and injury and near-miss outcomes for construction labourers." Work & Stress **17**(3): 218-240.
- Marais, K., J. H. Saleh, et al. (2006). "Archetypes for organizational safety." Safety science **44**(7): 565-582.
- Mearns, K., R. Flin, et al. (1998). "Measuring safety climate on offshore installations." Work and stress **12**(3): 238-254.
- Minami, N and S. Madnick (2007). "Reducing Combat Vehicle Accidents via Improved Organizational Processes." MIT Working Paper, available at <http://web.mit.edu/smadnick/www/wp/2007-02.pdf>
- Morecroft, J. D. W. (1985). "Rationality in the Analysis of Behavioral Simulation Models." Management Science **31**(7): 900-916.
- Naveh, E., T. Katz-Navon, et al. (2005). "Treatment Errors in Healthcare: A Safety Climate Approach." Management Science **51**(6): 948-960.
- Neal, A. and M. A. Griffin (2002). "Safety Climate and Safety Behaviour." Australian Journal of Management **27**(2).
- Nelson, Jennifer and B Hankinson. "High End Fare, with Side order of Violations," The Boston Globe, August 26, 2007
- Neuman, William (2007). "Safety Violations Cited for Deaths in Subway Work." The New York Times, August 2, 2007.
- Oliva, R. and J. D. Sterman (2001). "Cutting Corners and Working Overtime: Quality Erosion in the Service Industry." Management Science **47**(7): 894-914.
- Paich, M. and J. D. Sterman (1993). "Boom, Bust, and Failures to Learn in Experimental Markets." Management Science **39**(12): 1439-1458.

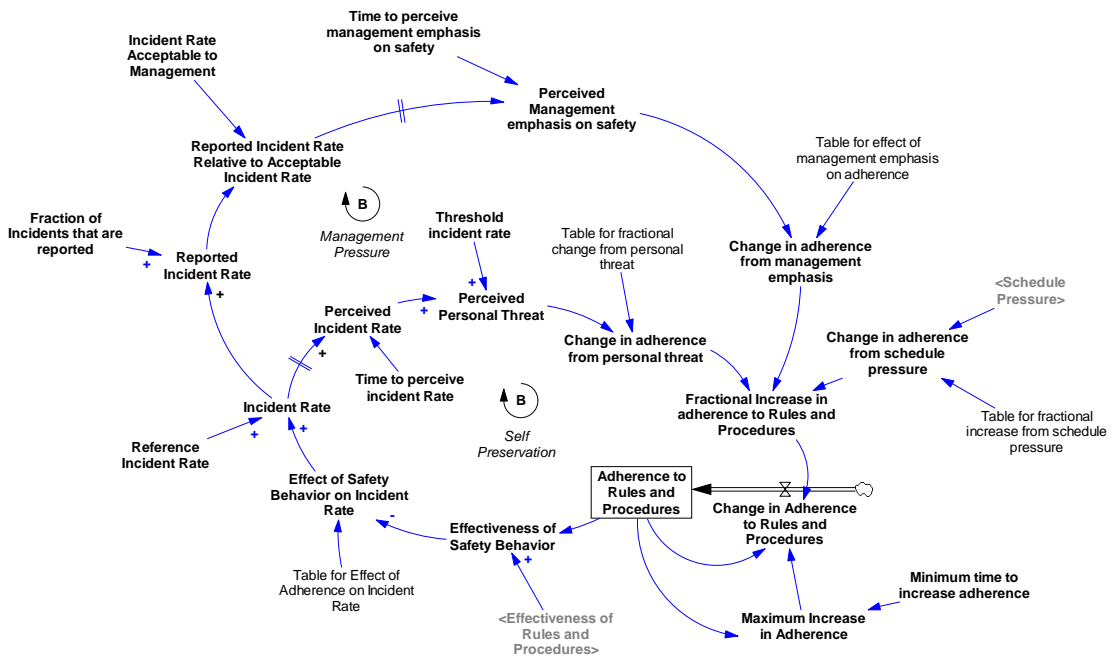
- Perrow, C. (1984). *Normal accidents: Living with high-risk technologies*, Basic Books, New York, NY.
- Pidgeon, N. (1997). "The limits to safety? Culture, politics, learning and man-made disasters." Journal of Contingencies and Crisis Management **5**(1): 1-14.
- Pidgeon, N. F. (1991). "Safety Culture and Risk Management in Organizations." Journal of Cross-Cultural Psychology **22**(1): 129.
- Reason, J. (1998). "Achieving a safe culture: theory and practice." Work and Stress **12**(3): 293-306.
- Repenning, N. (1999). "A Simulation-Based Approach to Understanding the Dynamics of Innovation Implementation."
- Repenning, N. P. and J. D. Sterman (2002). "Capability Traps and Self-Confirming Attribution Errors in the Dynamics of Process Improvement." Administrative Science Quarterly **47**(2): 265-295.
- Roberts, K. H. (1990). "Some Characteristics of One Type of High Reliability Organization." Organization Science **1**(2): 160-176.
- Roberts, K. H., S. K. Stout, et al. (1994). "Decision Dynamics in Two High Reliability Military Organizations." Management Science **40**(5): 614-624.
- Rudolph, J. W. and N. P. Repenning (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.
- Schein, E. H. (1985). Organizational culture and leadership. San Francisco, Jossey-Bass Publishers.
- Senge, P. M. (1992). The Fifth Discipline: The Art and Practice of the Learning Organization, Century Business.
- Sterman, J. D. (2000). Business Dynamics: Systems Thinking and Modeling for a Complex World, Irwin/McGraw-Hill.
- Sterman, J. D. (2000). "Learning In and About Complex Systems." Reflections: The SoL Journal **1**(3): 24-51.
- Sterman, J. D., R. Henderson, et al. (2007). "Getting Big Too Fast: Strategic Dynamics with Increasing Returns and Bounded Rationality." Management Science **53**(4): 683.
- Tucker, A. L. (2003). "Why hospitals don't learn from failures: Organizational and psychological dynamics." California Management Review **45**(2): 55-72.
- Tucker, A. L., I. M. Nembhard, et al. (2007). "Implementing New Practices: An Empirical Study of Organizational Learning in Hospital Intensive Care Units." Management Science **53**(6): 894.
- Weick, K. E. (1987). "Organizational culture as a source of high reliability." California Management Review **29**(2): 112-127.
- Wright, C. (1986). "Routine deaths: Fatal accidents in the oil industry." Sociological Review **34**(2): 265-288.
- Zohar, D. (1980). "Safety climate in industrial organizations: theoretical and applied implications." J Appl Psychol **65**(1): 96-102.
- Zohar, D. (2000). "A group-level model of safety climate: testing the effect of group climate on microaccidents in manufacturing jobs." J Appl Psychol **85**(4): 587-96.
- Zohar, D. (2002). "The effects of leadership dimensions, safety climate, and assigned priorities on minor injuries in work groups." Journal of Organizational Behavior **23**(1): 75-92.
- Zohar, D. and I. Erev (2007). "On the difficulty of promoting workers' safety behaviour: overcoming the underweighting of routine risks." International Journal of Risk Assessment and Management **7**(2): 122-136.

Appendix I: Model Diagrams


Full Causal Loop Diagram



Model Diagram: Incident and Adherence Sector



Model Diagram: Task Completion Sector



Appendix II: Complete Listing of Model Equations

Note: Model is written using the Vensim modeling software, available from www.vensim.com

| | |
|--|---|
| <p>Number of Rules and Procedures= 5.76 ~ Procedures/task ~</p> | <p>Time for learning to become obsolete= 48 ~ months ~</p> |
| <p>Decrease in Safety Knowledge= Safety Knowledge/Time for learning to become obsolete ~ 1/Month ~</p> | <p>Relative Incident Rate= Incident Rate/Reference Incident Rate for Learning ~ dmdl ~</p> |
| <p>Safety Requirements per Task= Number of Rules and Procedures*Adherence to Rules and Procedures*Time required per procedure ~ hours/task ~</p> | <p>Fraction of Incidents that are reported= 1 ~ dmdl ~</p> |
| <p>Reference Safety Knowledge= INITIAL(Safety Knowledge) ~ dmdl ~</p> | <p>Safety Knowledge= INTEG (Increase in Safety Knowledge-Decrease in Safety Knowledge, Initial Effectiveness of Learning* Time for learning to become obsolete) ~ dmdl ~</p> |
| <p>Table for Effect of Number of Incidents on Learning([(0,0)- (6,2)],(0,0.8),(1,1),(2,1.1),(3,1.2),(4,1.3),(5,1.4)) ~ dmdl ~</p> | <p>Table for Effect of Safety Knowledge on Effectiveness of Rules and Procedures([(0,0)- (3,1)],(0,0.05),(0.321101,0.192982),(0.678899,0. 359649),(1,0.5),(1.54128,0.714912\) ,(2.11927,0.877193),(2.54128,0.95614) ,(3,1)) ~ dmdl ~</p> |
| <p>Increase in Safety Knowledge= Effectiveness of Current Learning ~ 1/Month ~</p> | <p>Effectiveness of Rules and Procedures= Table for Effect of Safety Knowledge on Effectiveness of Rules and Procedures(Relative Safety Knowledge\) ~ dmdl ~</p> |
| <p>Relative Safety Knowledge= Safety Knowledge/Reference Safety Knowledge ~ dmdl ~</p> | <p>Effect of Number of Incidents on Learning= Table for Effect of Number of Incidents on Learning(Relative Incident Rate) ~ dmdl ~</p> |
| <p>Reference Incident Rate for Learning= INITIAL(Incident Rate) ~ Incidents/Month ~</p> | <p>Effectiveness of Safety Behavior= ~</p> |
| <p>Effectiveness of Current Learning= Strength of Learning*Effect of Number of Incidents on Learning ~ 1/Month ~</p> | |

Adherence to Rules and Procedures*Effectiveness of Rules and Procedures
 ~ dmn1
 ~ |

Initial Effectiveness of Learning=
 0.5
 ~ 1/Month
 ~ |

Planned Safety Requirements per task=
 Number of Rules and Procedures*Time
 required per procedure
 ~ hours/task
 ~ |

Strength of Learning=
 0.5
 ~ 1/Month
 ~ |

Desired Staff Level=
 (1-Switch for proactive
 management)*Total Desired Work Rate/Normal
 work amount+
 Switch for proactive
 management*Planned Desired Staff Level
 ~ ppl
 ~ |

Employee Productivity=
 Task completion/Staff Level
 ~ tasks/Month/ppl
 ~ |

Switch for proactive management=
 0
 ~ dmn1 [0,1,1]
 ~ |

Planned Desired Staff Level=
 (Desired Task completion rate*(Time
 per task+Planned Safety Requirements per
 task)/Normal work amount\
)
 ~ ppl
 ~ |

Perceived Management emphasis on safety=
 smooth(Reported Incident Rate Relative
 to Acceptable Incident Rate,Time to perceive
 management emphasis on safety\
)
 ~ dmn1
 ~ |

Fractional Increase in adherence to Rules and Procedures=
 Change in adherence from personal
 threat+Change in adherence from schedule
 pressure+\
 Change in adherence from
 management emphasis
 ~ 1/Month
 ~ |

Base Rate of new Tasks=
 1000
 ~ tasks/Month
 ~ |

Incident Rate Acceptable to Management=
 4
 ~ Incidents/Month
 ~ |

Reported Incident Rate Relative to Acceptable
 Incident Rate=
 Reported Incident Rate/Incident Rate
 Acceptable to Management
 ~ dmn1
 ~ |

Change in adherence from schedule pressure=
 Table for fractional increase from
 schedule pressure(Schedule Pressure)
 ~ 1/Month
 ~ |

Change in adherence from management
 emphasis=
 Table for effect of management
 emphasis on adherence(Perceived Management
 emphasis on safety\
)
 ~ 1/Month
 ~ |

Change in adherence from personal threat=
 Table for fractional change from
 personal threat(Perceived Personal Threat)
 ~ 1/Month
 ~ |

Change in Adherence to Rules and Procedures=
 MIN(Maximum Increase in
 Adherence,Adherence to Rules and
 Procedures*Fractional Increase in adherence to
 Rules and Procedures\
)
 ~ 1/Month

~ |
 Reported Incident Rate=
 Incident Rate*Fraction of Incidents that
 are reported

~ Incidents/Month
 ~ |

Staff Level=
 smooth3(Desired Staff Level,Time to
 adjust Staff Level)

~ ppl
 ~ |

Table for effect of management emphasis on
 adherence(
 [(0,0)-
 (2,1)],(0,0),(1,0),(1.19878,0.140351),(1.5474,0.2
 45614),(2,0.3))

~ 1/Month
 ~ |

Maximum Increase in Adherence=
 (1-Adherence to Rules and
 Procedures)/Minimum time to increase
 adherence

~ 1/Month
 ~ |

Rate of New Tasks=
 Base Rate of new Tasks*Input

~ tasks/Month
 ~ |

Minimum time to increase adherence=
 0.1

~ Month
 ~ |

Time to perceive management emphasis on
 safety=

~ 2
 ~ Month
 ~ |

Table for fractional change from personal threat(
 [(0,-0.06)-(2,1)],(0,-
 0.01),(1,0),(1.24159,0.5),(1.5107,0.8),(2,1))

~ 1/Month
 ~ |

Table for fractional increase from schedule
 pressure(
 [(0.9,-2)-(1.2,0.1)],(0.9,0.1),(1,0),(1.05,-
 0.8),(1.1,-1.8))

~ 1/Month

~ |
 Adherence to Rules and Procedures= INTEG (
 Change in Adherence to Rules and
 Procedures,

0.5)
 ~ dmdl
 ~ |

Capacity=
 Staff Level*Normal work amount

~ hours/Month
 ~ |

Capacity Utilization=
 Table for Capacity Utilization(Schedule
 Pressure)

~ dmdl
 ~ |

Desired Task completion rate=
 Task Backlog/Desired Time to
 complete tasks

~ tasks/Month
 ~ |

Desired Time to complete tasks=

1
 ~ Month
 ~ |

Effect of Safety Behavior on Incident Rate=
 Table for Effect of Adherence on
 Incident Rate(Effectiveness of Safety Behavior)

~ dmdl
 ~ |

Exponential Growth Rate=

0
 ~ 1/Month
 ~ The exponential growth rate in
 the input.
 ~ |

Exponential Growth Time=

0
 ~ Month
 ~ The time at which the
 exponential growth in the input begins.
 ~ |

Incident Rate=
 Reference Incident Rate*Effect of
 Safety Behavior on Incident Rate

~ Incidents/Month
 ~ |

Input=
 $1 + \text{STEP}(\text{Step Height}, \text{Step Time}) +$
 $(\text{Pulse Quantity} / \text{TIME}$
 $\text{STEP}) * \text{PULSE}(\text{Pulse Time}, \text{TIME STEP}) +$
 $\text{RAMP}(\text{Ramp Slope}, \text{Ramp Start}$
 $\text{Time}, \text{Ramp End Time}) + \text{STEP}(1, \text{Exponential}$
 $\text{Growth Time}) * (\text{EXP}(\$
 $\text{Exponential Growth}$
 $\text{Rate} * \text{Time}) - 1) +$
 $\text{STEP}(1, \text{Sine Start Time}) * \text{Sine}$
 $\text{Amplitude} * \text{SIN}(2 * 3.14159 * \text{Time} / \text{Sine}$
 $\text{Period}) + \text{STEP}(1, \text{Noise Start Time} \backslash$
 $) * \text{RANDOM NORMAL}(-4, 4$
 $, 0, \text{Noise Standard Deviation}, \text{Noise Seed}) \backslash$
 \sim Dimensionless
 \sim The test input can be
configured to generate a step, pulse, linear ramp,
 \backslash
exponential growth, sine wave,
and random variation. The initial value of \backslash
the input is 1 and each test
input begins at a particular start time. The \backslash
magnitudes are expressed as
fractions of the initial value.

Reference Incident Rate=
1
 \sim Incident/Month
 \sim The theoretical minimum
incident rate, or the incident rate arising due to \backslash
chance, assuming rules and
procedures are maximally effective and are \backslash
followed completely.

Noise Seed=
1000
 \sim Dimensionless
 \sim Varying the random number
seed changes the sequence of realizations for \backslash
the random variable.

Noise Standard Deviation=
0
 \sim Dimensionless
 \sim The standard deviation in the
random noise. The random fluctuation is \backslash
drawn from a normal
distribution with min and max values of +/- 4.
The \backslash
user can also specify the
random number seed to replicate simulations.
To \backslash

generate a different random
number sequence, change the random number
seed.

Noise Start Time=
0
 \sim Month
 \sim The time at which the random
noise in the input begins.

Normal work amount=
8*5*4
 \sim hours/Month/ppl

Perceived Incident Rate=
smooth(Incident Rate, Time to perceive
incident Rate)
 \sim Incident/Month

Perceived Personal Threat=
ZIDZ(Perceived Incident Rate,
Threshold incident rate)
 \sim dmdl

Pulse Quantity=
0
 \sim Dimensionless*Month
 \sim The quantity added to the input
at the pulse time.

Pulse Time=
0
 \sim Month
 \sim The time at which the pulse
increase in the input occurs.

Ramp End Time=
1e+009
 \sim Month
 \sim The end time for the ramp
input.

Ramp Slope=
0
 \sim 1/Month
 \sim The slope of the linear ramp in
the input.

Ramp Start Time=
 0
 ~ Month
 ~ The time at which the ramp in
 the input begins.
 |

Schedule Pressure=
 Total Desired Work Rate/Capacity
 ~ dmdl
 ~ |

Sine Amplitude=
 0
 ~ Dimensionless
 ~ The amplitude of the sine wave
 in the input.
 |

Sine Period=
 10
 ~ Month
 ~ The period of the sine wave in
 the input.
 |

Sine Start Time=
 0
 ~ Month
 ~ The time at which the sine
 wave fluctuation in the input begins.
 |

Step Height=
 0
 ~ Dimensionless
 ~ The height of the step increase
 in the input.
 |

Step Time=
 0
 ~ Month
 ~ The time at which the step
 increase in the input occurs.
 |

Table for Capacity Utilization(
 [(0,0)-
 (1.5,1.2)],(0,0),(0.5,0.5),(0.75,0.75),(1,1),(1.1284
 4,1.08947),(1.2844,1.13684
),(1.49541,1.15))
 ~ dmdl
 ~ |

Table for Effect of Adherence on Incident Rate(
 [(0,0)-
 (1,5)],(0,5),(0.183486,3.50877),(0.25,3),(0.5,2),(
 0.755352,1.27193),(1,1))
 ~ dmdl
 ~ |

Task Backlog= INTEG (
 Rate of New Tasks-Task completion,
 Base Rate of new
 Tasks*Desired Time to complete tasks)
 ~ tasks
 ~ |

Task completion=
 (Capacity*Capacity Utilization)/Total
 Time per task
 ~ tasks/Month
 ~ |

Threshold incident rate=
 6
 ~ Incidents/Month
 ~ The incident rate at which an
 individual perceives an imminent personal \
 threat
 |

Time per task=
 5
 ~ hours/task
 ~ |

Time required per procedure=
 0.1
 ~ hours/procedure
 ~ |

Time to adjust Staff Level=
 6
 ~ months
 ~ |

Time to change rules adherence habits=
 2
 ~ Month
 ~ |

Time to perceive incident Rate=
 0.5
 ~ Month
 ~ |

Total Desired Work Rate=
 Desired Task completion rate*(Time
 per task+Safety Requirements per Task)

```

~      hours/Month
~      |
Total Time per task=
  Time per task+Safety Requirements per
Task
~      hours/task
~      |
*****
*****
. Control
*****
*****~
      Simulation Control Parameters
|
FINAL TIME = 36
~      Month
~      The final time for the
simulation.

```

```

|
INITIAL TIME = 0
~      Month
~      The initial time for the
simulation.
|
SAVEPER =
  TIME STEP
~      Month [0,?]
~      The frequency with which
output is stored.
|
TIME STEP = 0.0625
~      Month [0,?]
~      The time step for the
simulation.
|

```