Modeling, Analyzing, and Engineering NASA's Safety Culture

Phase 1 Final Report
September 2004 to February 2005

PI: Nancy G. Leveson[1,2]

Co-Investigators (alphabetical):
Dr. Betty Barrett[2]
Prof. John Carroll[2,3]
Prof. Joel Cutcher-Gershenfeld[2,3]
Nicolas Dulac[1]
David Zipkin[4]

[1]Aeronautics and Astronautics
[2]Engineering Systems Division
[3]Sloan School of Management
[4]Technology and Policy Program
Massachusetts Institute of Technology
77 Massachusetts Ave.
Cambridge, MA 02139

Sub-Award Number 05115-CIPI-03

Contents

1 Phase 1 Research Results and Accomplishments

1.1 Description of Research

The basic hypothesis being tested in this research is that:

> *Safety culture can be modeled, analyzed and engineered just like physical systems. The models will be useful in designing and validating improvements to the risk management and safety culture, in evaluating the potential impact of changes and policy decisions, in assessing risk, in detecting when risk is increasing to unacceptable levels, and in performing root cause analysis.*

A culture is commonly defined by sociologists as the shared set of norms and values that govern appropriate individual behavior. Safety culture is the subset of organizational culture that reflects the general attitude and approaches to safety and risk management. While safety is sometimes narrowly defined in terms of human death and injury, we use a more inclusive definition that also considers mission loss as a safety problem and is thus applicable to all the NASA enterprises and missions. The accident reports and investigations into the loss of the two Mars 98 missions and other NASA mission failures (for example, WIRE, Huygens, and the SOHO mission interruption) point to cultural problems very similar to those identified by the CAIB in the more visible manned space program and the need for similar cultural and organizational improvements. Although for practical reasons we will focus on the manned space program for this grant, the results will be applicable to all NASA missions and enterprises.

The literature on organizational culture draws on many disciplines. Social anthropologists and social psychologists emphasize the socially constructed nature of culture. Gareth Morgan, for example, defines culture as an ongoing, proactive process of reality construction [7]. Organizations then are, in essence, socially constructed realities that rest as much in the heads and minds of their members as they do in concrete sets of rules and regulations. Morgan asserts that organizations are "sustained by belief systems that emphasize the importance of rationality." This myth of rationality "helps us to see certain patterns of action as legitimate, credible, and normal, and hence to avoid the wrangling and debate that would arise if we were to recognize the basic uncertainty and ambiguity underlying many of our values and actions" [7, pp.134-135]. Another related view of culture is that it is a way of looking at and interpreting the world and events around us (our mental model) and taking action in a social context.

Culture is embedded in and arises from the routine aspects of everyday practice as well as organizational structures and rules. It includes the underlying or embedded operating assumptions under which actions are taken and decisions are made. Management, resources, capabilities, and culture are intertwined, and trying to change the culture without changing the environment within which the culture operates is doomed to failure. At the same time, simply changing the organizational structures—including policies, goals, missions, job descriptions, and standardized operating procedures related to safety—may lower risk over the short term but superficial fixes that do not address the set of shared values and social norms are very likely to be undone over time. The changes and protections instituted at NASA after the Challenger accident slowly degraded to the point where the same performance pressures and unrealistic expectations implicated in the Challenger accident contributed also to the Columbia loss. To achieve lasting results requires making broad changes that provide protection from and appropriate responses to the continuing environmental influences and pressures that tend to degrade the safety culture. "Sloganeering" is not enough—all aspects of the culture that affect safety must be engineered to be in alignment with the organizational safety principles.

We believe the following are all important social system aspects of a strong safety culture:

- The *formal organizational safety structure* including safety groups, such as the headquarters Office of Safety and Mission Assurance, the S&MA offices at each of the NASA centers and facilities, and now NESC (the new NASA Engineering and Safety Center), as well as the formal safety roles and responsibilities of managers, engineers, civil servants, contractors, etc. This formal structure has to be approached not as a static organization chart, but as a dynamic, constantly evolving set of formal relationships.
- *Organizational subsystems* impacting the safety culture and risk management including open and multi-directional communication systems; safety information systems to support planning, analysis, and decision making; reward and reinforcement systems that promote organizational learning; selection and retention systems that promote safety knowledge, skills, and ability; learning and feedback systems from incidents or hazardous events, in-flight anomalies (IFA's), and other aspects of operational experience; and channels and procedures for expressing safety concerns and resolving conflicts.
- *Individual behavior*, including knowledge, skills, and ability; group dynamics; and many psychological factors including fear of surfacing safety concerns, learning from mistakes without blame, commitment to safety values, and so on.
- *Safety rules and procedures* along with their underlying values and assumptions and a clearly expressed system safety vision. The vision must be shared among all the stakeholders, not just articulated by the leaders.

There are several assumptions about the NASA safety culture that underlie this research:

The Gap Between Vision and Reality: NASA as an organization has always had high expectations for safety and appropriately visible safety values and goals. Unfortunately, the operational practices have at times deviated from the stated organizational principles due to political pressures (both internal and external), unrealistic expectations, and other social factors. Several of the findings in the CAIB and Rogers Commission reports involve what might be termed a "culture of denial" where risk assessment was unrealistic and where credible risks and warnings were dismissed without appropriate investigation. Such a culture of denial is common where embedded, operating assumptions do not match the stated organizational policies. To "engineer" a safety culture, or, in other words, to bring the operational practices and values into alignment with the stated safety values, requires first identifying the desired organizational safety principles and values and then establishing and engineering the organizational infrastructure to achieve those values and to sustain them over time. Successfully achieving this alignment process requires understanding why the organization's operational practices have deviated from the stated principles and not only making the appropriate adjustments but instituting protections against future misalignments. These are the goals of this research.

No One Single Culture: NASA (and any other large organization) does not have a single "culture." Each of the centers, programs, projects, engineering disciplines within projects, and workforce groupings have their own subcultures. Creating an oversimplified view of the NASA "safety culture" and then trying to change that will be ineffective: Understanding and modeling efforts must be capable of differentiating among subcultures. Another inherent danger or risk in attempting to change cultures is that the unique aspects of an organization that contribute to, or are essential for, its success are changed or negatively influenced by the attempts to make the culture "safer." Culture change efforts must not negatively impact those aspects of NASA's culture that has made it great.

Mitigation of Risk, Not Elimination of Risk: Risk is an inherent part of space flight and exploration and other NASA missions. While risk cannot be eliminated from these activities, some practices involving *unnecessary* risk can be eliminated without impacting on NASA's success. The problem is to walk a tightrope between (1) a culture that thrives on and necessarily involves risks by the unique nature of its mission and (2) eliminating unnecessary risk that is detrimental to the overall NASA goals. Neither the Challenger nor the Columbia accidents involved *unknown unknowns*, but simply failure to handle known risks adequately. The goal should be to create a culture and organizational infrastructure that can resist pressures that militate against applying good safety engineering practices and procedures without requiring the elimination of the necessary risks of space flight. Most major accidents do not result from a unique set of proximal events but rather from the drift of the organization to a state of heightened risk over time as safeguards and controls are relaxed due to conflicting goals and tradeoffs. The challenge in preventing accidents is to establish safeguards and metrics to prevent and detect such changes before an accident occurs. NASA must establish the structures and procedures to ensure a healthy safety culture is established and sustained.

We believe that safety culture can be modeled, analyzed, and engineered. The goal of our research is to create a model of the current safety control structure and dynamic safety decision-making and review processes in NASA. Such modeling could potentially be used to evaluate and assess risk, to detect when risk is increasing, to evaluate the potential impact of changes and policies on risk, and to design organizational structures, including feedback loops, that will eliminate unnecessary risk from NASA missions. In addition, it can be used to determine the information each decision-maker needs to manage risk and the communication requirements for coordinated decision-making across large projects.

Phase 1 of the grant involved modeling the current NASA safety culture in the NASA manned space program, including the safety engineering, assessment, and review processes, as well as the pressures and influences that created the conditions that existed prior to Challenger and again before Columbia, in order to determine how to "engineer" lasting improvement. Such modeling and analysis can provide insight into the implications and relationships among the causal factors of the Shuttle accidents and into the long-term effectiveness of various possible changes in the safety culture and organization. Gaining this fundamental insight is the first step in effective safety culture transformation at NASA or any other organization. A further use for the models is as a living feedback tool for management at all levels of the program. Safety culture problems have the unique property of being very clear in retrospect, but hard to see when they are emerging. Cultural problems such as drift in focus, desensitization to dangerous conditions, and others are hard to see when you are part of the culture where these incremental shifts are taking place. Our models have the potential to serve like the "canary in the coal mine"—helping to make visible potentially dangerous patterns earlier than might otherwise be the case. Phase 1 demonstrated the feasibility of building such models.

The rest of this section describes (1) the relation of this research to the present state of the art, (2) the unique and potentially paradigm-changing approach used in this research to advance the state of the art, and (3) the models that were created. The following section describes the principle findings and expected outcomes and their potential benefits.

1.1.1 Relation to Present State of the Art

Current system safety and safety culture and management approaches are based on assumptions that do not fit the systems we are attempting to build today: They were created in an era of mechanical systems and then adapted for electro-mechanical systems, all of which do not begin to approach the levels of complexity and technological innovation in today's systems. We believe that to make significant progress we need new models and conceptions of how accidents occur that more accurately and completely reflect the types of accidents we are experiencing today.

At the foundation of the current limitations in engineering for safety is the almost exclusive use of a model of accidents that assumes they arise from a chain of failure events and human errors. The causal relationships between the events in the chain are direct and linear, representing the notion that the preceding event or condition must have been present for the subsequent event to occur, i.e., if event X had not occurred, then the following event Y would not have occurred. As such, event chain models encourage limited notions of linear causality, and they cannot account for indirect, non-linear, and feedback relationships. Unfortunately, such a model does not explain *system accidents*, where no components fail but the problem arises in the interaction among operating components. As such, it is inappropriate for today's software-intensive, human–machine systems where such system accidents are common. It also does not handle the complex human decision-making required to operate today's highly automated systems or the organizational and social aspects of safety and safety culture.

Because the theoretical model underlying safety engineering today does not include organizational factors, little has been accomplished to include them. A few people doing probabilistic risk assessment have tried to include human and management factors (for example, Pate-Cornell and Apostolakis) but their approaches, again based on event-chain models of accidents (which forms the basis for probabilistic risk assessment), do not handle the complexities involved in studying and modeling safety culture. The use of component reliability approaches that require everything to be reduced to a probability density function also reduces the effectiveness of this approach when human decision-making and social behavior is involved. Johnson tried to include management in his MORT (Management Oversight Risk Tree) accident investigation method [1], but simply ended up with a checklist of 1500 very general management practices that apply to everything.

Much has been written by sociologists about safety culture, but their views are limited by their lack of understanding of engineering problems and environment [6]. Normal Accident theorists and High Reliability Organization researchers are the primary contributors in this area. Perrow and his Normal Accident Theory [8] does a good job of understanding the problems of engineering complex systems, but he considers only a very limited set of possible solutions (primarily redundancy) and therefore reaches the pessimistic conclusion that nothing can be done and accidents are inevitable. The High Reliability Organization researchers do provide positive suggestions, but their studies have focused on relatively simple and loosely coupled systems, and their bottom-up component reliability approaches do not work for complex systems. The systems they have studied do not stretch the technological envelope nor do they operate in areas of engineering having high uncertainty (technically, socially, and politically) as does NASA. Another sociologist, Diane Vaughn, has written extensively about the NASA safety culture with respect to the Challenger accident, but her theory of "normalization of deviance" again oversimplies the problems of engineering this type of system and does not provide much practical guidance in how to improve safety culture.

1.1.2 Approach Used

The approach we used rests on a new way of thinking about accidents, called STAMP or Systems-Theoretic Accident Modeling and Processes [3], that integrates all aspects of risk, including organizational and social aspects. STAMP can be used as a foundation for new and improved approaches to accident investigation and analysis, hazard analysis and accident prevention, risk assessment and risk management, and devising risk metrics and performance monitoring. In this research, we will concentrate on its uses for risk assessment and management. One unique aspect of this approach to risk management is the emphasis on the use of visualization and building shared mental models of complex system behavior among those responsible for managing risk.

Another important difference between STAMP and other common approaches is the lack of focus on blame. The goal is not to identify the root cause or causes of an accident, but to understand "why" the accident occurred in terms of *all* contributors to the accident process and how to reengineer the socio-technical system as a whole to lower risk.

Systems are viewed in STAMP as interrelated components that are kept in a state of dynamic equilibrium by feedback loops of information and control. A socio-technical system is not treated as just a static design, but as a dynamic process that is continually adapting to achieve its ends and to react to changes in itself and its environment. The original design must not only enforce constraints on behavior to ensure safe operations, but it must continue to operate safely as changes and adaptations occur over time.

Accidents then are viewed as the result of flawed processes involving interactions among people, societal and organizational structures, engineering activities, and physical system components. The process leading up to an accident can be described in terms of an adaptive feedback function that fails to maintain safety as performance changes over time to meet a complex set of goals and values. The accident itself results not simply from component failure (which is treated as a symptom of the problems) but from inadequate control of safety-related constraints on the development, design, construction, and operation of the socio-technical system.

Safety in this model is treated as a *control* problem: Accidents occur when component failures, external disturbances, and/or dysfunctional interactions among system components are not adequately handled. In the Space Shuttle Challenger accident, for example, the O-rings did not adequately control the propellant gas release by sealing a tiny gap in the field joint. In the Mars Polar Lander loss, the software did not adequately control the descent speed of the spacecraft—it misinterpreted noise from a Hall effect sensor as an indication the spacecraft had reached the surface of the planet.

Accidents such as these, involving engineering design errors, may in turn stem from inadequate control of the development process, i.e., risk is not adequately managed in design, implementation, and manufacturing. Control is also imposed by the management functions in an organization—the Challenger and Columbia accidents, for example, involved inadequate controls in the launch-decision process and in the response to external pressures—and by the political system within which the organization exists.

While events reflect the *effects* of dysfunctional interactions and inadequate enforcement of safety constraints, the inadequate control itself is only indirectly reflected by the events—the events are the *result* of the inadequate control. The control structure itself must be carefully designed and evaluated to ensure that the controls are adequate to maintain the constraints on behavior necessary to control risk. This definition of risk management is broader than definitions that define it in terms of particular activities or tools. STAMP, which is based on systems and control theory, provides the theoretical foundation to develop the techniques and tools, including modeling tools, to assist managers in managing risk in this broad context.

Note that the use of the term "control" does not imply a strict military command and control structure. Behavior is controlled not only by direct management intervention but also indirectly by policies, procedures, shared values, and other aspects of the organizational culture as defined above. All behavior is influenced and at least partially "controlled" by the social and organizational context in which the behavior occurs. Engineering this context can be an effective way of creating and changing a safety culture.

STAMP is constructed from three fundamental concepts: constraints, hierarchical levels of control, and process models. These concepts, in turn, give rise to a classification of control flaws that can lead to accidents. Each of these is described only briefly here; for more information see [3].

The most basic component of STAMP is not an event, but a constraint. In systems theory and control theory, systems are viewed as hierarchical structures where each level imposes constraints on the activity of the level below it—that is, constraints or lack of constraints at a higher level allow or control lower-level behavior.

Safety-related constraints specify those relationships among system variables that constitute the non-hazardous or safe system states—for example, the power must never be on when the access to the high-voltage power source is open, the descent engines on the lander must remain on until the spacecraft reaches the planet surface, and two aircraft must never violate minimum separation requirements.

Instead of viewing accidents as the result of an initiating (root cause) event in a chain of events leading to a loss, accidents are viewed as resulting from interactions among components that violate the system safety constraints. The control processes that enforce these constraints must limit system behavior to the safe changes and adaptations implied by the constraints. Preventing accidents requires designing a control structure, encompassing the entire socio-technical system, that will enforce the necessary constraints on development and operations. Figure 1 shows a generic hierarchical safety control structure. Accidents result from inadequate enforcement of constraints on behavior (e.g., the physical system, engineering design, management, and regulatory behavior) at each level of the socio-technical system. Inadequate control may result from missing safety constraints, inadequately communicated constraints, or from constraints that are not enforced correctly at a lower level. Feedback during operations is critical here. For example, the safety analysis process that generates constraints always involves some basic assumptions about the operating environment of the process. When the environment changes such that those assumptions are no longer true, the controls in place may become inadequate.

The model in Figure 1 has two basic hierarchical control structures—one for system development (on the left) and one for system operation (on the right)—with interactions between them (not all shown in order to declutter the diagram). A spacecraft manufacturer, for example, might only have system development under its immediate control, but safety involves both development and operational use of the spacecraft, and neither can be accomplished successfully in isolation: Safety must be designed into the physical system, and safety during operation depends partly on the original system design and partly on effective control over operations. Manufacturers must communicate to their customers the assumptions about the operational environment upon which their safety analysis and design was based, as well as information about safe operating procedures. The operational environment, in turn, provides feedback to the manufacturer about the performance of the system during operations.

Between the hierarchical levels of each control structure, effective communication channels are needed, both a downward *reference* channel providing the information necessary to impose constraints on the level below and a *measuring* channel to provide feedback about how effectively the constraints were enforced. For example, company management in the development process structure may provide a safety policy, standards, and resources to project management and in return receive status reports, risk assessment, and incident reports as feedback about the status of the project with respect to the safety constraints.

The safety control structure often changes over time, which accounts for the observation that accidents in complex systems frequently involve a migration of the system toward a state where a small deviation (in the physical system or in human behavior) can lead to a catastrophe. The foundation for an accident is often laid years before. One event may trigger the loss, but if that event had not happened, another one would have. As an example, Figure 2 shows the changes over time that led to a water contamination accident in Canada where 2400 people became ill and 7 died (most of them children) [4]. The reasons why this accident occurred would take too many

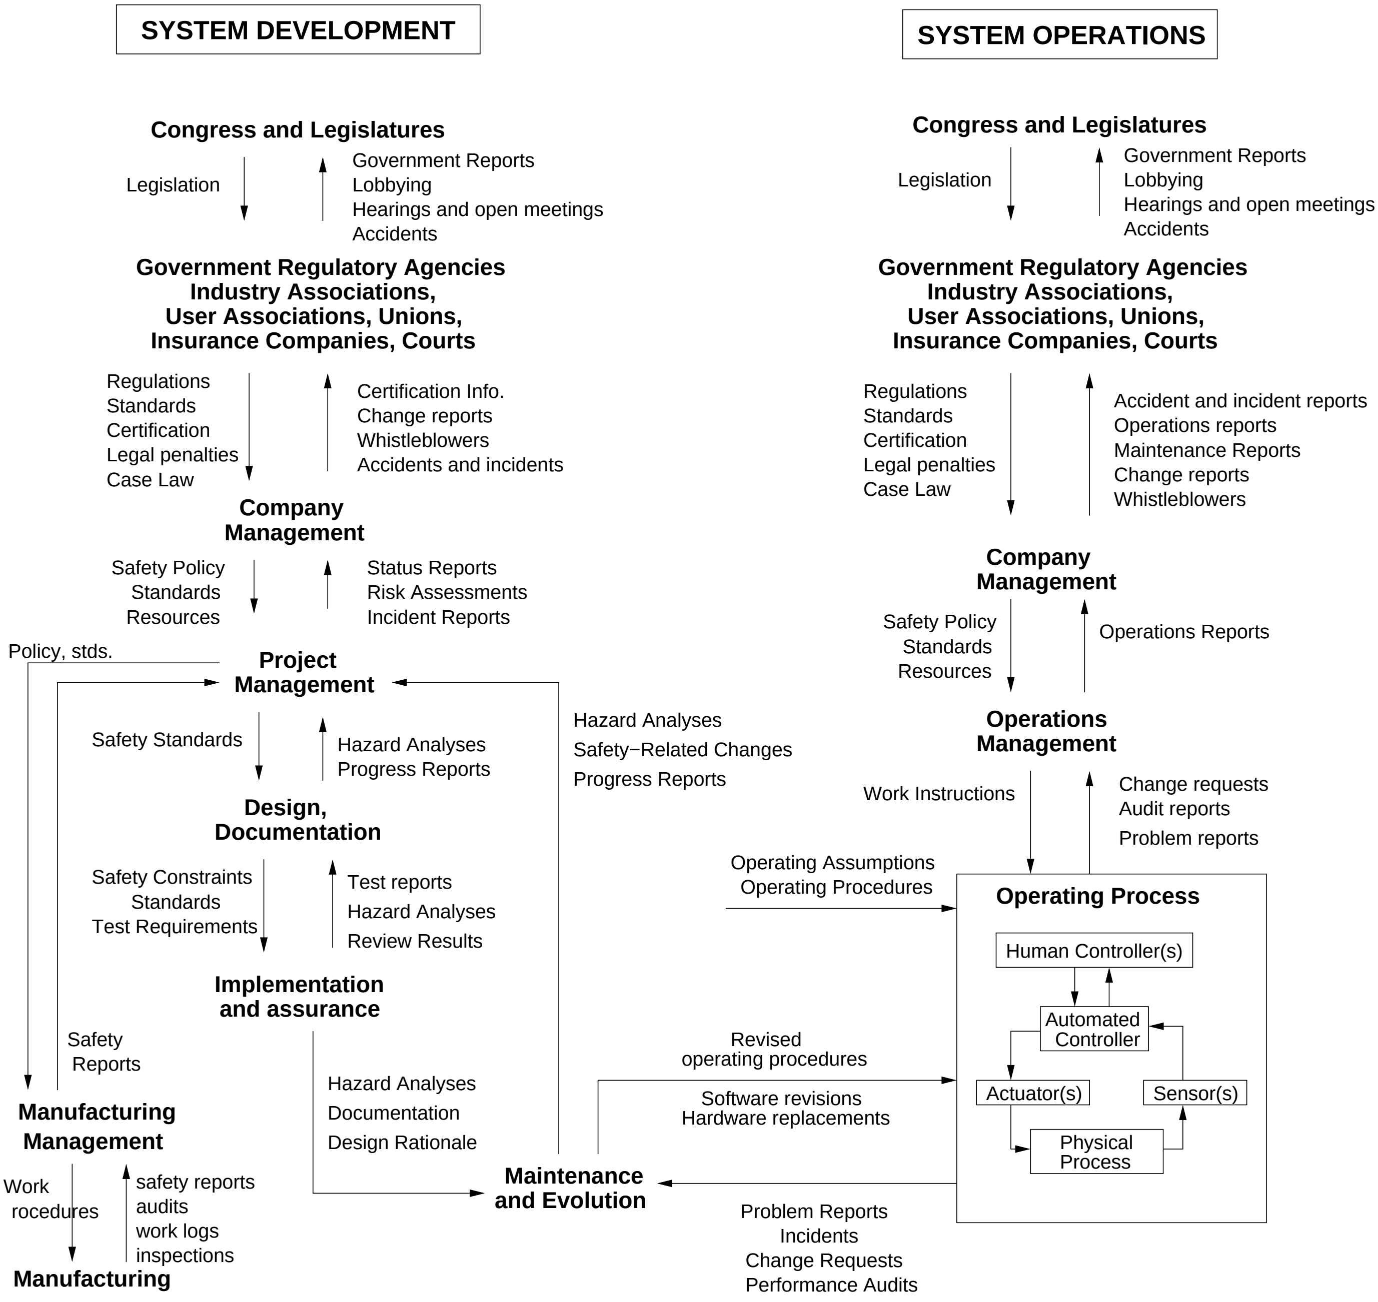

Figure 1: General Form of a Model of Socio-Technical Control.

pages to explain and only a small part of the overall STAMP model is shown. Each component of the water quality control structure played a role in the accident. The model at the top shows the control structure for water quality in Ontario Canada as designed. The figure at the bottom shows the control structure as it existed at the time of the accident. One of the important changes that contributed to the accident is the elimination of a government water testing laboratory. The private companies that were substituted were not required to report instances of bacterial contamination to the appropriate government ministries. Essentially, the elimination of the feedback loops made it impossible for the government agencies and public utility managers to perform their oversight duties effectively. Again note that the goal here is not to identify individuals to blame for the accident but to understand why they made the mistakes they made (none were evil or wanted children to die) and what changes are needed in the culture and water quality control structure to reduce risk in the future.

In this accident, and in most accidents, degradation in the safety margin occurred over time and without any particular single decision to do so but simply as a series of decisions that individually seemed safe but together resulted in moving the water quality control system structure slowly toward a situation where any slight error would lead to a major accident. An effective risk-management system must ensure that controls do not degrade as happened with both the Challenger and Columbia losses.

Figure 2 shows static models of the safety control structure. But models are also needed to understand *why* the structure changed over time in order to build in protection against unsafe changes. For this goal, we use system dynamics models. The field of system dynamics, created at MIT in the 1950s by Forrester, is designed to help decision makers learn about the structure and dynamics of complex systems, to design high leverage policies for sustained improvement, and to catalyze successful implementation and change. System dynamics provides a framework for dealing with dynamic complexity, where cause and effect are not obviously related. Like the other STAMP models, it is grounded in the theory of non-linear dynamics and feedback control, but also draws on cognitive and social psychology, organization theory, economics, and other social sciences [10]. System dynamics models are formal and can be executed, like our other models.

System dynamics is particularly relevant for complex systems and systems of systems. The world is dynamic, evolving, and interconnected, but we tend to make decisions using mental models that are static, narrow, and reductionist. Thus decisions that might appear to have no effect on safety—or even appear to be beneficial—may in fact degrade safety and increase risk. System dynamics makes it possible, for example, to understand and predict instances of policy resistance or the tendency for well-intentioned interventions to be defeated by the response of the system to the intervention itself. In related but separate research, Marais and Leveson are working on defining archetypical system dynamics models often associated with accidents to assist in creating the models for specific systems [5].

Figure 3 shows a simple systems dynamics model of the Columbia accident. This model is only a hint of what a complete model might contain. The loops in the figure represent feedback control loops where the "+" or "-" on the loops represent polarity or the relationship (positive or negative) between state variables: a positive polarity means that the variables move in the same direction while a negative polarity means that they move in opposite directions. There are three main variables in the model: safety, complacency, and success in meeting launch rate expectations.

The control loop in the lower left corner of Figure 3, labeled R1 or *Pushing the Limit*, shows how as external pressures increased, performance pressure increased, which led to increased launch rates and thus success in meeting the launch rate expectations which in turn led to increased expectations and increasing performance pressures. This, of course, is an unstable system and cannot be maintained indefinitely—note the larger control loop, B1, in which this loop is embedded,

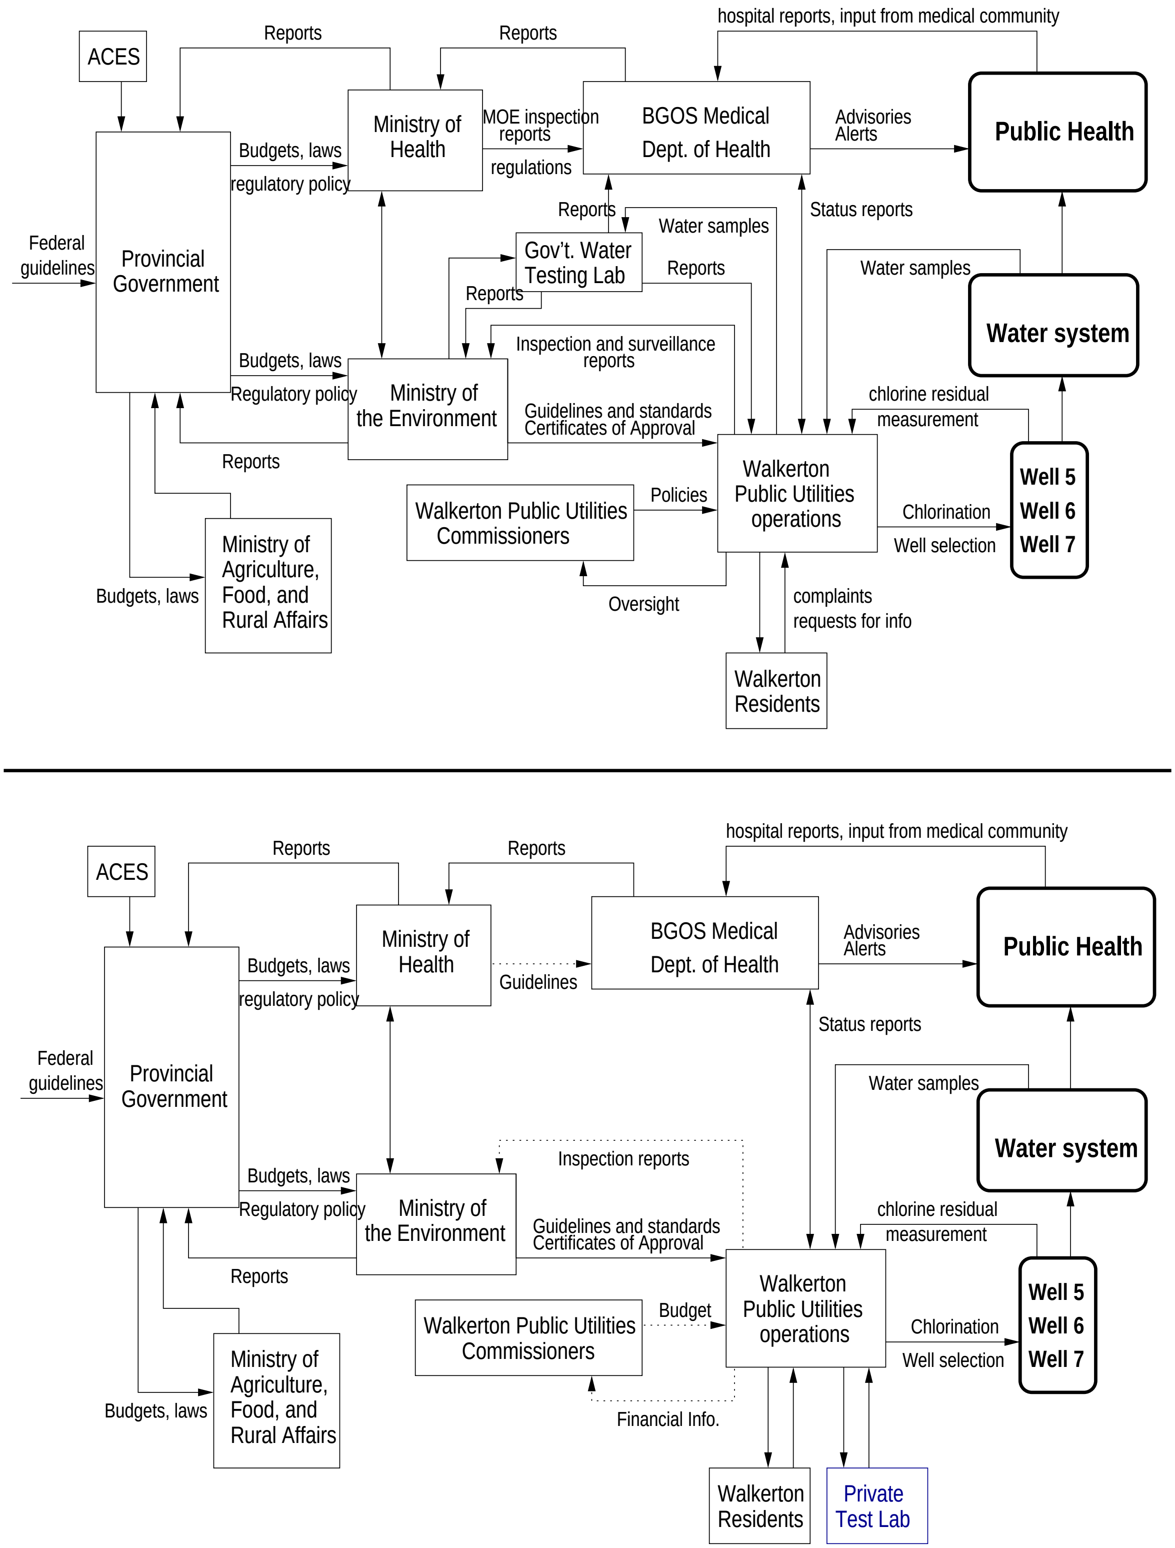

Figure 2: The Safety Control Structure in the Walkerton Water Contamination Accident. The structure is drawn in the form commonly used for control loops. Lines going into the left of a box are control lines. Lines from or to the top or bottom of a box represent information, feedback, or a physical flow. Rectangles with sharp corners are controllers while rectangles with rounded corners represent plants.

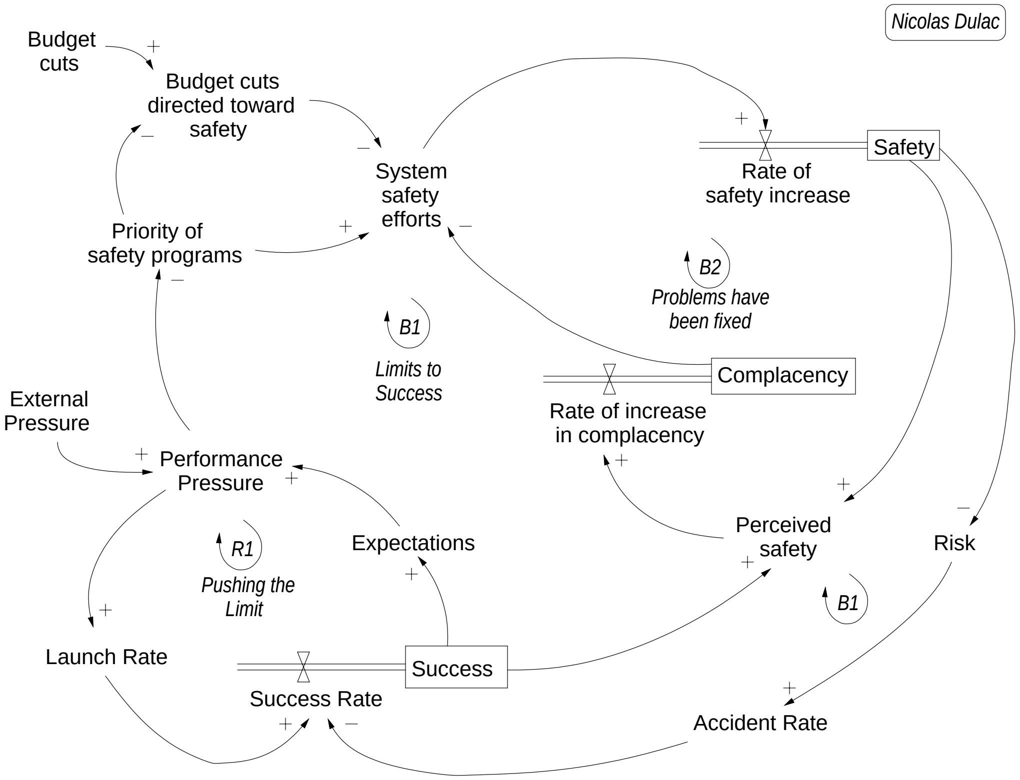

Figure 3: Simplified Model of the Dynamics Behind the Shuttle Columbia Loss.

is labeled *Limits to Success*. The upper left loop represents part of the safety program loop. The external influences of budget cuts and increasing performance pressures that reduced the priority of safety procedures led to a decrease in system safety efforts. The combination of this decrease along with loop B2 in which fixing problems increased complacency, which also contributed to reduction of system safety efforts, eventually led to a situation of (unrecognized) high risk. One thing not shown in the diagram is that these models also can contain delays. While reduction in safety efforts and lower prioritization of safety concerns may lead to accidents, accidents usually do not occur for a while so false confidence is created that the reductions are having no impact on safety and therefore pressures increase to reduce the efforts and priority even further as the external performance pressures mount.

The models can be used to devise and validate fixes for the problems. For example, one way to eliminate the instability of the model in Figure 3 is to anchor the safety efforts by, perhaps, externally enforcing standards in order to prevent schedule and budget pressures from leading to reductions in the safety program. Other solutions are also possible. Alternatives can be evaluated for their potential effects using a more complete system dynamics model, like the one produced for this grant and described later in this report.

Often degradation of the control structure involves *asynchronous evolution* where one part of a system changes without the related necessary changes in other parts. Changes to subsystems may be carefully designed, but consideration of their effects on other parts of the system, including the control aspects, may be neglected or inadequate. Asynchronous evolution may also occur when one part of a properly designed system deteriorates. The Ariane 5 trajectory changed from that of the Ariane 4, but the inertial reference system software did not. One factor in the loss of contact with the SOHO (SOlar Heliospheric Observatory) spacecraft in 1998 was the failure to communicate to operators that a functional change had been made in a procedure to perform gyro spin-down.

Besides constraints and hierarchical levels of control, a third basic concept in STAMP is that of process models. *Any* controller—human or automated—must contain a model of the system being controlled. The figure below shows a typical control loop where an automated controller is supervised by a human controller.

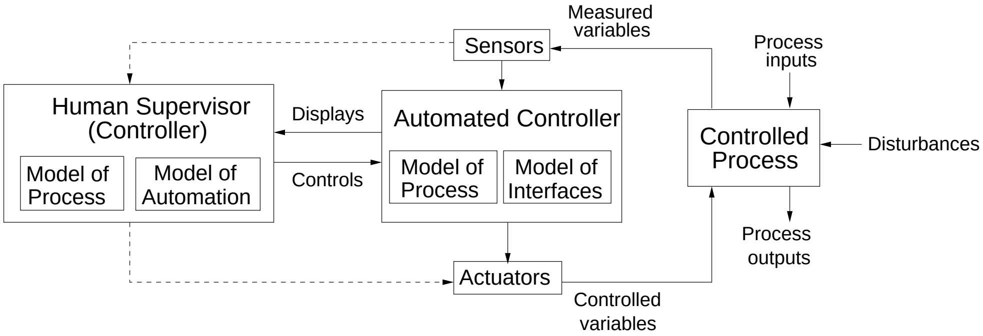

Any controller must have a model (for human controllers this is a mental model) of (1) the current state of the system being controlled, (2) the required relationship between system variables, and (3) the ways the process can change state. Accidents, particularly system accidents, frequently result from inconsistencies between the model of the process used by the controllers and the actual process state; for example, the lander software thinks the lander has reached the surface and shuts down the descent engine; the Minister of Health has received no reports about water quality problems and believes the state of water quality in the town is better than it actually is; or a mission manager believes that foam shedding is a maintenance or turnaround issue only. Part of our modeling efforts involve creating the process models, examining the ways that they can become inconsistent with the actual state (e.g., missing or incorrect feedback), and determining what feedback loops are necessary to maintain the safety constraints.

When there are multiple controllers and decision makers, system accidents may also involve inadequate control actions and unexpected side effects of decisions or actions, again often the result of inconsistent process models. For example, two controllers may both think the other is making the required control action, or they make control actions that conflict with each other. Communication plays an important role here. Leplat suggests that accidents are most likely in *boundary* or *overlap* areas where two or more controllers control the same process [2].

A STAMP analysis involves creating a model of the organizational safety structure including the static safety control structure and the safety constraints that each component is responsible for maintaining, a model of the dynamics and pressures that can lead to degradation of this structure over time, process models representing the view of the process by those controlling it, and a model of the cultural and political context in which decision making occurs. The information that results from the modeling and analysis of the models can be used to understand and analyze the risk in both the current organizational culture and structure and in potential changes, to devise policies and changes that can decrease risk and evaluate their implications with respect to other important goals, and to create metrics and other performance measures to identify when risk is increasing to unacceptable levels.

Our goal for Phase 1 was to build the models for the manned space program. To accomplish our Phase 1 goals, we were able to take advantage of the PIs long-term association with NASA to interview current and former employees of the manned space program. We combined the results of these conversations with information from books on NASA's safety culture (such as Howard McCurdy's *Inside NASA: High Technology and Organizational Change in the U.S. Space Program*), books on the Challenger and Columbia accidents, NASA mishap reports (CAIB, Mars Polar Lander, Mars Climate Orbiter, WIRE, SOHO, Huygens, etc.), other NASA reports on the manned space program (SIAT or Shuttle Independent Assessment Team Report and others) as well as many of the better researched magazine and newspaper articles.

2 Results, Principle Findings, and Anticipated Outcomes/Benefits

Our results fall into two categories: the models built and the preliminary analysis of the aspects of the NASA safety culture that were modeled. We first describe some challenges we faced and the solutions we adopted for them and then the models and analysis results.

2.1 Problems Encountered and Remedial Steps Taken

The biggest problem we encountered was in understanding the complex management relationships and roles within the Shuttle program. We had anticipated this problem would occur: The CAIB report noted the Manned Space Flight program has confused lines of authority, responsibility, and accountability in a "manner that almost defies explanation." The PI had similar experiences while trying to understand the control structure in her role on the NASA Aerospace Safety Advisory Panel (ASAP) and on various other NASA Advisory committees. We modeled this structure as accurately as possible without excessively bothering people during the return to flight process. In the end, we decided that the current control structure was too much of a mess to provide a useful model for analysis and decided to focus most of our analysis effort on the behavioral dynamics of the NASA safety culture using system dynamics models, which we believe we were able to model accurately.

2.2 The Models

Most of our focus was on building the system dynamics model as we felt it was the most difficult to construct and would be the most helpful in evaluating the feasibility and usefulness of our new approach to risk management.

System behavior in system dynamics is modeled by using feedback (causal) loops, stock and flows (levels and rates), and the non-linearities created by interactions among system components. In this view of the world, behavior over time (the dynamics of the system) can be explained by the interaction of positive and negative feedback loops [9]. The models are constructed from three basic building blocks: positive feedback or reinforcing loops, negative feedback or balancing loops, and delays. Positive loops (called reinforcing loops) are self-reinforcing while negative loops tend to counteract change. Delays introduce potential instability into the system.

Figure 4a shows a *reinforcing loop*, which is a structure that feeds on itself to produce growth or decline. Reinforcing loops correspond to positive feedback loops in control theory. An increase in variable 1 leads to an increase in variable 2 (as indicated by the "+" sign), which leads to an increase in variable 1 and so on. The "+" does not mean the values necessarily increase, only that variable 1 and variable 2 will change in the same direction. If variable 1 decreases, then variable 2 will decrease. A "-" indicates that the values change in opposite directions. In the absence of external influences, both variable 1 and variable 2 will clearly grow or decline exponentially. Reinforcing loops generate growth, amplify deviations, and reinforce change [10].

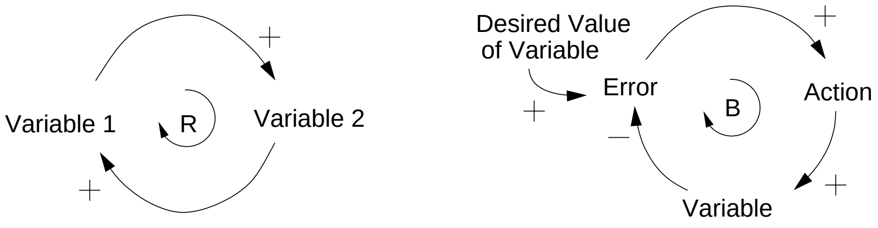

a. A Reinforcing Loop

b. A Balancing Loop

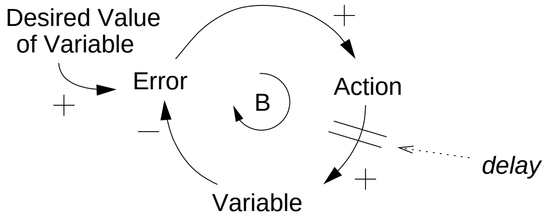

c. A Balancing Loop with a Delay

Figure 4: The Three Basic Components of System Dynamics Models

A *balancing* loop (Figure 4b) is a structure that changes the current value of a system variable or a desired or reference variable through some action. It corresponds to a negative feedback loop in control theory. The difference between the current value and the desired value is perceived as an error. An action proportional to the error is taken to decrease the error so that, over time, the current value approaches the desired value.

The third basic element is a delay, which is used to model the time that elapses between cause and effect. A delay is indicated by a double line as shown in Figure 4c. Delays make it difficult to link cause and effect (dynamic complexity) and may result in unstable system behavior. For example, in steering a ship there is a delay between a change in the rudder position and a corresponding course change, often leading to over-correction and instability.

The simple "News Sharing" model in Figure 5 is helpful in understanding the stock and flow syntax and the results of our modeling effort. The model shows the flow of information through a population over time. The total population is fixed and includes 100 people. Initially, only one person knows the news, the other 99 people do not know it. Accordingly, there are two *stocks* in the model: *People who know* and *People who don't know.* The initial value for the *People who know* stock is 1 and that for the *People who don't know* stock is 99. Once a person learns the news, he or she moves from the left-hand stock to the right-hand stock through the double arrow flow called *Rate of sharing the news.* The rate of sharing the news at any point in time depends on the number of *Contacts between people who know and people who don't*, which is function of the value of the two stocks at that time. This function uses a differential equation, i.e., the rate of change of a variable V, i.e., dV/dt, at time t depends on the value of V(t). The results for each stock and variable as a function of time are obtained through a standard numerical integration routine using the following formulations:

$$People\ who\ know(t) = \int_0^t Rate\ of\ sharing\ the\ news \tag{1}$$

$$People\ who\ know(0) = 1 \tag{2}$$

$$People\ who\ don't\ know(0) = 99 \tag{3}$$

$$People\ who\ don't\ know(t) = \int_0^t -rate\ of\ sharing\ the\ news \tag{4}$$

$$Total\ People = People\ who\ don't\ know(t) + People\ who\ know(t) \tag{5}$$

$$\begin{gathered} Rate\ of\ sharing\ the\ news(t) = \\ Contacts\ between\ people\ who\ know\ and\ people\ who\ don't(t) \end{gathered} \tag{6}$$

$$\begin{gathered} Contacts\ between\ people\ who\ know\ and\ people\ who\ don't(t) = \\ \frac{People\ who\ don't\ know(t) \times People\ who\ know(t)}{Total\ People} \end{gathered} \tag{7}$$

The graph in Figure 5 shows the numerical simulation output for the number of people who know, the number of people who don't know, and the rate of sharing the news as a function of time.

One of the significant challenges associated with modeling a socio-technical system as complex as the Shuttle program is creating a model that captures the critical intricacies of the real-life

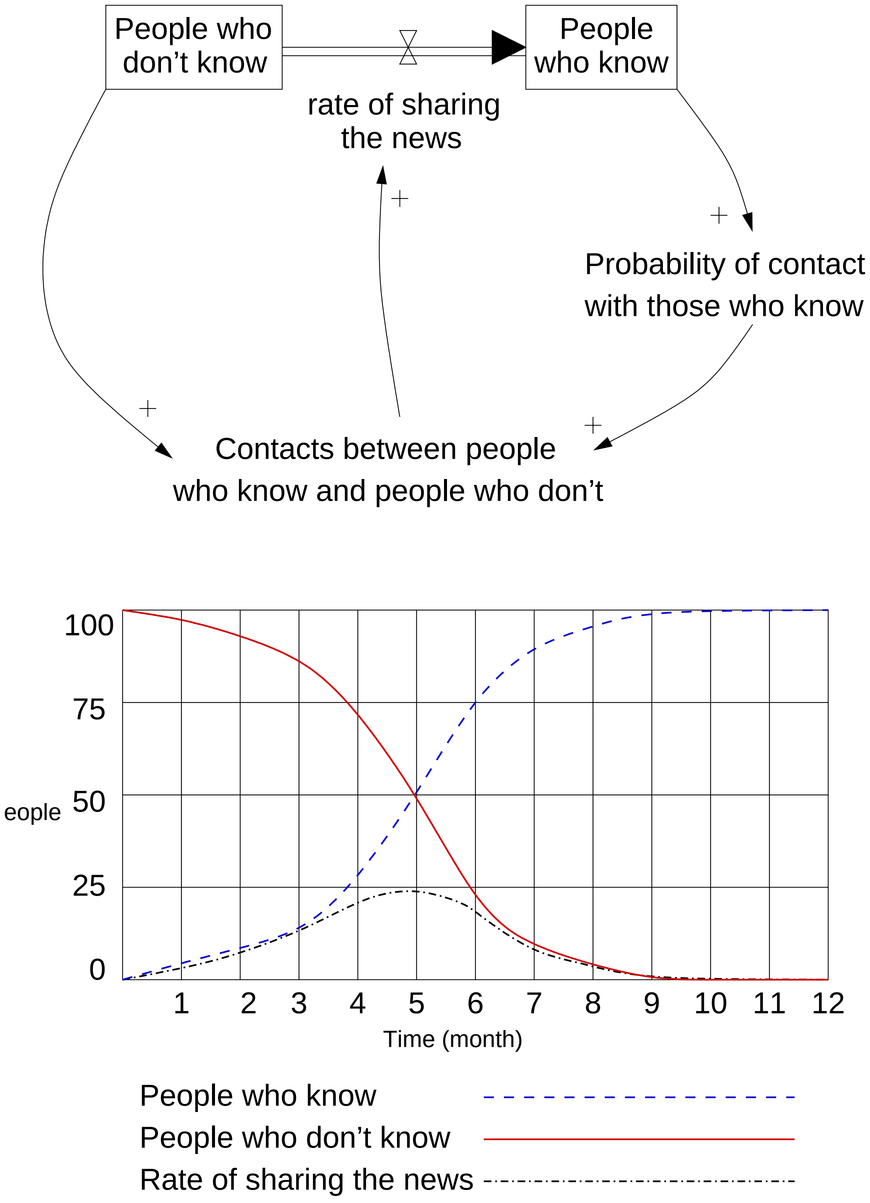

Figure 5: An Example Output from a Systems Dynamics Model

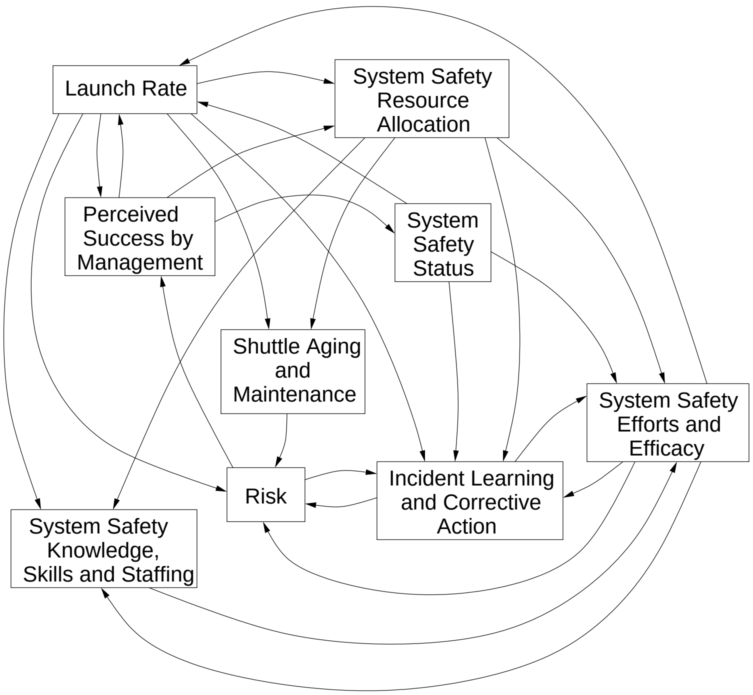

Figure 6: The Nine Submodels and Their Interactions

system, but is not so complex that it cannot be readily understood. To be accepted and therefore useful to risk decision makers, a model must have the confidence of the users and that confidence will be limited if the users cannot understand what has been modeled. We addressed this problem by breaking the overall system model into nine logical subsystem models, each of a intellectually manageable size and complexity. The subsystem models can be built and tested independently and then, after validation and comfort with the correctness of each subsystem model is achieved, the subsystem models can be connected to one another so that important information can flow between them and emergent properties that arise from their interactions can be included in the analysis. Figure 6 shows the nine model components along with the interactions among them.

As an example, our Launch Rate model uses a number of internal factors to determine the frequency at which the Shuttle can be launched. That value—the “output” of the Launch Rate model—is then used by many other subsystem models including the Risk model and the Perceived Success by High-Level Management models.

The nine subsystem models are:

- Launch Rate
- System Safety Resource Allocation
- System Safety Status
- Incident Learning and Corrective Action
- Technical Risk
- System Safety Efforts and Efficacy

- Shuttle Aging and Maintenance
- System Safety Knowledge Skills and Staffing
- Perceived Success by High-Level Management

Each of these submodels is described in more detail below, including both the outputs of the submodel and the factors used to determine the results. The models themselves are included in the Appendix.

Technical Risk: The purpose of the technical risk model is to determine the level of occurrence of anomalies and hazardous events, as well as the interval between accidents. The assumption behind the risk formulation is that once the system has reached a state of high risk, it is highly vulnerable to small deviations that can cascade into major accidents. The primary factors affecting the technical risk of the system are the effective age of the Shuttle, the quantity and quality of inspections aimed at uncovering and correcting safety problems, and the proactive hazard analysis and mitigation efforts used to continuously improve the safety of the system. Another factor affecting risk is the response of the program to anomalies and hazardous events (and, of course, mishaps or accidents).

The response to anomalies, hazardous events, and mishaps can either address the symptoms of the underlying problem or the root causes of the problems. Corrective actions that address the symptoms of a problem have insignificant effect on the technical risk and merely allow the system to continue operating while the underlying problems remain unresolved. On the other hand, corrective actions that address the root cause of a problem have a significant and lasting positive effect on reducing the system technical risk.

System Safety Resource Allocation: The purpose of the resource allocation model is to determine the level of resources allocated to system safety. To do this, we model the factors determining the portion of NASA's budget devoted to system safety. The critical factors here are the priority of the safety programs relative to other competing priorities such as launch performance and NASA safety history. The model assumes that if performance expectations are high or schedule pressure is tight, safety funding will decrease, particularly if NASA has had past safe operations.

System Safety Status: The safety organization's status plays an important role throughout the model, particularly in determining effectiveness in attracting high-quality employees and determining the likelihood of other employees becoming involved in the system safety process. Additionally, the status of the safety organization plays an important role in determining their level of influence, which in turn, contributes to the overall effectiveness of the safety activities. Management prioritization of system safety efforts plays an important role in this submodel, which in turn influences such safety culture factors as the power and authority of the safety organization, resource allocation, and rewards and recognition for raising safety concerns and placing emphasis on safety. In the model, the status of the safety organization has an impact on the ability to attract highly capable personnel; on the level of morale, motivation, and influence; and on the amount and effectiveness of cross-boundary communication.

Safety Knowledge, Skills, and Staffing: The purpose of this submodel is to determine both the overall level of knowledge and skill in the system safety organization and to determine if the number of NASA system safety engineers is sufficient to oversee the contractors. These two values are used by the System Safety Effort and Efficacy submodel.

In order to determine these key values, the model tracks four quantities: the number of NASA employees working in system safety, the number of contractor system safety employees, the aggregate experience of the NASA employees, and the aggregate experience of the system safety contractors such as those working for United Space Alliance (USA) and other major Shuttle contractors.

The staffing numbers rise and fall based on the hiring, firing, attrition, and transfer rates of the employees and contractors. These rates are determined by several factors, including the amount of safety funding allocated, the portion of work to be contracted out, the age of NASA employees, and the stability of funding.

The amount of experience of the NASA and contractor system safety engineers relates to the new staff hiring rate and the quality of that staff. An organization that highly values safety will be able to attract better employees who bring more experience and can learn faster than lower quality staff. The rate at which the staff gains experience is also determined by training, performance feedback, and the workload they face.

Shuttle Aging and Maintenance: The age of the Shuttle and the amount of maintenance, refurbishments, and safety upgrades affects the technical risk of the system and the number of anomalies and hazardous events. The effective Shuttle age is mainly influenced by the launch rate. A higher launch rate will accelerate the aging of the Shuttle unless extensive maintenance and refurbishment are performed. The amount of maintenance depends on the resources available for maintenance at any given time. As the system ages, more maintenance may be required; if the resources devoted to maintenance are not adjusted accordingly, accelerated aging will occur.

The original design of the system also affects the maintenance requirements. Many compromises were made during the initial phase of the Shuttle design, trading off lower development costs for higher operations costs. Our model includes the original level of design for maintainability, which allows the investigation of scenarios during the analysis where system maintainability would have been a high priority from the beginning.

While launch rate and maintenance affect the rate of Shuttle aging, refurbishment and upgrades *decrease* the effective aging by providing complete replacements and upgrade of Shuttle systems such as avionics, fuel systems, and structural components. The amount of upgrades and refurbishment depends on the resources available, as well as on the perception of the remaining life of the system. Upgrades and refurbishment will most likely be delayed or canceled when there is high uncertainty about the remaining operating life. Uncertainty will be higher as the system approaches or exceeds its original design lifetime, especially if there is no clear vision and plan about the future of the manned space program.

Launch Rate: The Launch Rate submodel is at the core of the integrated model. Launch rate affects many parts of the model, such as the perception of the level of success achieved by the Shuttle program. A high launch rate without accidents contributes to the perception that the program is safe, eventually eroding the priority of system safety efforts. A high launch rate also accelerates system aging and creates schedule pressure, which hinders the ability of engineers to perform thorough problem investigation and to implement effective corrective actions that address the root cause of the problems rather than just the symptoms.

The launch rate in the model is largely determined by three factors:

1. Expectations from high-level management: Launch expectations will most likely be high if the program has been successful in the recent past. The expectations are reinforced through

a "Pushing the Limits" phenomenon where administrators expect ever more from a successful program, without necessarily providing the resources required to increase launch rate;

2. Schedule pressure from the backlog of flights scheduled: This backlog is affected by the launch commitments, which depend on factors such as ISS commitments, Hubble servicing requirements, and other scientific mission constraints;
3. Launch delays that may be caused by unanticipated safety problems: The number of launch delays depends on the technical risk, on the ability of system safety to uncover problems requiring launch delays, and on the power and authority of system safety personnel to delay launches.

System Safety Efforts and Efficacy: This submodel captures the effectiveness of system safety at identifying, tracking, and mitigating Shuttle system hazards. The success of these activities will affect the number of hazardous events and problems identified, as well as the quality and thoroughness of the resulting investigations and corrective actions. In the model, a combination of reactive problem investigation and proactive hazard mitigation efforts leads to effective safety-related decision making that reduces the technical risk associated with the operation of the Shuttle. While effective system safety activities will improve safety over the long run, they may also result in a decreased launch rate over the short run by delaying launches when serious safety problems are identified.

The efficacy of the system safety activities depends on various factors. Some of these factors are defined outside this submodel, such as the availability of resources to be allocated to safety and the availability and effectiveness of safety processes and standards. Others depend on characteristics of the system safety personnel themselves, such as their number, knowledge, experience, skills, motivation, and commitment. These personal characteristics also affect the ability of NASA to oversee and integrate the safety efforts of contractors, which is one dimension of system safety effectiveness. The quantity and quality of lessons learned and the ability of the organization to absorb and use these lessons is also a key component of system safety effectiveness.

Hazardous Event (Incident) Learning and Corrective Action: The objective of this submodel is to capture the dynamics associated with the handling and resolution of safety-related anomalies and hazardous events. It is one of the most complex submodels, reflecting the complexity of the cognitive and behavioral processes involved in identifying, reporting, investigating, and resolving safety issues. Once integrated into the combined model, the amount and quality of learning achieved through the investigation and resolution of safety problems impacts the effectiveness of system safety efforts and the quality of resulting corrective actions, which in turn has a significant effect on the technical risk of the system.

The structure of this model revolves around the processing of incidents or hazardous events, from their initial identification to their eventual resolution. The number of safety-related incidents is a function of the technical risk. Some safety-related problems will be reported while others will be left unreported. The fraction of safety problems reported depends on the effectiveness of the reporting process, the employee sensitization to safety problems, the possible fear of reporting if the organization discourages it, perhaps due to the impact on schedule. Problem reporting will increase if employees see that their concerns are considered and acted upon, that is, if they have previous experience that reporting problems led to positive actions. The reported problems also varies as a function of the perceived safety of the system by engineers and technical workers. A problem-reporting positive feedback loop creates more reporting as the perceived risk increases, which is influenced by the number of problems reported and addressed. Numerous studies have shown that the risk perceived by engineers and technical workers is different from high-level management perception of risk. In our model, high-level management and engineers use different cues to evaluate risk and safety, which results in very different assessments.

A fraction of the anomalies reported are investigated in the model. This fraction varies based on the resources available, the overall number of anomalies being investigated at any time, and the thoroughness of the investigation process. The period of time the investigation lasts will also depend on these same variables.

Once a hazardous event or anomaly has been investigated, four outcomes are possible: (1) no action is taken to resolve the problem, (2) a corrective action is taken that only addresses the symptoms of the problem, (3) a corrective action is performed that addresses the root causes of the problem, and (4) the proposed corrective action is rejected, which results in further investigation until a more satisfactory solution is proposed. Many factors are used to determine which of these four possible outcomes results, including the resources available, the schedule pressure, the quality of hazardous event or anomaly investigation, the investigation and resolution process and reviews, and the effectiveness of system safety decision-making. As the organization goes through this ongoing process of problem identification, investigation, and resolution, some lessons are learned, which may be of variable quality depending on the investigation process and thoroughness. In our model, if the safety personnel and decision-makers have the capability and resources to extract and internalize high-quality lessons from the investigation process, their overall ability to identify and resolve problems and do effective hazard mitigation will be enhanced.

Perceived Success by High-Level Management The purpose if this submodel is to capture the dynamics behind the success of the Shuttle program as perceived by high-level management and NASA administration. The success perceived by high-level management is a major component of the Pushing the Limit reinforcing loop, where much will be expected from a highly successful program, creating even higher expectations and performance pressure. High perceived success also creates the impression by high-level management that the system is inherently safe and can be considered operational, thus reducing the priority of safety, which affects resource allocation and system safety status. Two main factors contribute to the perception of success: the accumulation of successful launches positively influences the perceived success while the occurrence of accidents and mishaps have a strong negative influence.

2.3 Principle Findings and Anticipated Outcomes/Benefits

The models we constructed can be used in many ways, including understanding how and why accidents have occurred, testing and validating changes and new policies (including risk and vulnerability assessment of policy changes), learning which "levers" have a significant and sustainable effect, and facilitating the identification and tracking of metrics to detect increasing risk. But in order to trust the models and the results from their analysis, the users need to be comfortable with the models and their accuracy.

We first validated each model individually, using (1) review by experts familiar with NASA and experts on safety culture in general and (2) execution of the models to determine whether the results were reasonable.

Once we were comfortable with the individual models, we ran the integrated model using baseline parameters. In the graphs that follow, the arrows on the x-axis (timeline) indicate when accidents occur during the model execution (simulation). Also, it should be noted that we are *not* doing risk *assessment*, i.e., quantitative or qualitative calculation of the likelihood and severity of

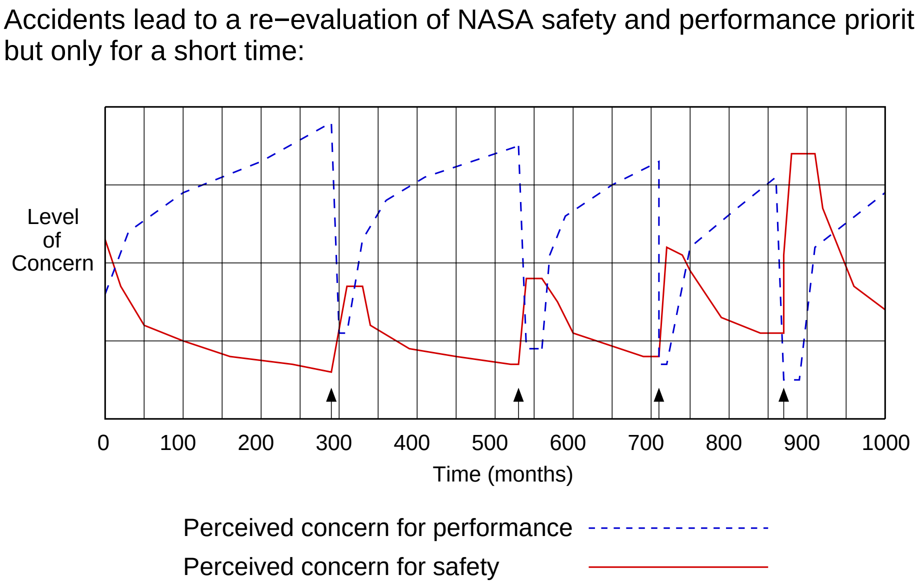

Figure 7: Relative level of concern between safety and performance.

an accident or mishap. Instead, we are doing risk *analysis*, i.e., trying to understand the static causal structure and dynamic behavior of risk or, in other words, identifying what technical and organizational factors contribute to the level of risk and their relative contribution to the risk level, both at a particular point in time and as the organizational and technical factors change over time.

The first example analysis of the baseline models evaluates the relative level of concern between safety and performance (Figure 7). In a world of fixed resources, decisions are usually made on the perception of relative importance in achieving overall (mission) goals. Immediately after an accident, the perceived importance of safety rises above performance concerns for a short time. But performance quickly becomes the dominant concern.

A second example looks at the fraction of corrective action to fix systemic safety problems over time (Figure 8): Note that after an accident, there is a lot of activity devoted to fixing systemic factors for a short time, but as shown in the previous graph, performance issues quickly dominate over safety efforts and less attention is paid to fixing the safety problems. The length of the period of high safety activity basically corresponds to the return to flight period. As soon as the Shuttle starts to fly again, performance becomes the major concern as shown in the first graph.

The final example examines the overall level of technical risk over time (Figure 9). In the graph, the level of risk decreases only slightly and temporarily after an accident. Over longer periods of time, risk continues to climb due to other risk-increasing factors in the model such as aging and deferred maintenance, fixing symptoms and not root causes, limited safety efforts due to resource allocation to other program aspects, etc.

The analysis described so far simply used the baseline parameters in the integrated model. One of the important uses for our system dynamics models, however, is to determine the effect of changing those parameters. As the last part of our Phase 1 model construction and validation efforts, we ran three scenarios that evaluated the impact of varying some of the model factors.

In the first scenario, we examined the relative impact on level of risk from fixing symptoms only after an accident (e.g., foam shedding or O-ring design) versus fixing systemic factors (Figure

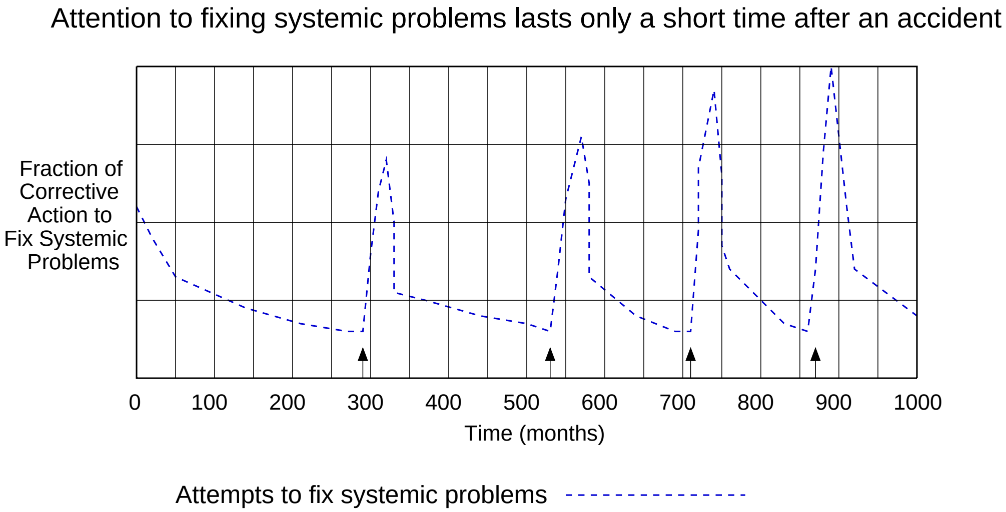

Figure 8: Fraction of corrective action to fix systemic safety problems over time.

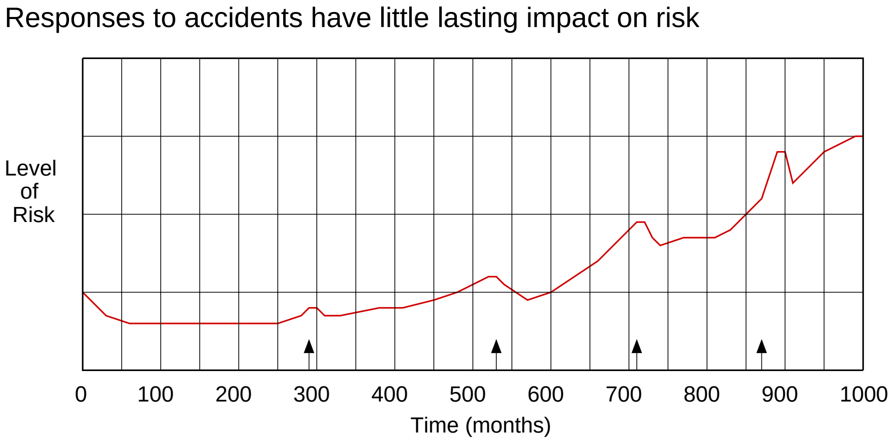

Figure 9: Level of Technical Risk over Time.

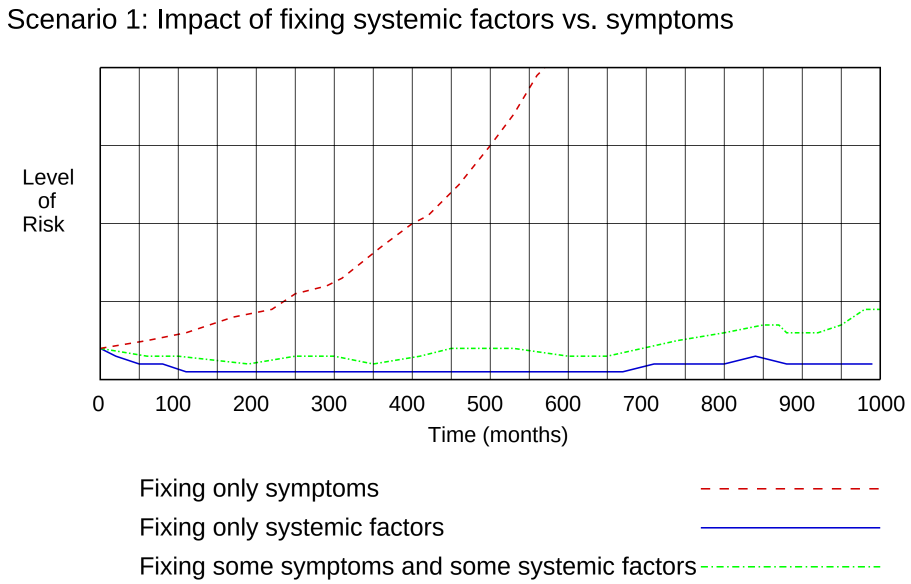

Figure 10: Fixing Symptoms vs. Fixing Systemic Factors

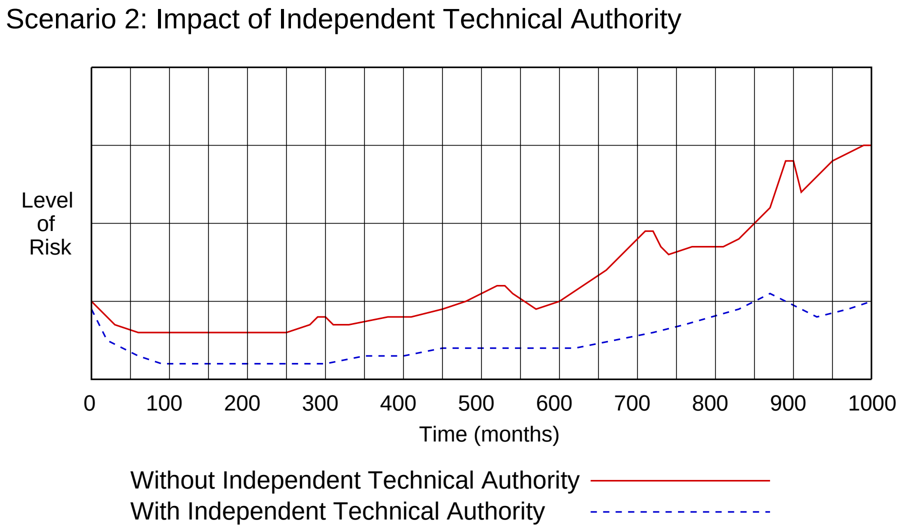

Figure 11: The Impact of Introducing an Independent Technical Authority.

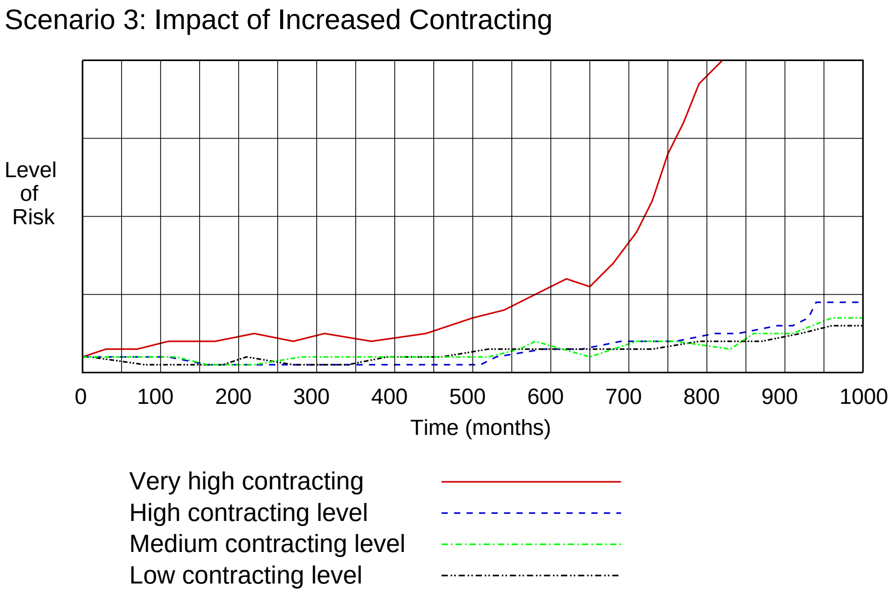

Figure 12: Relative Impact on Risk of Various Levels of Contracting.

10). Risk quickly escalates if symptoms only are fixed and not the systemic factors involved in the accident. In the graph, the combination of fixing systemic factors and symptoms comes out worse than fixing only systemic factors because we assume a fixed amount of resources and therefore in the combined case only partial fixing of symptoms and systemic factors is accomplished.

The second scenario looks at the impact on the model results of increasing the independence of safety decision makers through an organizational change like the Independent Technical Authority (Figure 11). The decreased level of risk arises from our assumptions that the ITA will involve:

- The assignment of high-ranked and highly regarded personnel as safety decision-makers;
- Increased power and authority of the safety decision-makers;
- The ability to report problems and concerns without fear of retribution, leading to an increase in problem reporting and increased investigation of anomalies; and
- An unbiased evaluation of proposed corrective actions that emphasize solutions that address systemic factors.

Note that although the ITA reduces risk, risk still increases over time. This increase occurs due to other factors that tend to increase risk over time such as aging and complacency.

The final scenario we ran during Phase 1 examined the relative effect on risk of various levels of contracting. We found that increased contracting did not significantly change the level of risk until a “tipping point” was reached where NASA was not able to perform the integration and safety oversight that is their responsibility. After that point, risk escalates substantially.

3 Progress and Plans for NASA Engagement

Dr. Stan Fishkind, Chief Engineer of the NASA Space Operations Division (formerly Code M), helped us at the beginning of this research by including us in a trip to KSC and arranging for us to interview key NASA employees.

More recently, this work was briefed at NASA Headquarters, and we have arranged to support the ITA assessment process being performed in the NASA Chief Engineer's office. The assessment involves both a risk and vulnerability analysis of the ITA design. While assisting this NASA activity, we plan to further validate and extend the model where necessary and provide appropriate analyses. More specifically, the ITA assessment group at HQ is currently running focus groups that will rank the various risks and vulnerabilities associated with the ITA program. We will take those identified risks and vulnerabilities and analyze them using our model to better understand their implications and how to avoid them and provide NASA with what we learn. An additional task will be to use our model to develop an appropriate set of metrics to measure the effectiveness of the program and to detect when the identified risks and vulnerabilities are starting to weaken the impact and lead to increased risk. Further tasks will be identified as the assessment process progresses.

4 References

1. William G. Johnson, *MORT Safety Assurance* , New York: Marcel Dekker, 1980.
2. Jacques Leplat, "Occupational accident research and systems approach. in Jens Rasmussen, Keith Duncan, and Jacques Leplat, editors, *New Technology and Human Error*, pages 181–191, John Wiley & Sons, New York, 1987.
3. Nancy Leveson, "A New Accident Model for Engineering Safer Systems," *Safety Science*, 42:4, 2004, pp. 237–270.
4. Nancy Leveson, Mirna Daouk, Nicolas Dulac, and Karen Marais. "Applying STAMP in Accident Analysis," *Workshop on the Investigation and Reporting of Accidents*, Sept. 2003.
5. Karen Marais and Nancy Leveson. "Archetypes for Organizational Safety," *Workshop on the Investigation and Reporting of Accidents*, Sept. 2003.
6. Karen Marais, Nicolas Dulac, and Nancy Leveson. "Beyond Normal Accidents and High Reliability Organizations: The Need for an Alternative Approach to Safety in Complex Systems," *ESD Symposium*, March 2004.
7. Gareth Morgan, *Images of Organization*, Sage Publications, 1986.
8. Charles Perrow, *Normal Accidents: Living with High-Risk Technology*, Basic Books, 1984.
9. Peter M. Senge, *The Fifth Discipline: The Art and Practice of the Learning Organization*, Doubleday Currency, New York, 1990.
10. John Sterman, *Business Dynamics: Systems Thinking and Modeling for a Complex World*, McGraw-Hill, 2000.

5 Appendix: System Dynamics Models

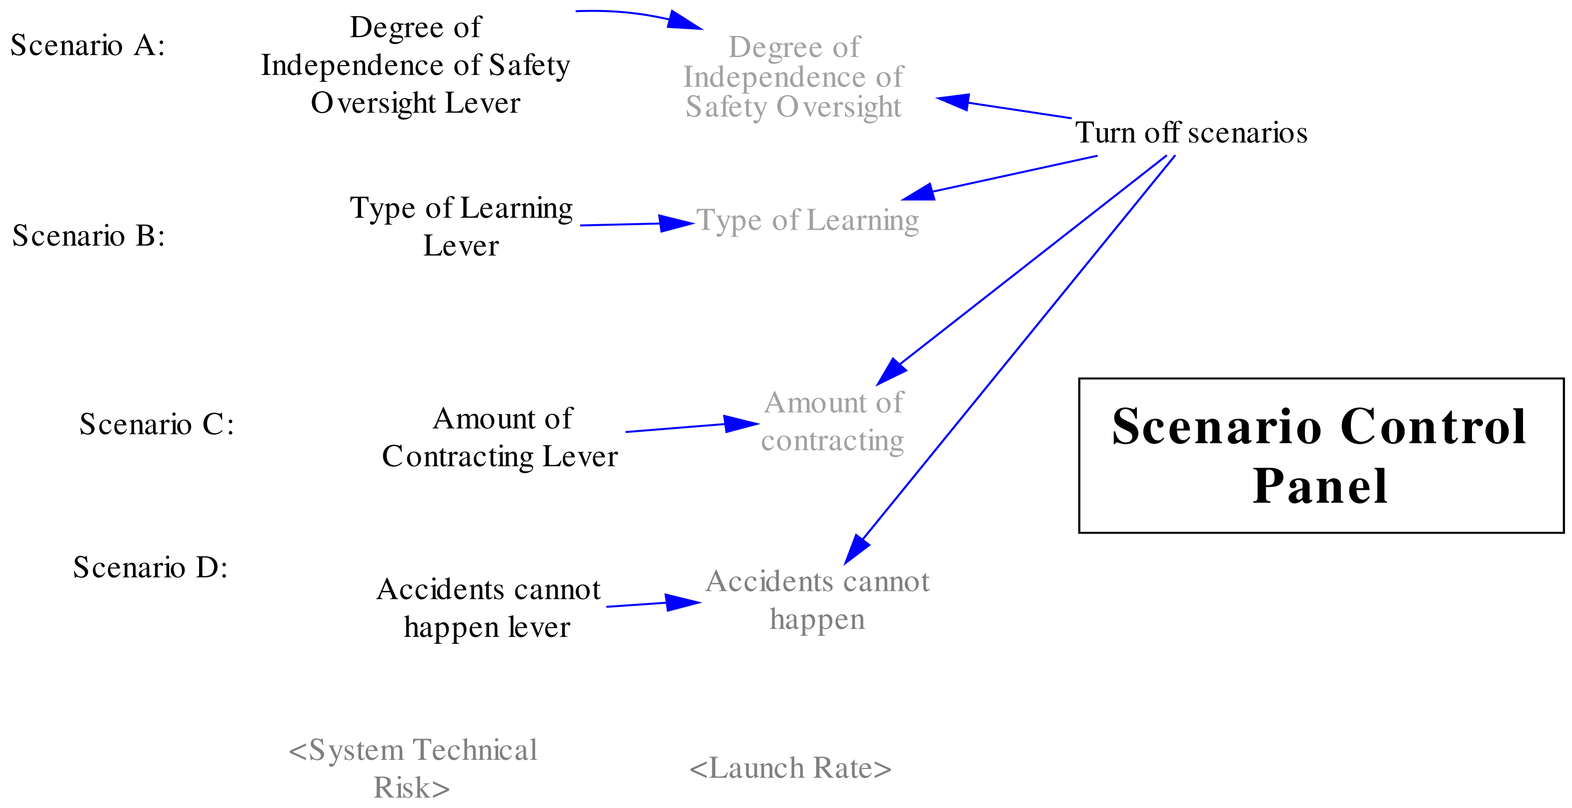
Scenario Control Panel
Scenario A:
Degree of Independence of Safety Oversight Lever
Degree of Independence of Safety Oversight
Turn off scenarios
Scenario B:
Type of Learning Lever
Type of Learning
Scenario C:
Amount of Contracting Lever
Amount of contracting
Scenario D:
Accidents cannot happen lever
Accidents cannot happen
<System Technical Risk>
<Launch Rate>

Integration Control Panel

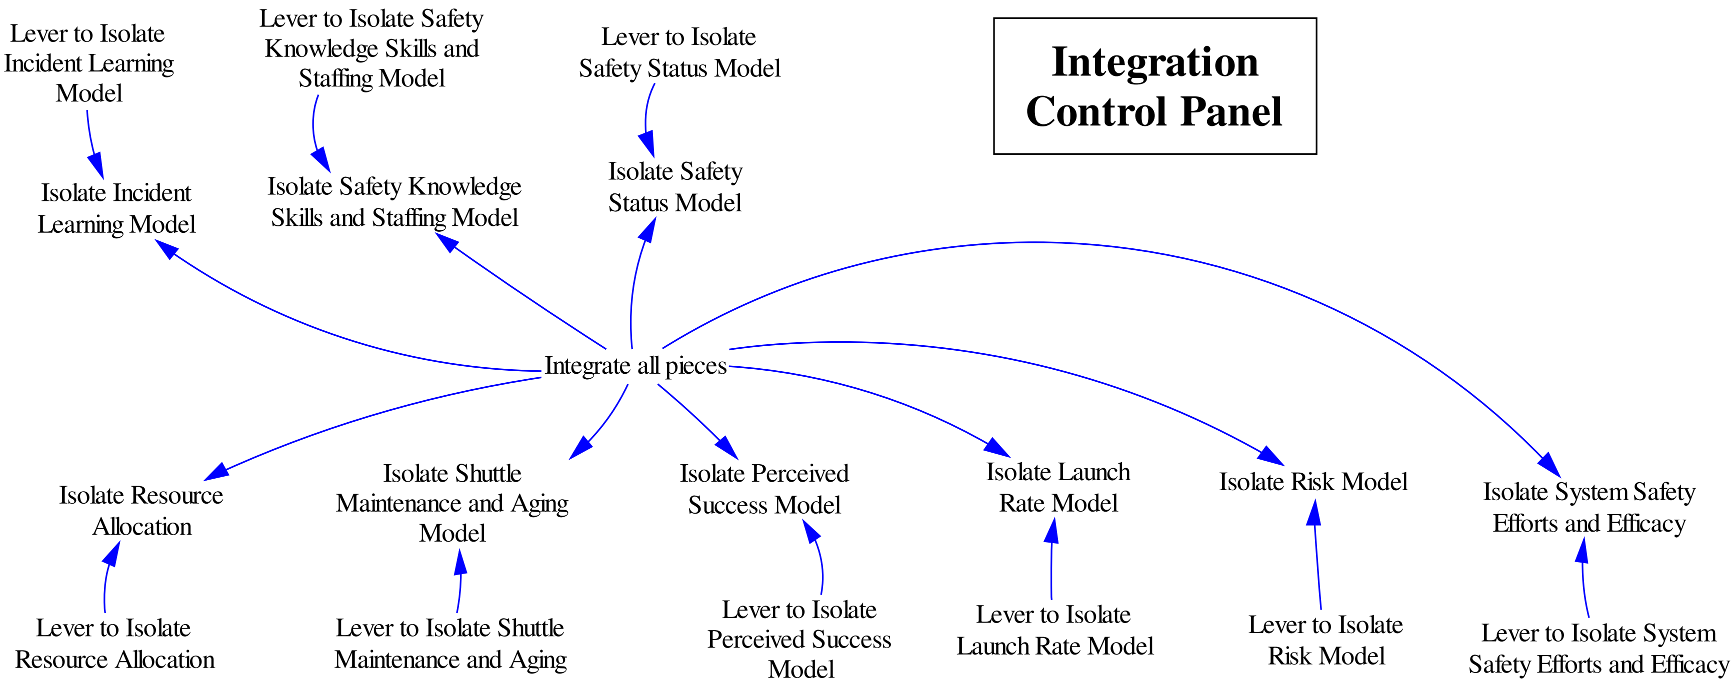

Resource Allocation Model

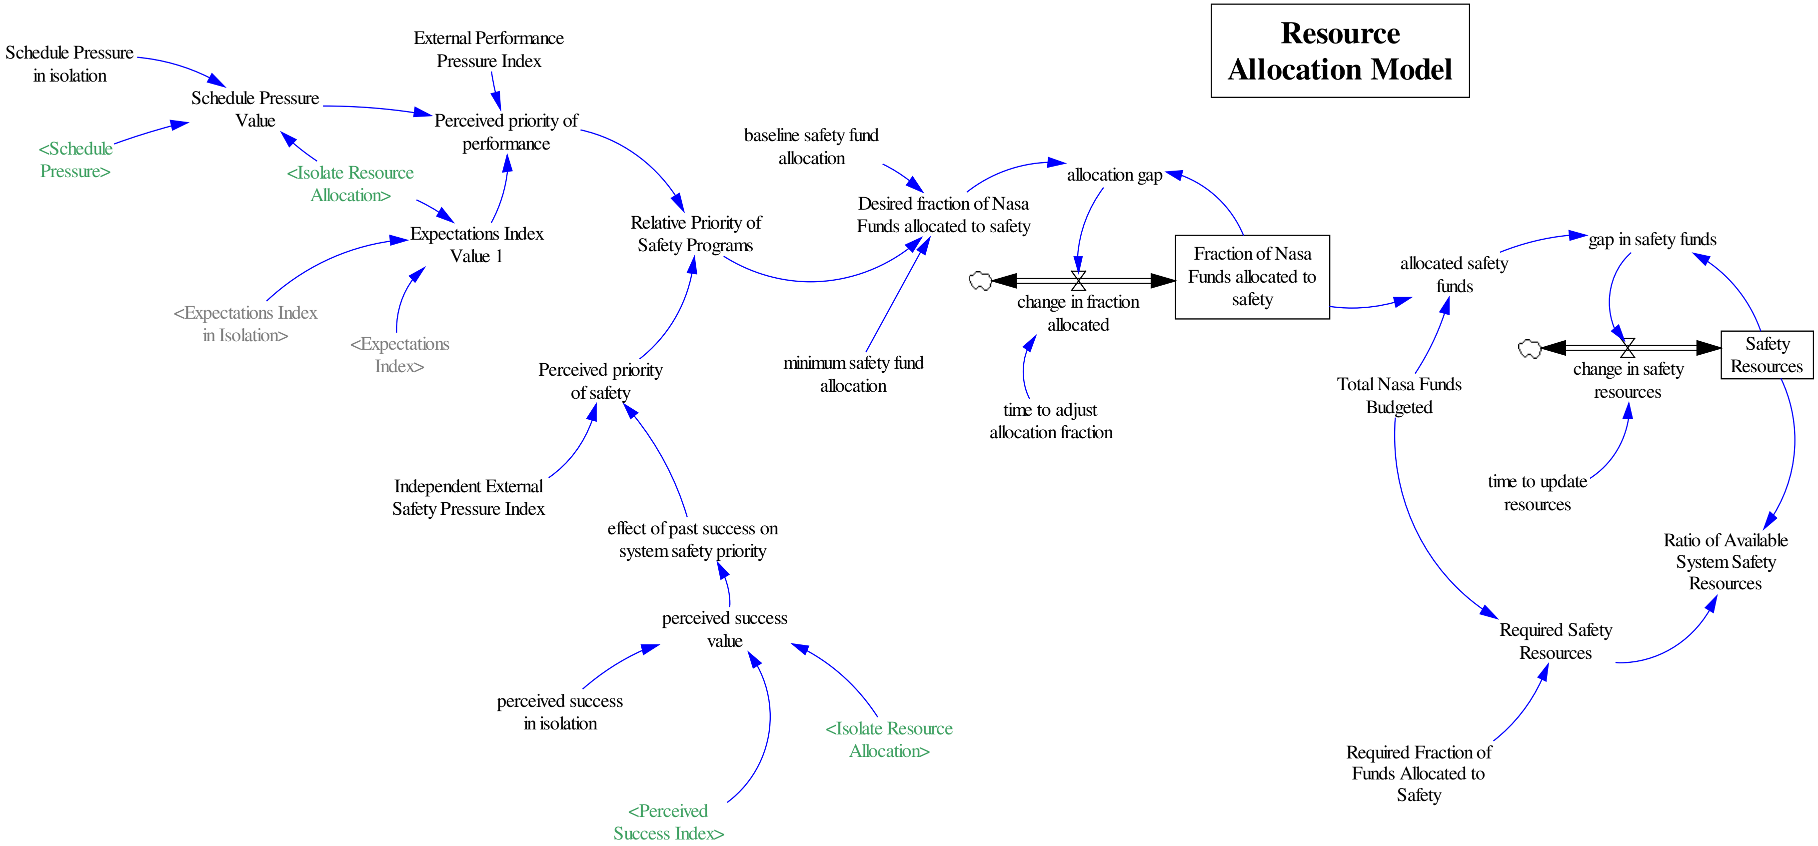

Shuttle Aging and Maintenance Model

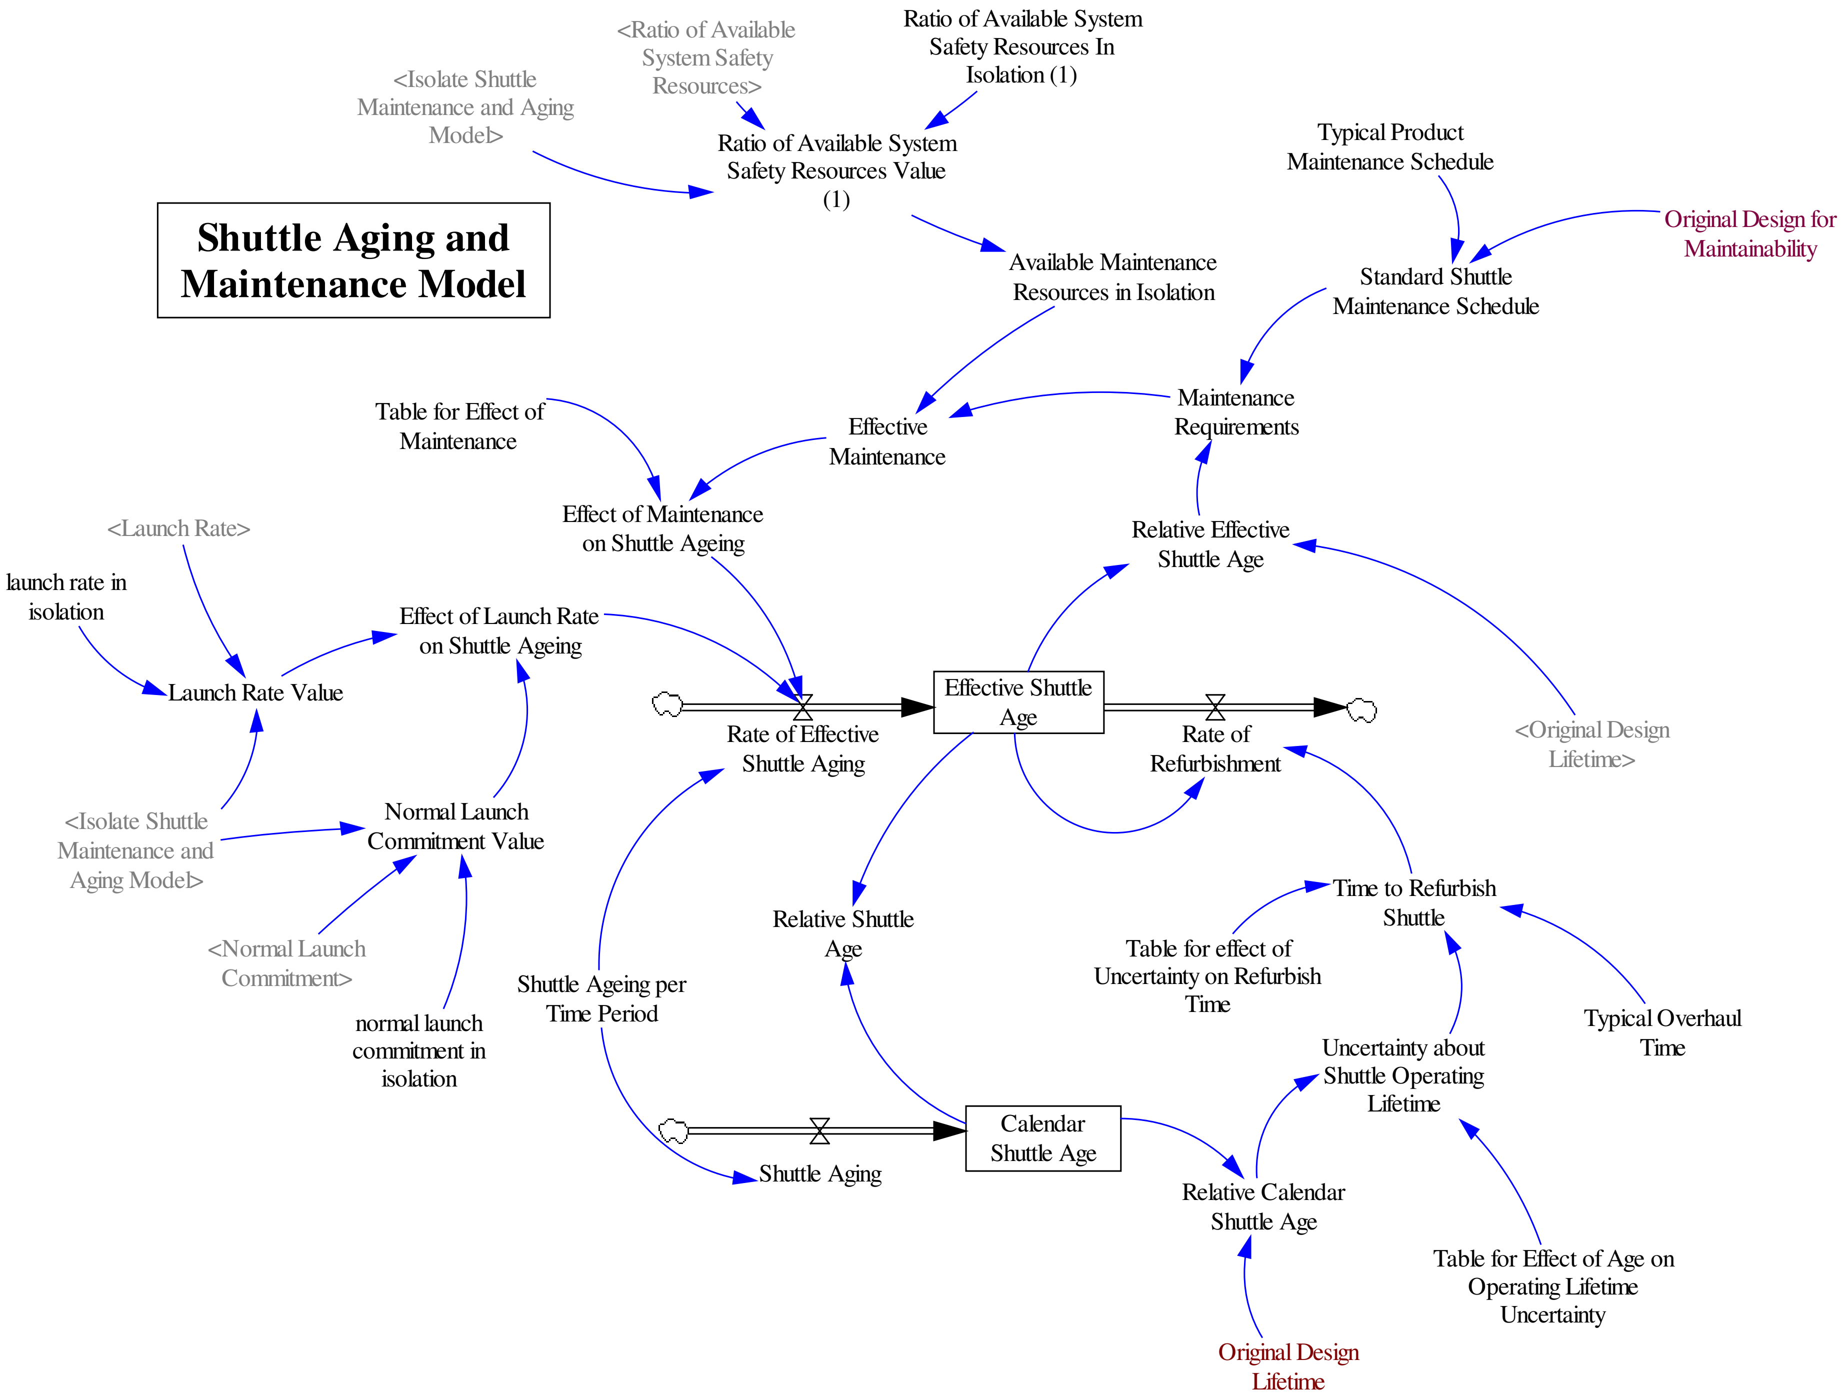

Perceived Success by High-Level Management

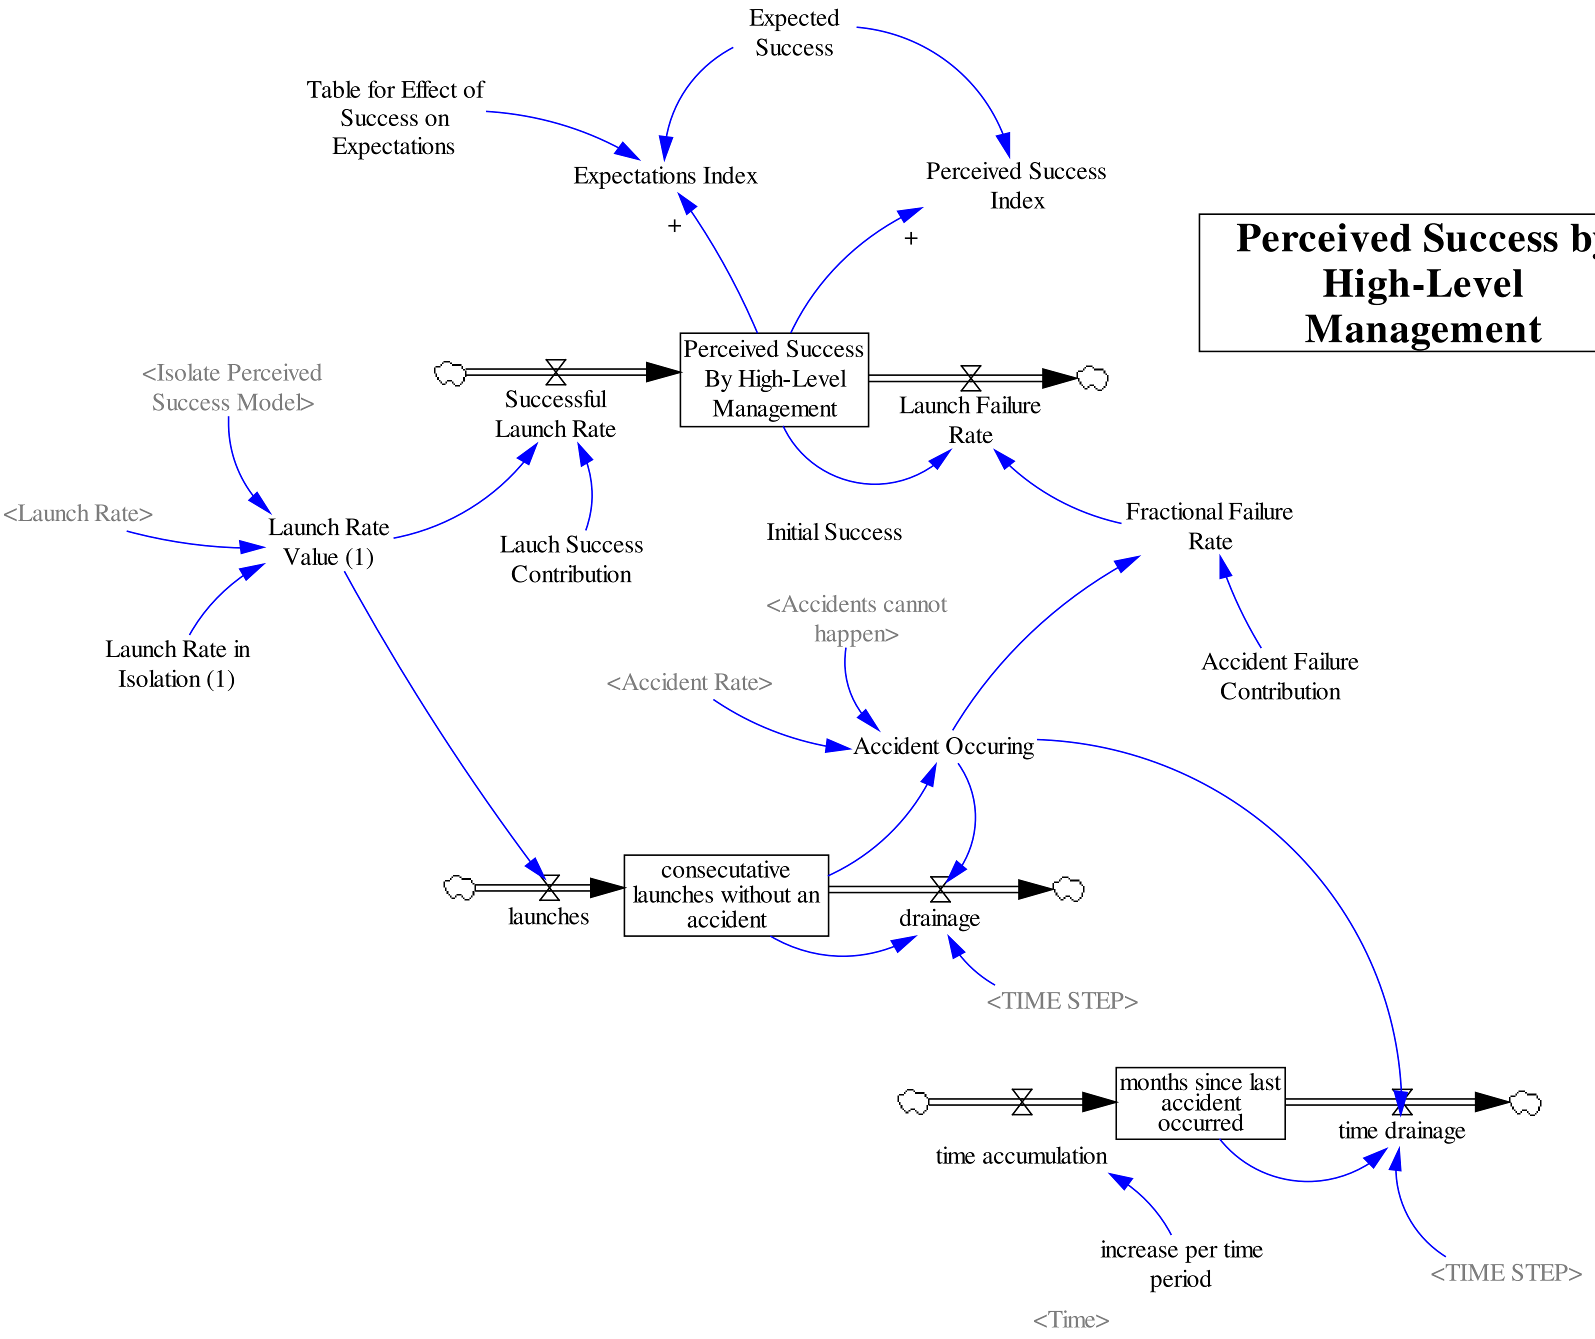

Launch Rate Model

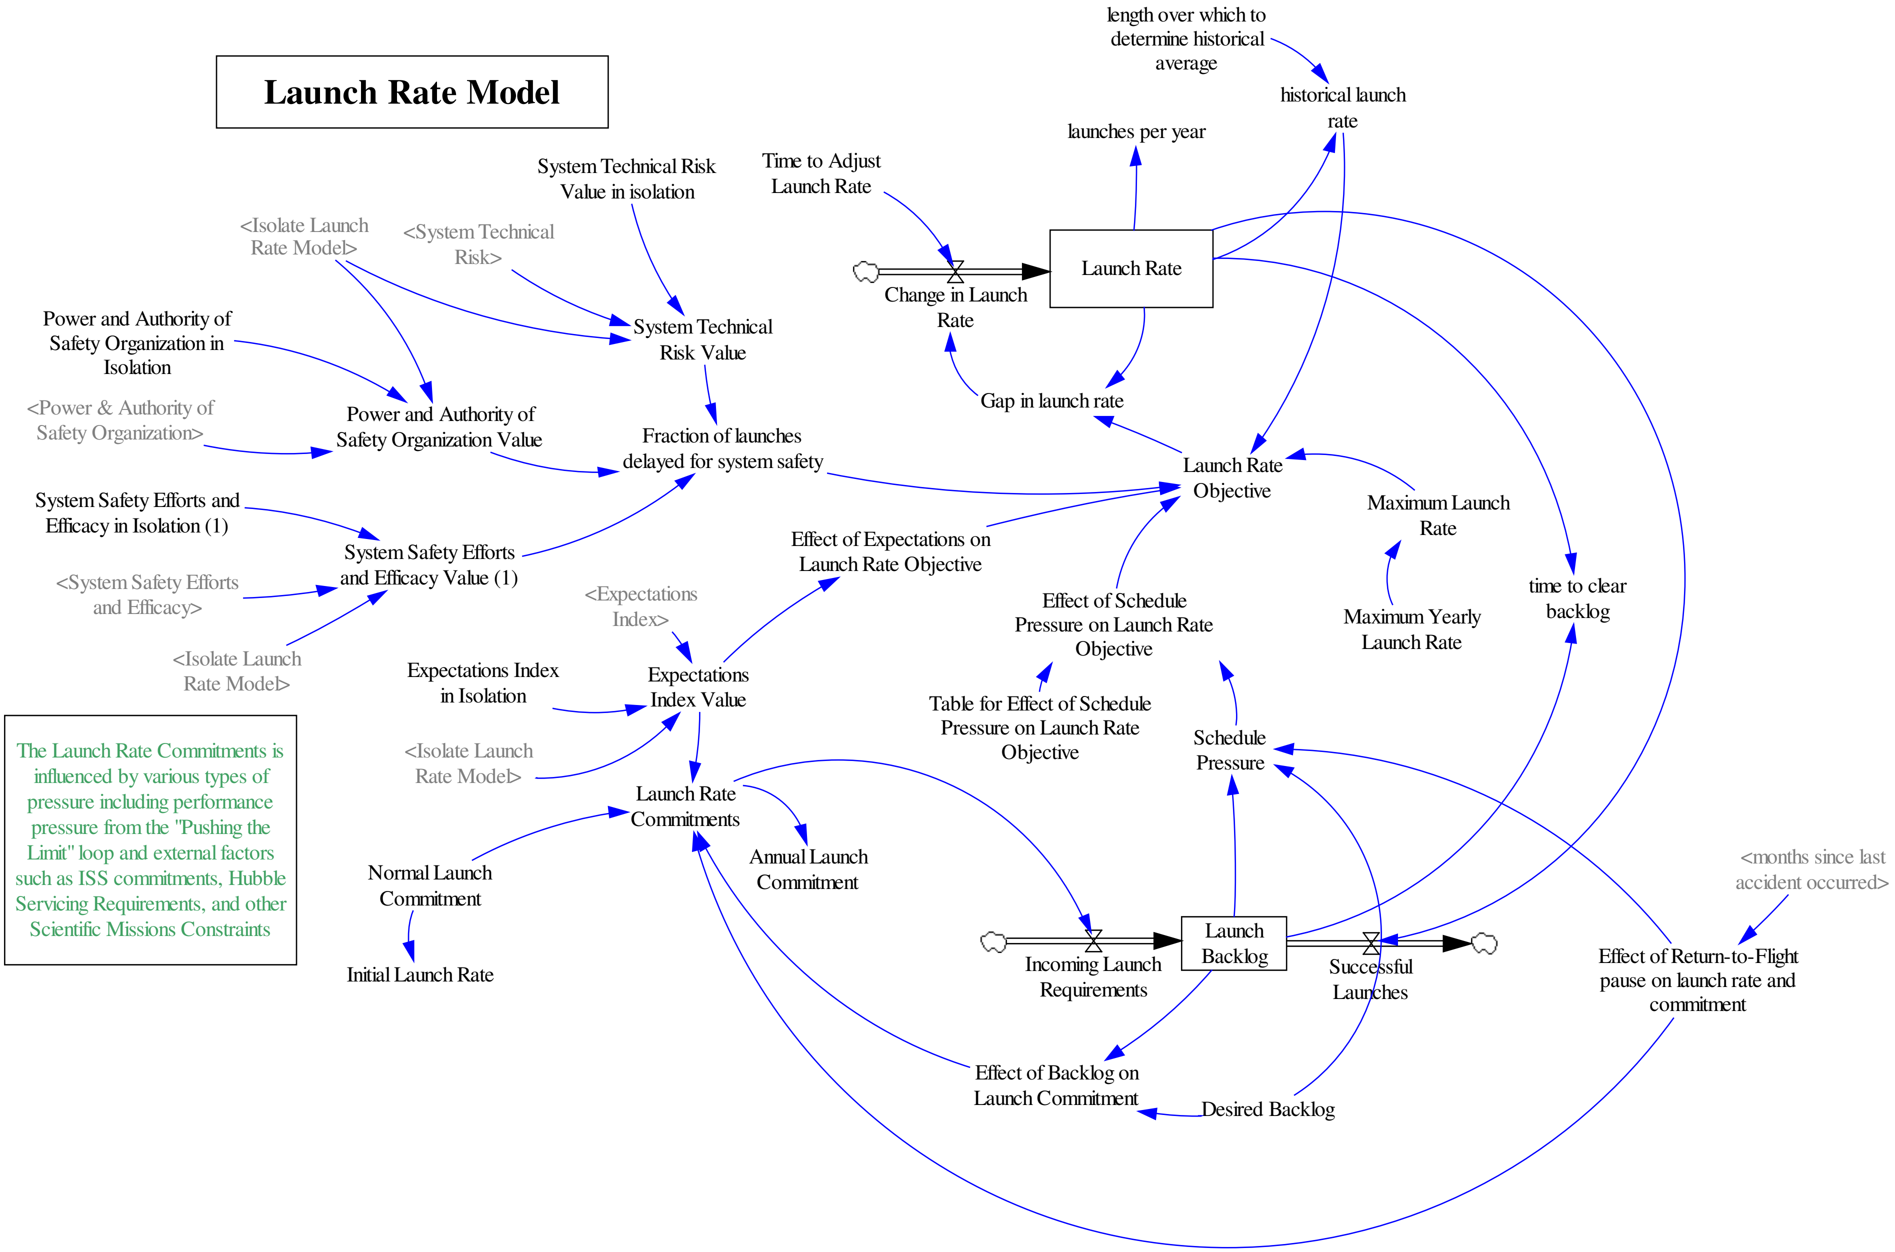

The Launch Rate Commitments is influenced by various types of pressure including performance pressure from the "Pushing the Limit" loop and external factors such as ISS commitments, Hubble Servicing Requirements, and other Scientific Missions Constraints

Risk Model

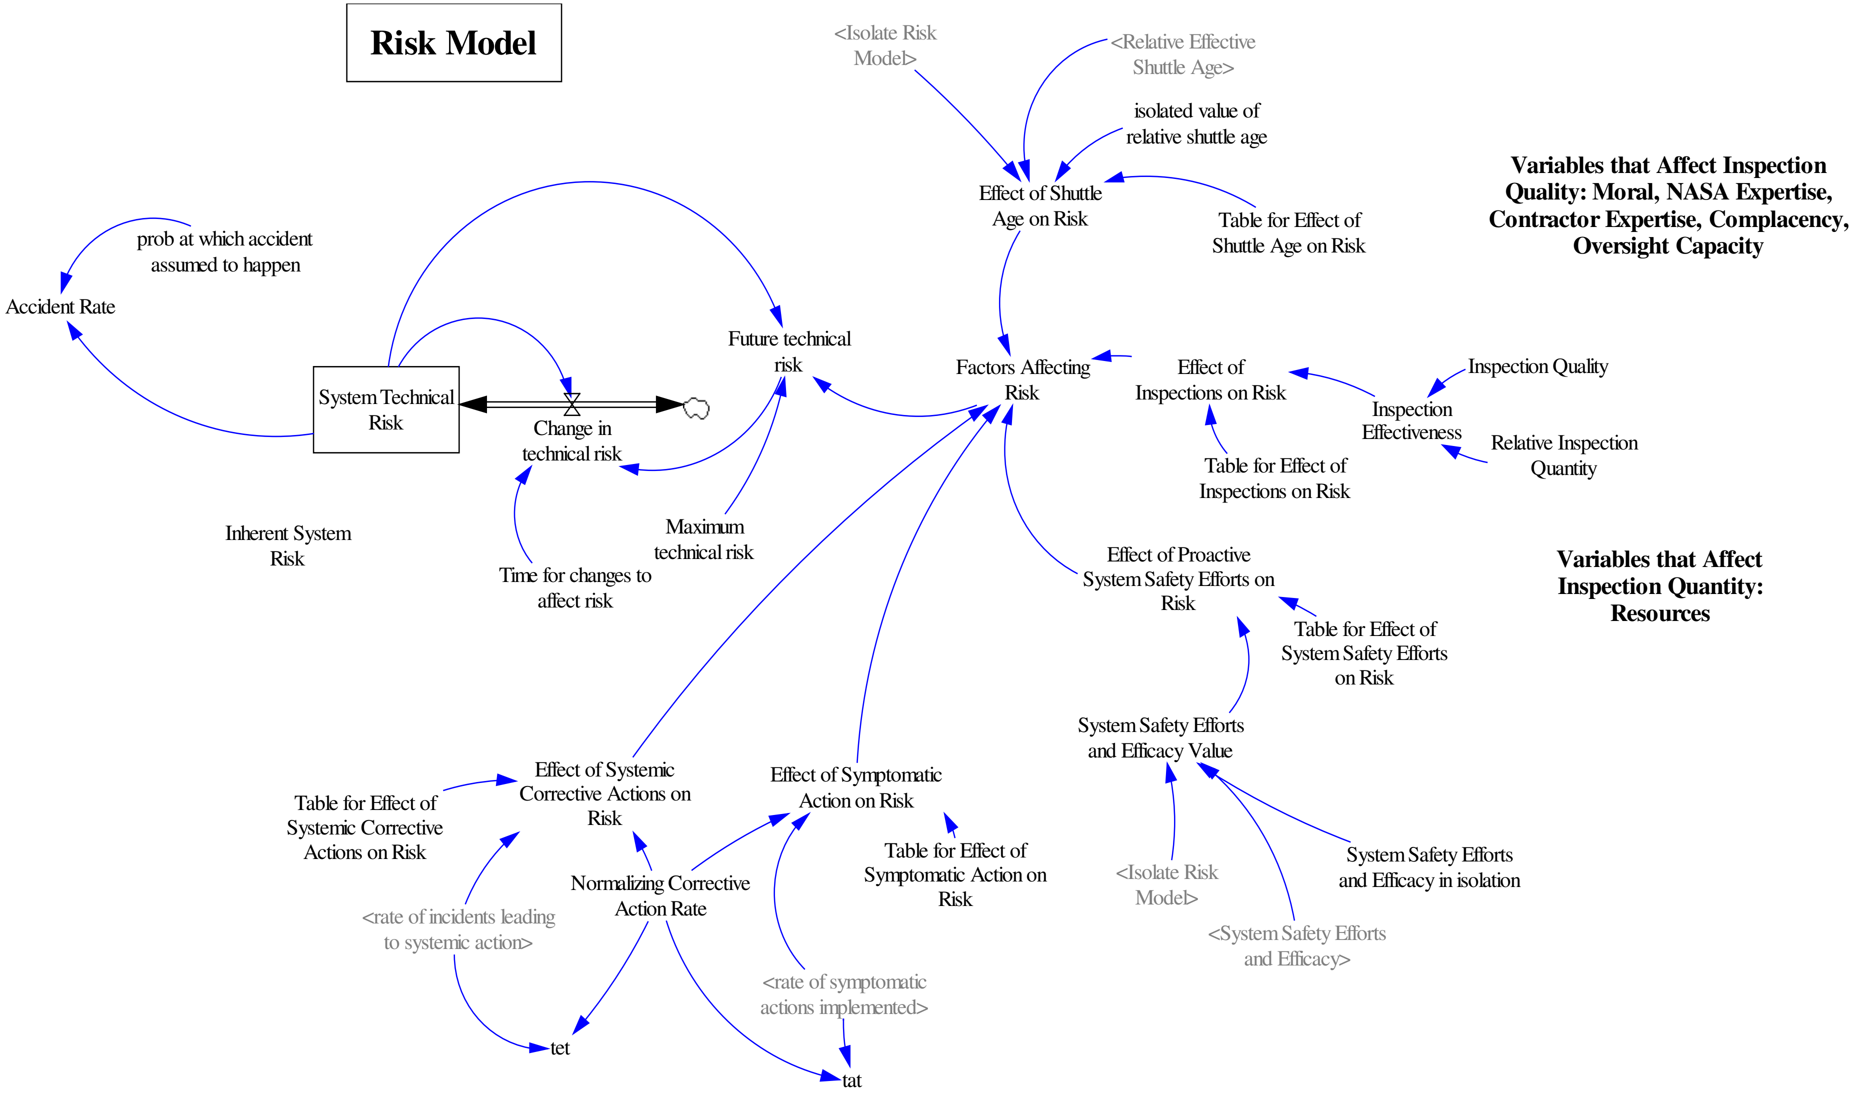

Variables that Affect Inspection Quality: Moral, NASA Expertise, Contractor Expertise, Complacency, Oversight Capacity

Variables that Affect Inspection Quantity: Resources

System Safety Efforts and Efficacy Model

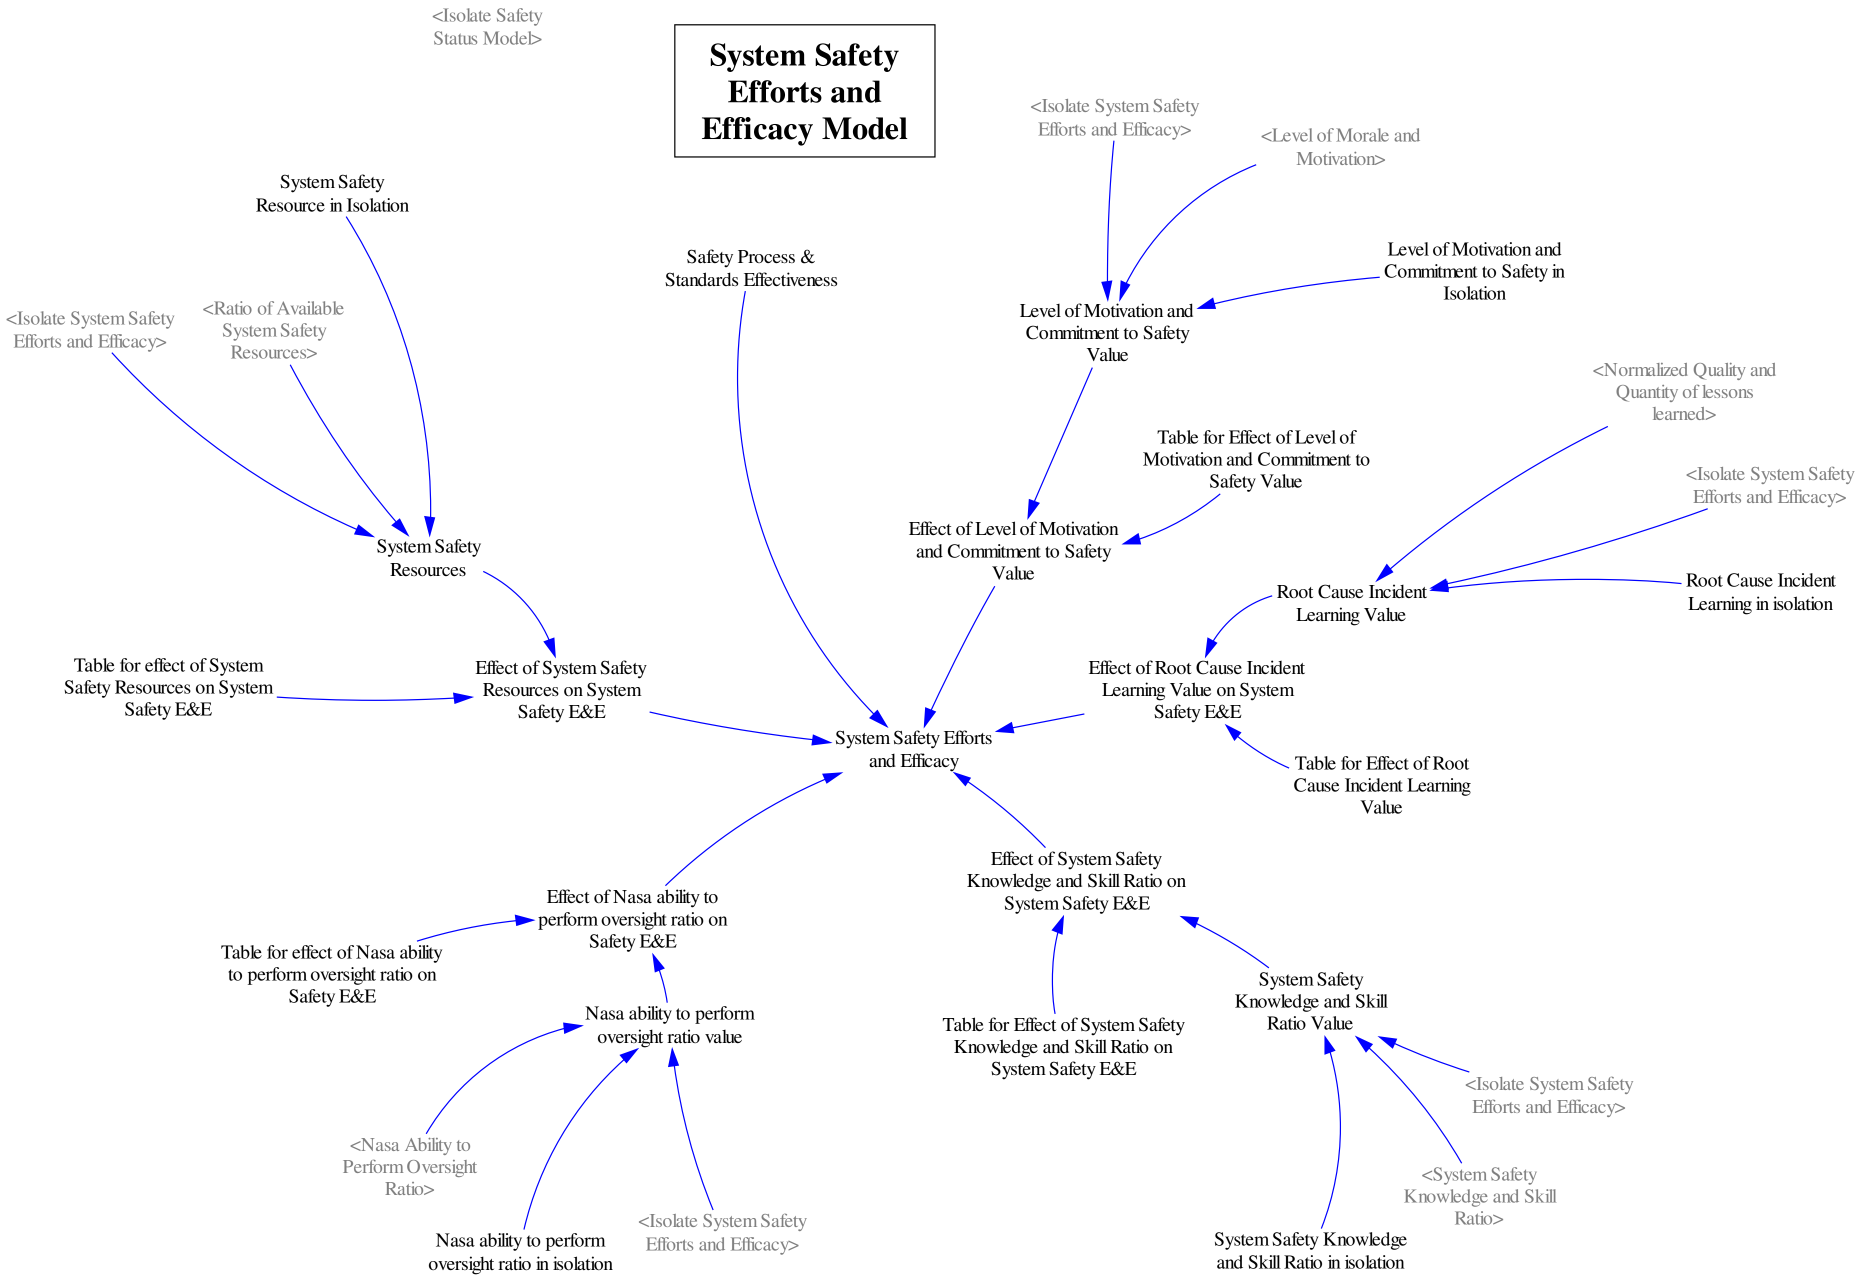

Incident Learning and Corrective Actions Model

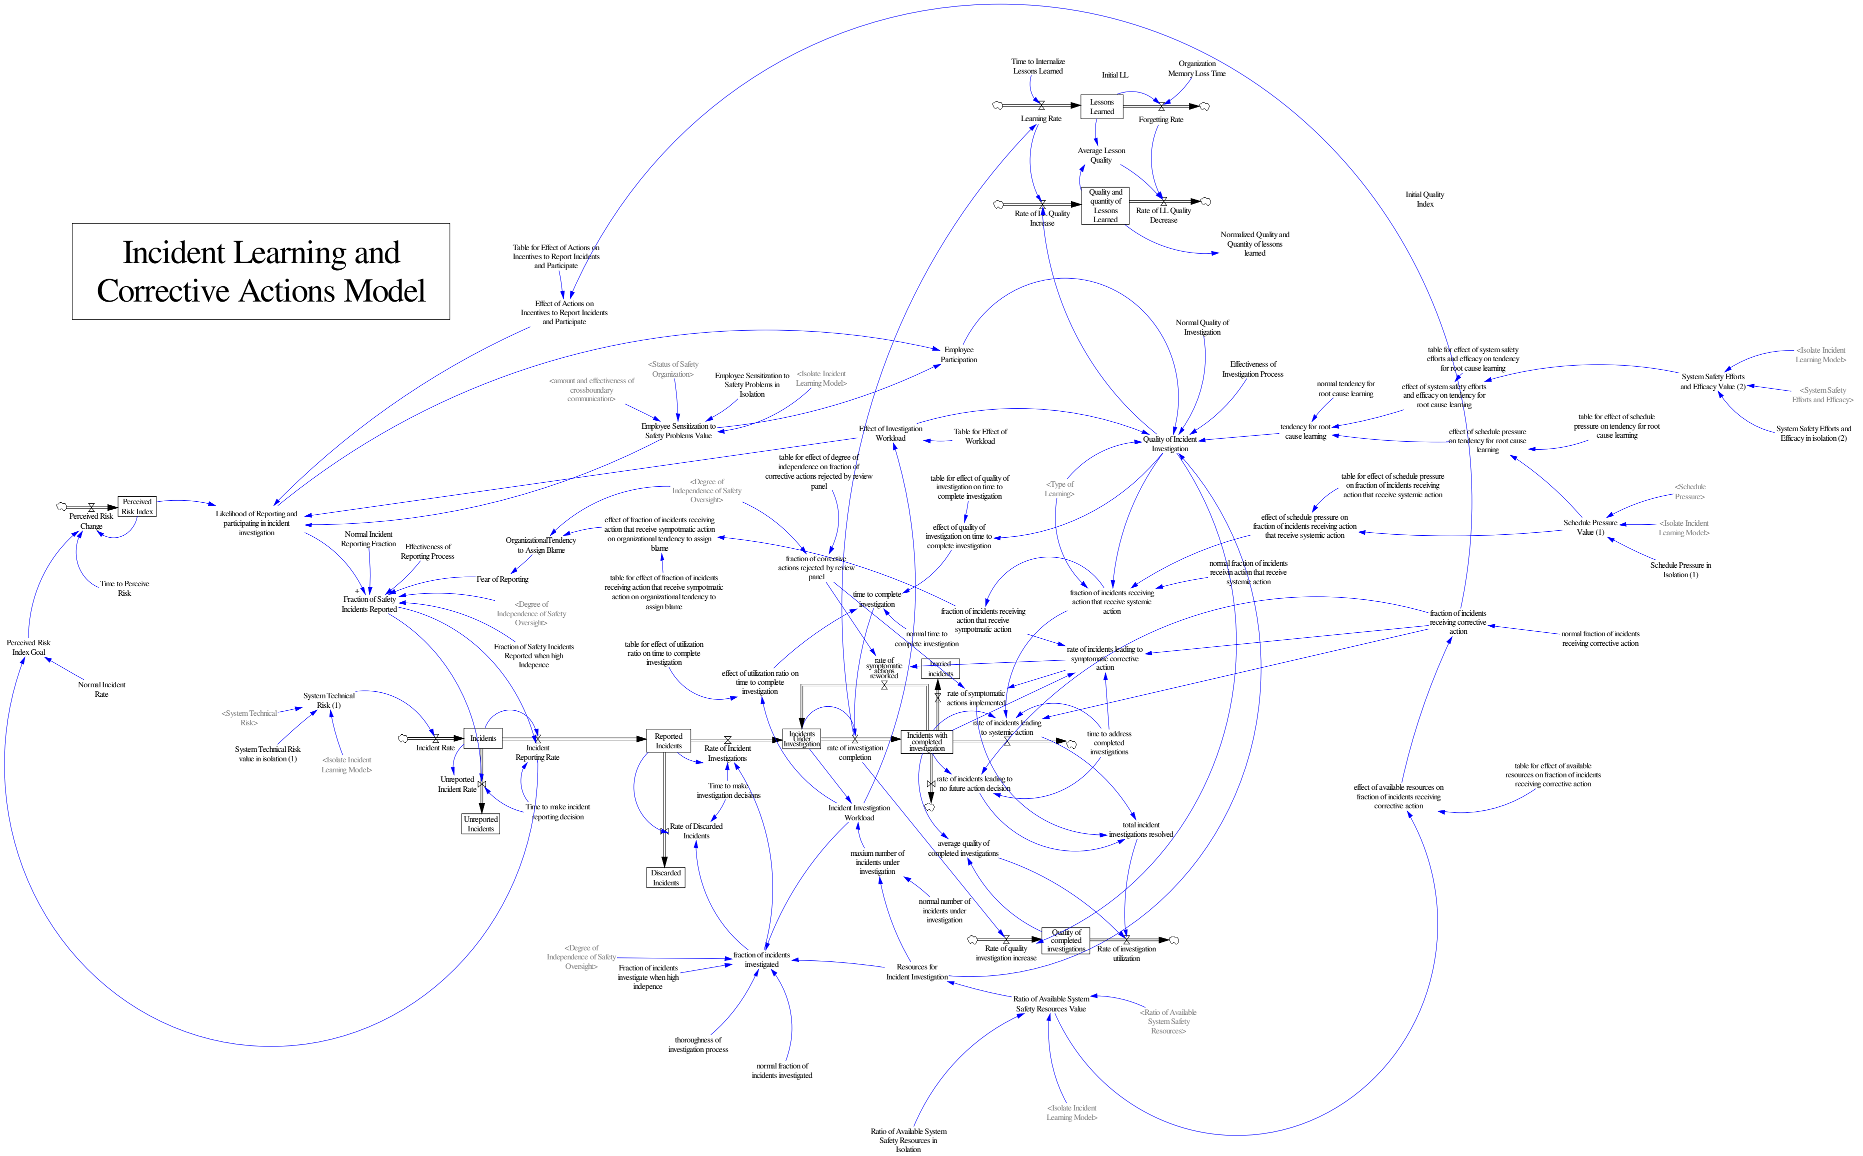

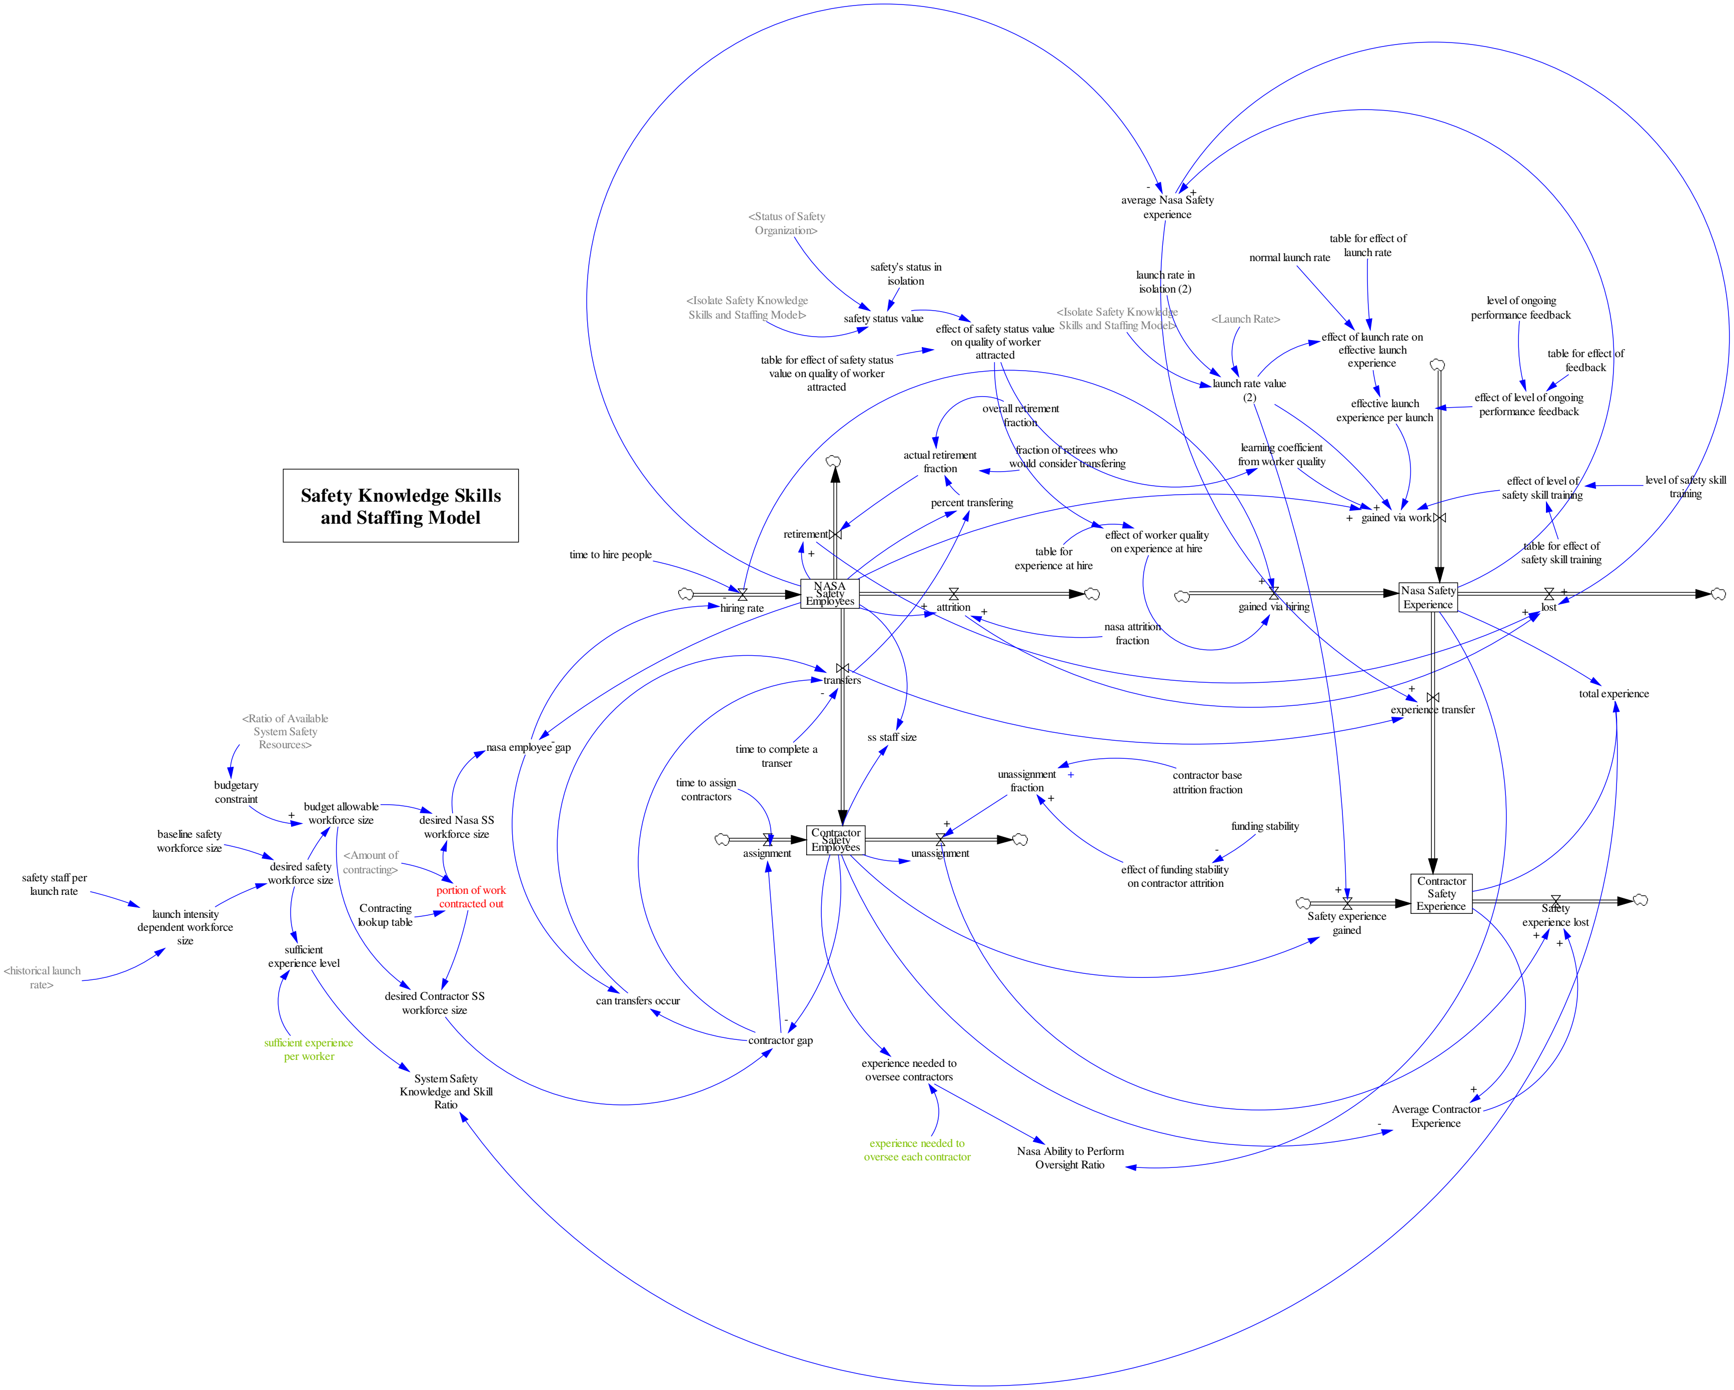
Safety Knowledge Skills and Staffing Model

System Safety Status Model