CAST Handbook

A "Systems Thinking" Approach to the Investigation of Healthcare Adverse Events

Nancy Leveson, MIT

January 28, 2024
(Draft)

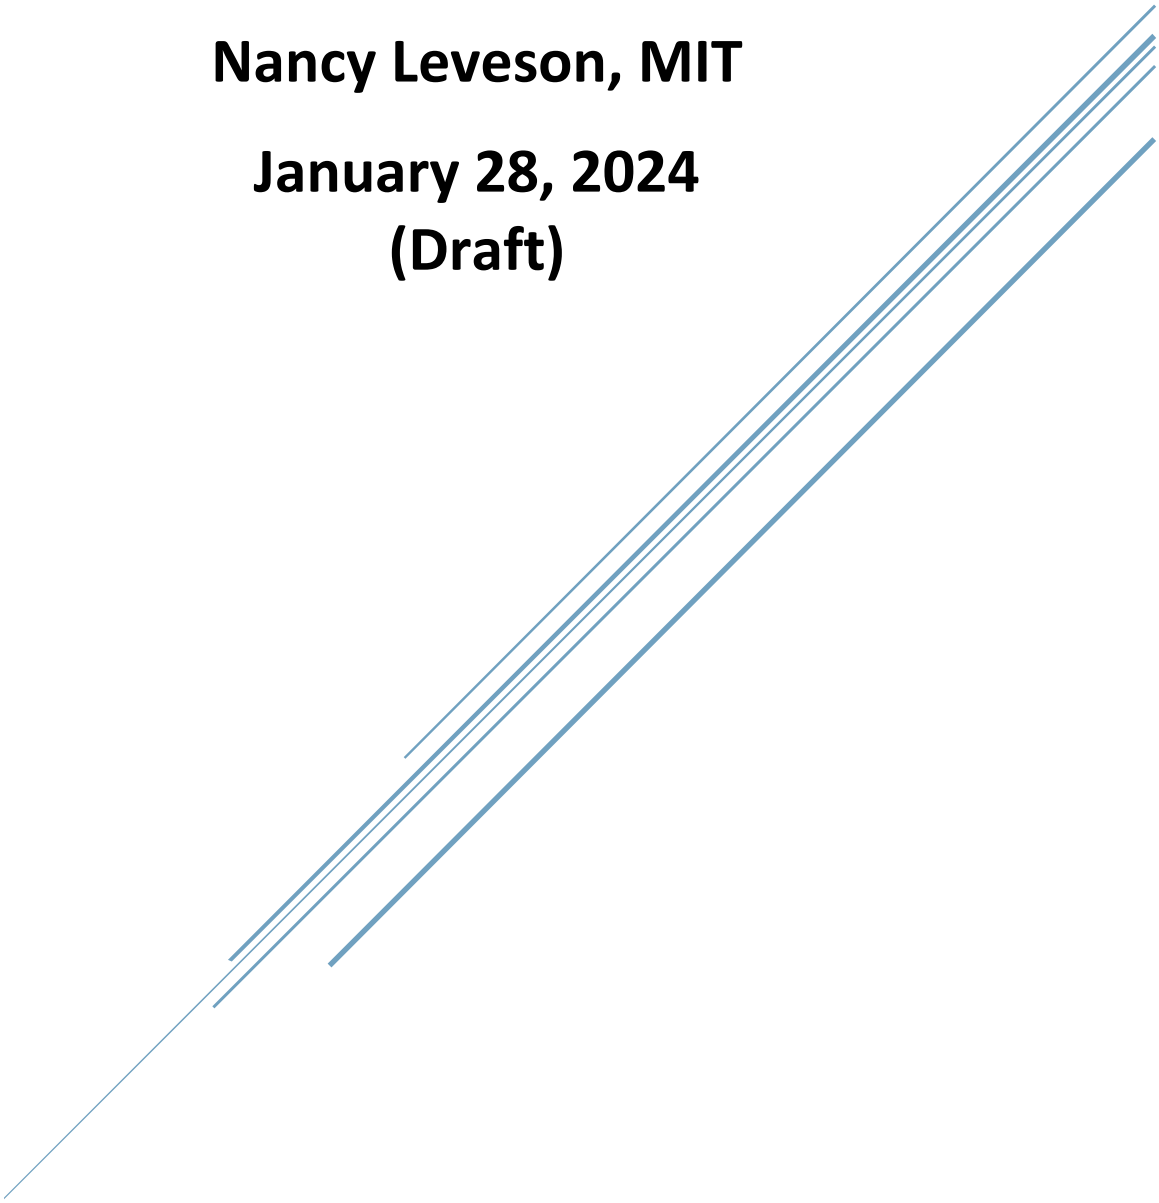

CAST Handbook

Abstract

The use of systems theory as a basis for improving safety in complex systems has been very effective in reducing losses in a wide variety of industries, including aviation. It has the potential to do the same for healthcare safety. While the primary focus should be on preventing adverse events before they occur, we can also learn a great deal about how to prevent adverse events by understanding the cause of those that do happen.

The current approach in healthcare to investigating adverse events relies too much on finding someone or something to blame rather than understanding what features of the system allowed the events to occur. Labeling someone as responsible provides little useful information in reducing adverse events. Instead, the goal should be to identify clear, actionable, and effective measures to redesign the system in order to change behavior and thus to eliminate or reduce future adverse events. This new goal requires a different way of thinking about causality in complex systems and why adverse events occur.

Systems theory provides this different way of thinking about causality. CAST (Causal Analysis based on Systems Theory) is a structured, step-by-step method for identifying the causal factors in the occurrence of an adverse event. CAST has the potential to greatly improve adverse event investigation and causal analysis over the standard RCA methods based on older and simpler causality models. Its efficacy has been evaluated on real systems and real accidents in a large number of industries, including healthcare [1, 8, 9, 19, 20, 22, 24, 31].

This handbook describes how to use CAST, using a real cardiac transplant adverse event as an example.

Contents

Introduction

The IOM report brought attention and a national spotlight on adverse events in medicine and has led to much effort to improve patient safety [13]. An HHS report in May 2014 suggested that in some targeted areas much improvement has been found [10], which suggests that the problem is not totally intractable.

One of the interesting aspects of the 2014 HHS report is that most of the successful interventions cited were attributed to changes in the overall healthcare system structure, such as changing incentives and payment structures and improving information sharing and patient transitions. These interventions are consistent with a systems-theoretic approach.

Industries with very low accident rates, such as commercial aviation and nuclear submarines, use a systems approach to safety [14]. Basically, the systems approach argues that a flawed system, rather than flawed individuals, is responsible for accidents and adverse events. Improving safety, therefore, requires changes in the system design.

All medicine is practiced within a system. A hospital is a dynamic and complex system, interacting as a structured functional unit to achieve its goals (e.g., treating patients). One system may be nested within another; for example, a hospital is nested within a larger healthcare system and an ICU exists inside a hospital. The behavior of a system reflects the linkages and interactions among the components or entities that make up the entire system. The behavior of the components or entities that exist within the system is influenced by the system design and structure. Changing that design and structure can result in improving safety and other system properties by changing the behavior of the system [14].

Aviation is an example of an industry that has used a systems approach to reduce accident rates enormously. Like healthcare, aviation is a complex, adaptive system with many interacting components, including manufacturers of aircraft, airlines, airports, the FAA, air traffic control, maintenance organizations, etc. Each plays a role in maintaining the low accident rates. Safety control is not concentrated only with the manufacturer or the airlines or the pilots, but instead starts at the governmental level with policies, standards, oversight, and accident investigation, and it continues down to the aircraft manufacturers, airlines, and airports, with each level and component in this complex system playing a role and having specific responsibilities for ensuring safety.

Learning from adverse events is only one part of improving patient safety, but it is an essential--component of an overall solution. Some limitations in such learning from past events include (1) wide variability exists in how adverse events are investigated and preventive measures are implemented, (2) investigation of adverse events is done locally and lessons stay local, (3) investigation is often done by clinicians in their “spare time” with no team of safety experts to assist them, and (4) the information in national incident databases, if reported, tends to be superficial at best. Changes in these factors could improve learning from adverse events immensely.

Improving Adverse Event Investigation using Systems Theory

This handbook describes an approach to adverse event investigation and analysis based on systems theory and a systems approach. Root cause analyses in healthcare currently are heterogeneous in methodology and quality with little evidence of efficacy in preventing future incidents [29, 33]. Procedures range from haphazardly investigating accidents to rigorously following protocols for identifying root causes [1, 20, 22, 31].

Almost all of these procedures are based on a notion of causality as a linear sequence of failures, each failure resulting in a subsequent failure. Common analogies are a series of dominos falling and knocking the next one down in order or holes lining up in pieces of Swiss cheese. An example of this type of thinking of causality as a linear sequence of failure events is that a mistake might be made in entering data in the EHR (Electronic Health Record) system, which leads the healthcare professional to provide

incorrect treatment, which leads to patient injury. Blame is then placed on those performing the events in the chain of events, which in turn limits the amount that can be learned from the events. Blame and linear causality lead to ineffective solutions, such as getting rid of the "bad apples," and the number of adverse events usually continues unabated.

Human behavior is always influenced by the design of the system in which it occurs. Understanding why someone behaved the way they did requires examining the larger system in which the person was working. Reducing or eliminating that behavior, in turn, requires changing the system design. Simply replacing one fallible human being with another one inside a system design that leads to human mistakes and flawed decision making is akin to moving deck chairs around on the Titanic.

While a simple linear model of causality can be imposed on any set of events preceding an adverse event, it omits critical information such as the reasons ***why*** the events occurred—which are usually much more complicated than just the existence of a single preceding event. The linear model and investigation methods based on it, therefore, do not provide enough information to most effectively prevent large categories of adverse events. Using the prior example, why were mistakes made in entering EHR data? Was this related to some property of the person entering or retrieving the data? Was it perhaps instead (or in addition) related to the design of the EHR system? Was it related to the procedures for using the EHR system? Was it related to the environment in which the person entering or retrieving information was working, including stress, time pressures, incentive structures, and so on? Or perhaps it was related to all of these to some extent. Without a deeper understanding of why each event occurred, the information needed to eliminate adverse events will remain out of reach.

CAST (Causal Analysis based on Systems Theory)

CAST is a structured process for obtaining a deep understanding of why the adverse event occurred. That information can then be used to design effective changes to prevent future adverse events. Causality is always complex. CAST provides a set of steps that investigators can follow to answer the questions that need to be answered, and, in fact, to assist in generating the right questions to ask. Simply telling investigators to identify the complex causality usually involved in adverse events is not enough; we need to provide structured methods and techniques for them to apply in order to accomplish that goal. Having good intentions is usually not enough for success; humans need well-designed procedures and processes to assist them in finding answers to difficult questions.

An adverse event investigation can be thought of as a process of generating and asking questions about why the adverse event occurred. CAST is a structured way to generate the questions and organize the answers in order to understand the entire loss process. It enables enhanced learning from the adverse events that do occur.

CAST differs from most such investigation processes in several ways:

- The goal is not to identify who is responsible for the adverse event—a blame perspective—but instead to understand why those involved thought they were behaving in a safe manner. While this "no blame" approach is not original with CAST [4, 5], there are few suggestions for how it can be effectively achieved in an adverse event investigation. CAST provides such a process.
- A more general causal process is assumed than a simple linear chain of causes. Instead, causality is envisioned as a complex interacting process.
- The focus is not placed only on individual human behavior but instead considers the system design as a whole and how it contributed to the events.

CAST does not necessarily require greater resources than traditional RCA processes, but it does identify the causes that can be used to redesign the system to change the behavior leading to the adverse events.

Figure 4 illustrates the CAST process. The first step establishes an understanding of "what" happened. The next three steps investigate "why." The final step provides a way to use the information obtained to improve safety.

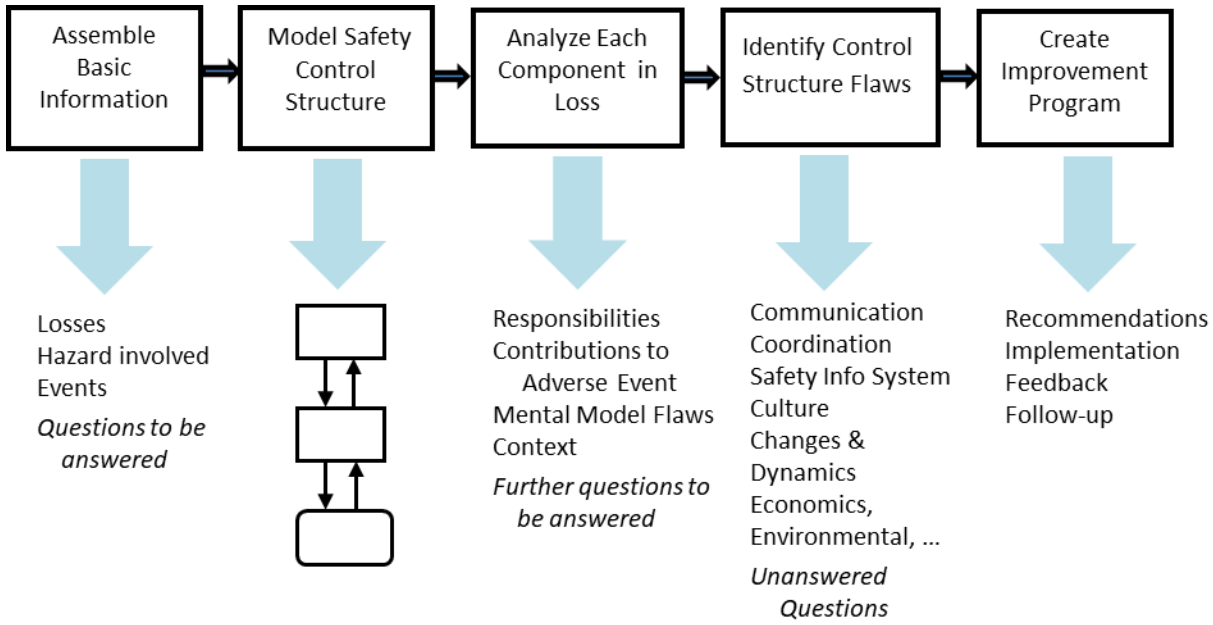

How is CAST performed?

CAST has five steps. These steps are not rigid: At any time, a previous step can be revisited and the results updated or parts of the later steps may be performed, but the overall goal is to use a structured process to produce a final report. Such a process makes it easier to ensure that everything important was examined and analyzed, and that (1) the result is complete (within the boundaries that were defined for the system and the investigation) and (2) the result is effective in identifying recommendations that will improve the system as a whole. In the end, the investigation results should provide an overall organized and thorough explanation of why the events occurred and what can be done to prevent a repetition.

The steps in CAST are explained in this handbook using a real adverse event involving a cardiac patient at a midwestern hospital. Dr. Aubrey Samost performed this CAST analysis during an experimental evaluation of checklists in cardiac surgery [22]. Thirty adverse events not prevented by the use of a surgical checklist were reanalyzed using CAST. The goal was to determine whether a change in the checklist might have prevented these adverse events or whether additional measures are required going beyond checklists. One of these adverse events is used in this handbook as an example of how CAST is performed.

Assemble Basic Information

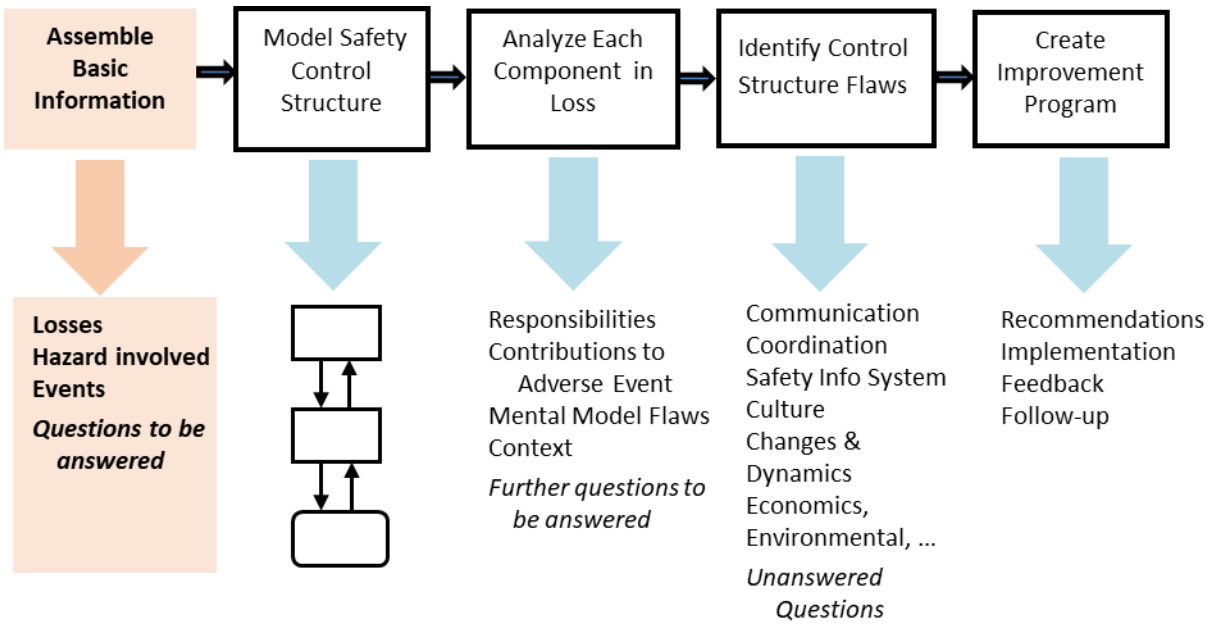

As with any investigation, CAST starts with assembling the basic information about what happened. At this stage, the scope of the investigation starts to be established. This step is necessary no matter what investigation process is involved. CAST emphasizes the activities in this step that will lay the foundation for the following steps, such as the basic questions to be answered in the investigation process.

Step 1a. Identify the loss and the hazard involved:

Safety engineering starts from the concept of a loss, in healthcare called an *adverse event* but in other industries called an accident or mishap. In CAST, the concept of a loss may include any result that is unacceptable to the system stakeholders, such as death or injury to patients or personnel, equipment or property damage, environmental pollution, mission loss, negative business impact (damage to reputation, etc.), loss or disclosure of protected information, financial losses, etc. Note that non-traditional losses beyond human death and injury can be considered, which makes CAST more general than most accident analysis processes.

The goals of the investigation are clarified by first identifying the losses that occurred and prioritizing them if necessary. In the example shown in this handbook, the losses involved include the death of a patient, reputational damage to the healthcare facility, and, perhaps, the psychological injury to the personnel involved (the "second victim" phenomenon [32]) by being incorrectly blamed for the death.

Next, the hazard involved is identified. Safety engineering focuses on hazards. These are states or conditions of the system that, along with a particular set of worst-case environmental conditions, will lead to a loss or adverse event. Two examples of hazards in healthcare are patients receiving less than acceptable standard of care and HIPAA privacy rule violations. Hazards may occur without any subsequent loss. But if there is some set of conditions under which their occurrence could lead to a loss, then they are considered to be hazards. Patients may receive the wrong medication without any subsequent adverse effects, but if there are conditions under which providing the wrong medication can

lead to patient harm, then it is a hazard. Steps should be taken to prevent hazards from occurring and dealing with them if, despite efforts to eliminate them, they do occur.

For the example considered here, the loss was the death of a patient and the hazard was that the patient did not receive immunosuppressant medication at an acceptable time to prevent organ rejection.

Loss: Death of a patient after a heart transplant

Hazard: Immunosuppressant medication not provided within an acceptable time to prevent organ rejection (less than the acceptable standard of care)

Next, the basic information about the events involved needs to be collected

Step 1b. Collect information about the basic events involved:

The events in this real case are shown in Table 1.

Table 1: Events occurring in the case study

A 56-year old male with a history of heart failure was stable with an implanted left ventricular assist device.
He moved to the top of the UNOS transplant list after experiencing persistent drive line infections.
An organ became available and he was transported to the operating room from the cardiac intensive care unit.
The transplant was completed without complication and the patient returned to the CICU in stable condition.
Several hours later, the patient showed signs of decreased cardiac function.
The patient was treated with immunosuppression for presumptive acute graft rejection.
Despite maximal support for several days, the patient's cardiac function never returned and he died.
Upon review of his chart several weeks later, it was discovered that pre-operative immunosuppression had been ordered but had never been given to the patient, contributing to his death.

After this adverse event, the clinical staff discovered two other previous incidents of missing pre-operative immunosuppression. Because the loss event was not detected until several other similar events had occurred (and thus an investigation was delayed) and the CAST analysis was performed even later, the only information readily available and used in the CAST analysis here was (1) the original adverse event report in the in-house incident reporting system, (2) the cardiac unit's subsequent investigation, and (3) personal observation of workflows and equipment at the medical center cardiac operating rooms.

Despite the fact that the CAST re-analysis of the incident did not use any additional information, the causal factors identified and recommendations were very different than those generated by the original investigation and by the typical RCA process.

Step 1c. Generate an initial set of questions to be answered in the analysis:

The goal of any investigation is to answer questions about why the events occurred. The CAST process provides a structured way to generate these questions, with more detailed questions created as the investigation and analysis proceeds. The first questions generated at this early stage of the process

will be very general and probably no different than the questions investigated using any process. The differences occur in later steps when answering the initial questions leads to generating additional ones.

For this case study, the initial questions that might be generated include:

- *Why was the immunosuppressant not administered before the patient was transported to the operating room?*
- *Why was it not discovered by the surgeons and other operating room personnel before the surgery that pre-operative immunosuppressant had not been administered (and then potentially administered before the transplant commenced)?*
- *Why did it take several weeks to discover the missing pre-operative immunosuppressant administration? While it might not have made a difference in this case as the patient was very quickly treated for graft rejection, the answer to the question might have important implications for other events involving medication ordered but not administered in a timely fashion.*
- *Why did it take several instances of this type of adverse event for clinical staff to learn about it and make recommendations to handle it?*

Notice that the focus in these questions is on "why" and not "who." For example, the first question was not (a) "Who did not administer the immunosuppressant" but (b) "Why was it not administered." Answering (b) will result in more actionable and effective recommendations. Answering (a) will simply provide a basis for blame and punishment.

By the end of the CAST-based investigation, these original questions along with all the later questions generated in the CAST process should have been answered. The answers will provide an actionable path forward in terms of preventing similar incidents in the future. Any questions left unanswered will suggest further studies or information that needs to be collected. For example, if an incident involves misinterpretation of information in the EHR, a thorough human factors evaluation of the design of the EHR interface may be appropriate as well as studies to see if similar misinterpretation is occurring widely.

Step 2: Create a Control Structure Model of the System Involved in the Adverse Event

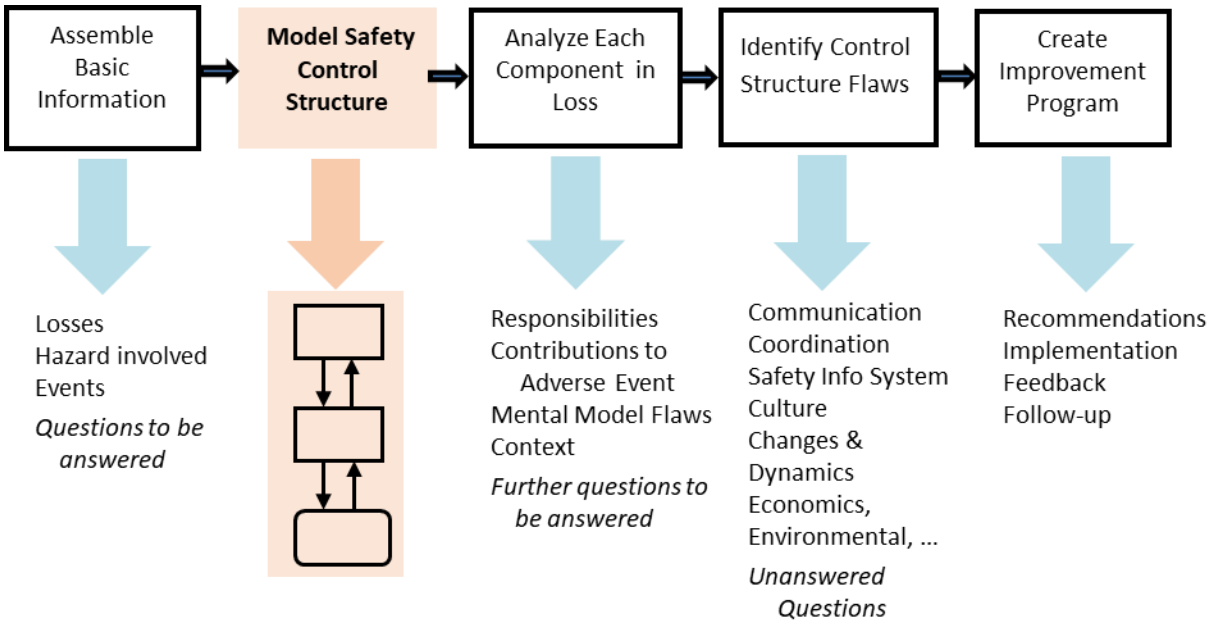

All accident investigation is based on an underlying theoretical causality model. This model determines what we look for as causes of the loss and what recommendations are generated. We may not be aware of using such a model, but we are.

The prevailing model used today in all fields is one that dates back several hundred years and assumes that accidents are caused by failures of system components, where the components include the humans in the system. More specifically, the focus in this model of causality is on identifying individual "failures" (events) by the front-line humans involved such as doctors, nurses, and technicians,[1] namely, what they did "wrong."

In systems theory, or more generally a systems approach, the cause of an adverse event is not a sequence of failure events; instead the cause is a system design that did not prevent those events from occurring. Human error is a symptom not of "bad" humans but of a system design that needs to be changed [16].

Human behavior is *contextual*, that is, it is influenced by the context in which it occurs. Understanding, predicting, or changing behavior requires examining and designing that context. For example, the process of healthcare personnel taking a sample involves not only the ability of the individual involved but also features of their environment such as the equipment available, potential distraction or time stresses, the process for communicating the need to take the sample to the people who have the expertise to do it, and so on.

In CAST, emphasis is not on identifying what people did wrong but on determining why the existing controls and system design contributed to or were not effective in preventing the behavior. Most people want to do a good job. While in hindsight their behavior may appear to involve mistakes, at the time they were trying to do the right thing [14, 16]. Usually, the reasons for their behavior will be perfectly reasonable at the time it occurred. Why did a nurse not administer the medication that had been ordered for the patient? Why did the physician not check that the medication had been administered? When we can answer those questions, we have the information necessary to prevent such mistakes in the future. If, instead, we simply declare them to be negligent, we have learned nothing useful.

To get that information, the goal of adverse event investigation should be to understand *why* it made sense to those involved to act the way they did when the behavior, in hindsight, turned out to be unsafe [4].

Design features or controls to enhance safety may be physical, procedural, or social. Examples include designs that prevent incorrectly connecting physical equipment (e.g., color coding or male-female connectors) or the use of order sets in hospitals to minimize forgetting steps in the orders.

Nurses and technicians, for example, may provide incorrect inputs to a machine because of flaws in the design of the displays or controls. Healthcare personnel may work in an environment that makes it difficult to get the information necessary to effectively treat patients except by breaking the rules and thus incentivizing unsafe behavior, alarms and warnings may be poorly designed, or the only way for personnel to do their job effectively may be to break the rules and use workarounds. Identifying the personnel involved and their behavior as the cause of the adverse event does not provide useful information for preventing adverse events in the future.

In summary, a systems approach treats safety not as a human reliability problem but as a control problem where the system design should prevent or control unsafe behavior. To design effective

[1] It is advantageous to avoid using the term "failure" when talking about causality involving human behavior. The term "failure" is pejorative and implies blame, namely, that the cause has been found and the discussion is ended. The term "human failure' does not appear in this CAST handbook. Humans may do the wrong thing or make errors, but they do not fail (unless their heart stops). The term "failure" usually permeates accident reports and contributes to a blame mindset and the widespread confusion between reliability and safety.

controls, we need to identify the contextual influences that determine or influence unsafe behavior. Then, to change behavior, we change that context.

This is not an "authoritarian" concept of control. The goal of a systems approach is not to reduce human behavior to rule-following. In fact, it is the exact opposite: simply enforcing compliance with rules is a very ineffective way to change behavior. Instead, the goal is to design a system where individual responsibility and competence can effectively help create desired outcomes [14]. Achieving this goal includes designing the system to reduce human errors by providing necessary resources and information in a timely manner, well designed equipment, coordinated activities and decision making, and so on.

To help identify missing or poorly designed controls during an adverse event investigation, CAST uses a model of the system involved. That model is then analyzed to identify the causal factors and the "why" underlying the events.

Modeling system behavior as an adaptive process rather than a chain of events:

Engineers analyze a system for a particular property, such as safety, by building formal models and analyzing those models for the properties of interest. When using a systems approach to safety, the model is a dynamic control model rather than a linear sequence of events (such as holes in Swiss cheese). Consider the model in Figure 1. In this model, the physician is treating a patient (shown by the downward arrow from the physician to the patient).

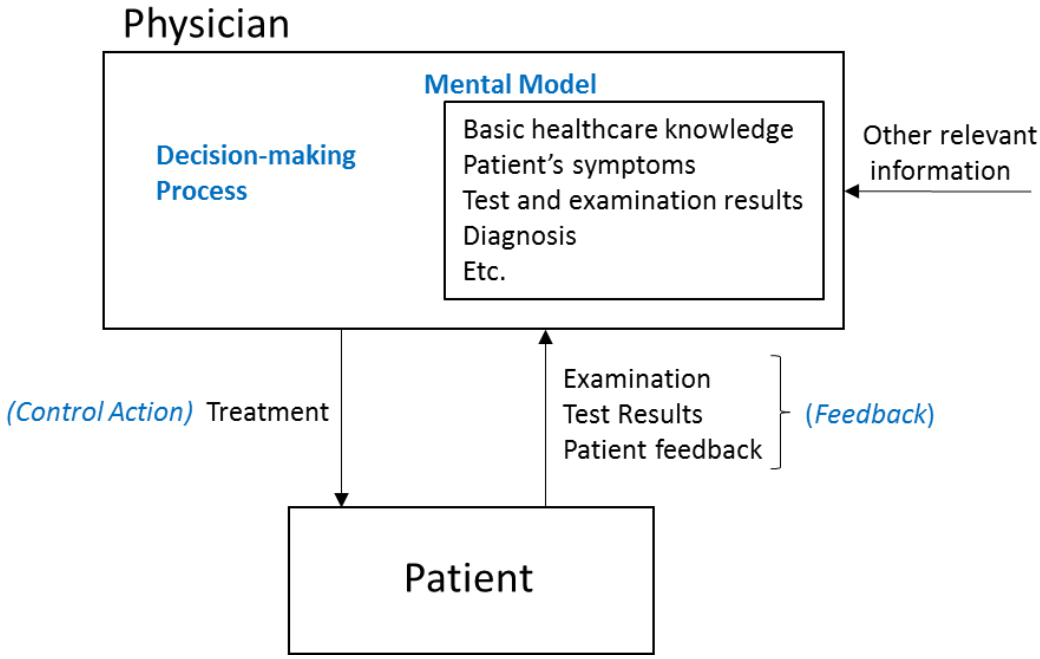

Figure 1: A basic engineering "feedback control" loop.

To make a decision about what treatment to prescribe, physicians use information in their mental models of the situation. Mental models may include such information as basic healthcare knowledge and experience, the patient's reported symptoms, test and examination results, the physician's current diagnosis, etc.

The downward arrow from the Physician to the Patient represents *control* actions, which may include treatment or even not doing anything but waiting to get more information.

After treating a patient, physicians get *feedback* about the effect of the treatment through an examination, test results, patient feedback about their symptoms after the treatment, etc. This information (and perhaps other new and relevant information) is used to update their mental model. The physician then decides whether further treatment is necessary and, if so, what that might entail.

The information in the physician's mental model of the patient's situation may come from sources other than the patient—shown in Figure 1 by the arrow going into the side of the physician box (labeled "Other relevant information"). Other information may come from an EHR system, consultation with other physicians, information about the current environment such as the existence of an epidemic, and so on.

Perception is affected by expectation, that is, we often interpret what we see through the lens of what we expect to see. If we expect to see the results from a diagnostic lab expressed in micrograms, we might not notice that the units reported differ from what we expected, even if the new units are noted somewhere on the lab results report.

Why does ineffective and unsafe behavior occur in this model? One cause may be that the physician's mental model—the model of the state of the controlled process, in this case, the patient's health—becomes inconsistent with the real process (patient) state, perhaps because of missing or incorrect feedback. The incorrect mental model can lead to ineffective and unsafe behavior (control actions). For example, the nurse or physician thinks that an ordered medication has been administered when, in fact, it has not. The controller's decision-making process could also be flawed. A third reason is that the control actions (in this example, the treatment) are not implemented correctly.

This general type of model is called a *feedback-control loop*, where the downward arrow represents the actions provided by the controller (in this case a physician) on the controlled object or process (in this case a patient). In the example in Figure 1, the control action is the "treatment" provided to the patient. Feedback (as well as additional information) is used to update the controller's mental model. The current mental model is used in the decision-making process to identify the next control action. The resulting model is a circular loop rather than just a chain of successive events.

In this example the controlled process is the patient's health, but it may be physical equipment such as an MRI, or a radiation therapy machine, or a pacemaker as well as social and physical processes. Healthcare examples of the latter include the healthcare facility administrators providing rules and procedures for the care of patients and the use of the healthcare facilities. Those are types of control actions. The controllers need not be humans but can be software or, in general, automation.

Note that the control structure model is an adaptive model, where behavior adapts to the current situation as the process unfolds. For example, the physician's behavior will change over time from learning about the response of this particular patient but also more generally by learning from experience and new information. Of course, adaptation or learning do not always imply improvement. For many reasons, adaptation may involve worse behavior or be maladaptive.

The ability to include adaptation and learning over time is an important difference between a systems approach and the traditional linear accident causality model. As stated in the introduction, healthcare, along with most interesting systems, is a complex, adaptive system.

Of course, the events resulting from the control loop process can be strung out on a linear time line. The events are linear because time is linear. But events occurring in a certain order does not imply causality between the sequential events, which is the assumption in linear causality.

In systems theory, the *process* of how the events are generated, i.e., *why* they occur, is the cause of the events. Causality is thus a process rather than a chain of events. While the traditional linear causality is simpler and easier to learn than the more sophisticated modern models of causality, the simple models do not usually provide the information necessary to understand and prevent losses in complex, adaptive systems like healthcare. A new model of accident causality, called STAMP, is based on these

concepts [14]. CAST, in turn, is an accident investigation technique based on the STAMP theoretical accident causality model.

If all of this seems unfamiliar or difficult to grasp, the good news is you do not need to understand the theoretical underpinnings in order to use CAST.

Creating More Complex Models:

The model in Figure 1 is too simple to be very useful in understanding healthcare adverse events. Instead, control loops are put together into *hierarchical control structures* that provide more comprehensive models for identifying the cause of adverse events. Figure 2 shows an example with more system components included. Here the healthcare provider is controlled in some respects by the healthcare facility administrators.

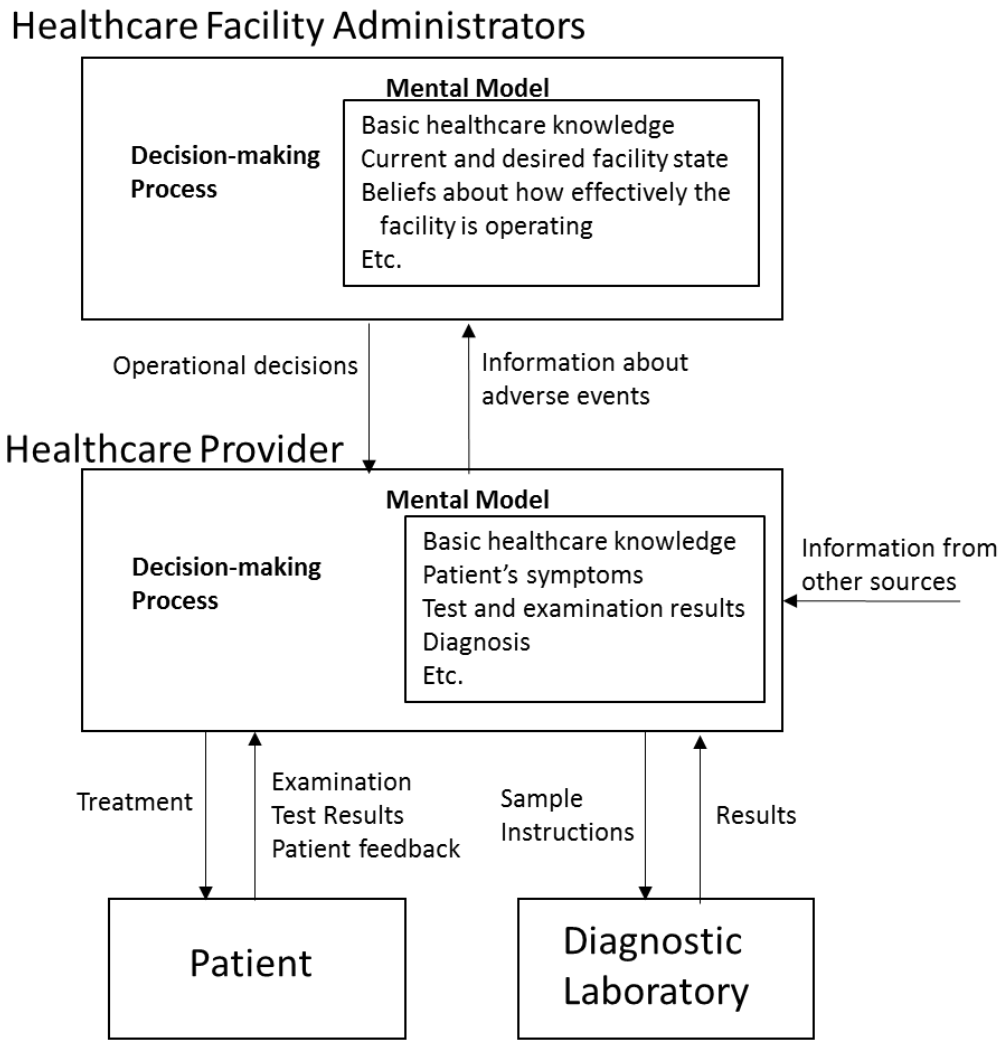

Figure 2: An example of a hierarchical healthcare control structure

An even more complex structure might have multiple controllers at each level. For example, the healthcare providers may obtain basic patient information from the EHR, the design of which is controlled by someone else.

Figure 3 shows the initial safety control structure for the cardiac transplant case study. The model shows the system as it is assumed to work under ideal conditions. It will differ for each hospital, depending on the particular processes used. Adverse events and near misses occur when the control structure (i.e., the designed controls) does not enforce safe overall behavior of the system. In this case study, the cardiac surgeon and the surgical fellow both have responsibility for ordering medications, which could potentially lead to confusion and omission of required actions.

While the groups in this model have many more responsibilities than shown here, only those related to the adverse event are needed to use CAST.

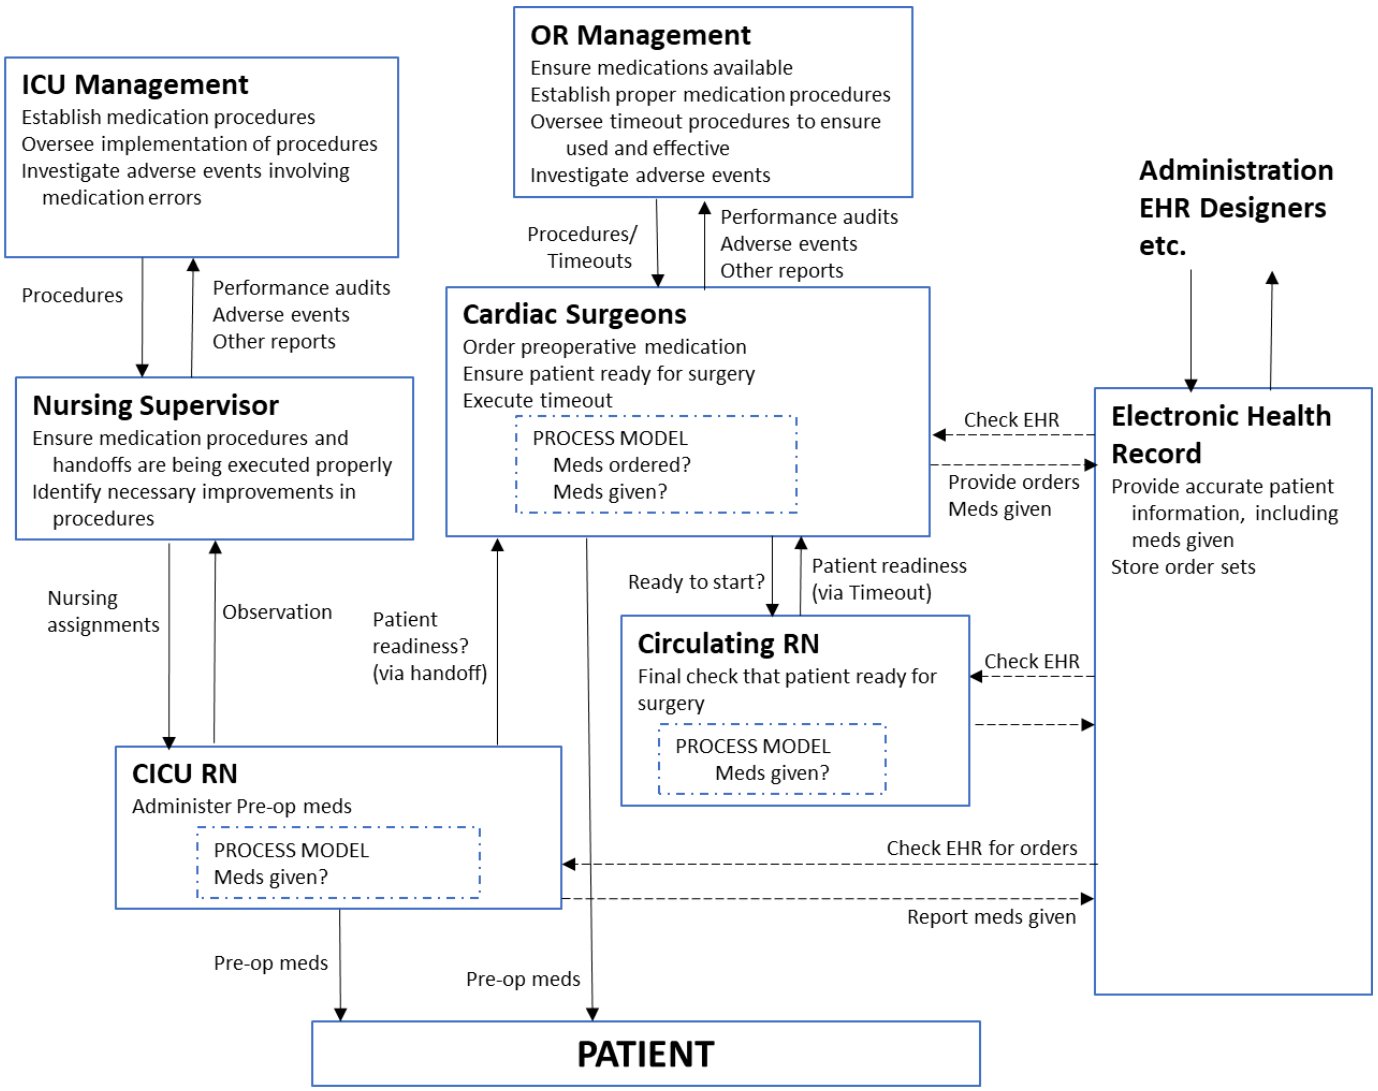

Figure 3. The initial control structure for the case study, adapted from Samost [22]. Additional parts of the system may be included as the investigation continues.

The hierarchical control structure for this incident is limited to the controllers in two clinical areas, the Surgical Intensive Care Unit/Critical Care Unit (SICU/CCU) and the operating room. These are modeled as two separate control hierarchies. Functionally, these two units operate independently of each other in this hospital with separate responsibilities, personnel, policies, and geographic space despite sharing ultimate responsibility for providing safe care to the patient. There may and probably

will be communication between them, however. Adverse events often arise due to flawed communication.

A third and independent hierarchy contains the hospital electronic health record (EHR). The EHR is included as it plays an important role in the behavior of those using it in this adverse event and thus is essential in understanding their actions. More commonly, software acts as a controller itself rather than simply a storage and communication mechanism, for example robotic surgical tools, radiation therapy machines, and pacemakers. In this particular adverse event, however, the EHR was used only to store and communicate information. Because of a lack of information, the control structure above the EHR itself is not detailed here but might have been included in the investigation.

While individual control structures were created for each of the 30 adverse OR events analyzed by Samost, they basically had a similar structure with OR and nursing management at the top and the patient at the bottom. In the middle, typically the physician filled the role of directing treatment goals through placing orders or surgical requests with the nurses carrying out those orders on the patient. A similar or the same control structure can be used for all of the CAST analyses in a hospital with the only differences being the specific responsibilities involved in particular adverse events. Therefore, the time required to create the control structure can be amortized over a large number of uses, including non-CAST uses of the control structure to do proactive analysis and prevent adverse events before they occur [14].

While a control structure is useful in a CAST adverse event investigation, there is no "correct" model. Different models of the control structure can all provide the information sought. In fact, CAST can be performed without creating a control structure model at all, but a model is useful in understanding and documenting the system behavior.

The analysis focuses on two goals:

(1) determining why the individuals involved behaved the way they did and, in fact, thought they were doing the right thing, and

(2) the role of the control structure design as a whole in the adverse event.

The overall goal is to identify the information that will be most important in preventing adverse events in the future. CAST provides a structured method to achieve this goal.

Step 3: Determine why individuals behaved the way they did

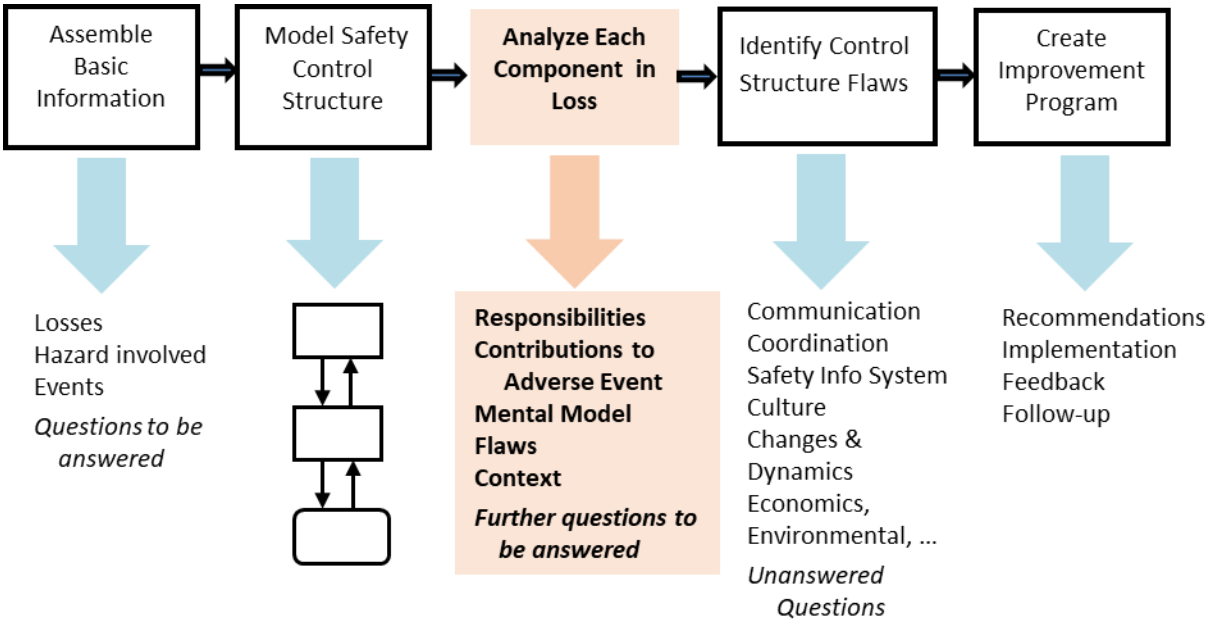

How are the initial components to analyze selected? The events and the initial set of questions to be answered in the analysis provides a starting point. It is not necessary to have a complete control structure to start. Clearly, the whole world cannot and need not be included. As more is learned about what happened and why, decisions can be revisited about what to include and what is less relevant and can be excluded.

For those used to identifying only one or two "guilty" parties, it may be surprising that in most cases when using CAST almost all the controllers in the control structure will be found to have contributed to the events. The solution to preventing a similar event or even a large number of future events is usually to make changes throughout the control structure.

Once again, the starting assumption is that those involved did not intend to contribute to an adverse event. Unless there is evidence that the behavior involved was intentionally malicious (very rare), in the vast majority of cases, the people involved thought they were behaving in a safe way. Even if those involved were grossly negligent, it is important to understand why negligent personnel were hired in the first place or why they continued to be employed in positions requiring conscientious behavior. In most uses of CAST, the controllers involved are found to be trying to do the right thing and have a history of conscientious and safe behavior.

The first step in the analysis, then, is to understand why the individuals involved thought their behavior was safe, that is, would not lead to an adverse event. There has to be some reason that they thought they were doing the correct and safe thing.

The first question to investigate that was identified during Step 1 of CAST is

1. *Why was the immunosuppressant not administered before the patient was transported to the operating room?*

Using the control structure in the previous section, each component with responsibilities related to answering this question is examined to identify:

What happened:
Responsibilities (related to the adverse event)
Contributions to the adverse event:

Why?
Mental model flaws
Contextual factors

Additional questions that need to be answered

The most reasonable place to start is at the bottom of the control structure and then work one's way up to answer the question. In our heart transplant example, the patient's heart was damaged when his immune system attacked it after the transplant. The patient was clearly not involved in any way and is not included in the analysis. If the patient contributed in some way (for example, withheld information or did not follow instructions), then they will need to be included in the analysis.

In this hospital, the CICU nurses are responsible for administering pre-operative medications.

CICU (Cardiac Intensive Care Unit) Nursing Personnel:
Responsibilities (related to adverse event):

- Administer pre-operative medications
- Report any concerns about the patient to the surgical team

Contribution to Adverse Event:

- Did not administer pre-operative immunosuppression
- Did not tell surgical team that the patient had not received the medication.

Stopping here, the nurse's behavior looks very bad. But this is just the start. We need to understand why the unsafe behavior occurred. Behavior is most directly explained by examining any mental model flaws (beliefs) that could have contributed to it and any contextual factors that can help to explain the behavior.

Mental Model Flaws

- The nurse did not think the CICU staff were responsible for administering the immunosuppression.

This, of course, raises an additional question that needs to be answered: *Why did she not believe that administering immunosuppressant was her responsibility?*

Contextual Factors:

- The order set in the EHR does specify that immunosuppression is to be given, but it does not specify who is responsible for carrying out the order.
- Antibiotics are ordered as part of the pre-operative order set but the floor (CICU) nurses do not give them; they are instead given in the OR. This could have[2] caused confusion about who was responsible for giving the immunosuppression. There is no distinction in the orders between pre-operative medications to be given in the CCU versus those administered in the OR.
- During the handoff to the surgical team, the nurse did not know that she was responsible for giving immunosuppressant (she believed it would be given in the OR) so she had no reason to mention it in the handoff.
- This EHR does not send any alert if an order has not been fulfilled.
- New leadership in cardiac surgery was pushing cardiac transplants after several years of doing very few. Nurses were not very familiar with that particular operation.
- This hospital was experiencing nursing shortages and inadequately trained nursing staff was sometimes assigned to the CICU without adequate supervision.

At this point, much has already been learned that can lead to actionable recommendations. But there is more to be learned.

Some new questions have been raised that need to be answered:

[2] The CAST analysis was not performed until long after the AE and its investigation. Therefore, those doing the CAST analysis were unable to talk to those involved about why they did or did not do something. The analysis instead attempts to understand why the people behaved as they did from general knowledge about the hospital and its procedures. In most other situations, the investigators will be able to talk to those involved and get a direct answer at the time of the investigation.

- *Why did the order set not specify who was responsible for giving the immunosuppression?*
- *Why was a new type of operation introduced without ensuring the staff understood their new responsibilities?*
- *There is a handoff from the SICU staff to the operating staff when the patient is transferred to the OR. Should or could the surgical staff have found the medication omission at that time?*
- *What was the background of the nurse involved and was she qualified for her position? If not, why was there no supervision provided?*

As the analysis proceeds, these questions need to be answered. Some will be answered while examining individual components in this step while others are answered in step 4 of CAST, which examines the interactions among the system components—not just individual control loops. Obviously, the questions and answers need to be documented throughout the CAST process, with the final set providing a thorough explanation of the events and the information necessary to generate the most important recommendations for related adverse events in the future,

The next question that might be investigated is:

2. *Why was it not discovered by the surgeons and other operating room personnel before the surgery that pre-operative immunosuppressant had not been administered in the CICU (and then potentially administered in the OR before the transplant commenced)?*

Answering this question requires examining the OR personnel, which consisted of the surgical team and the OR nursing staff (circulating nurse).

Surgical Team (Attending Surgeon and Surgical Fellow or Physician's Assistant):

Responsibilities:

- Order pre-operative antibiotics and immunosuppression
- Ensure that patient is ready for surgery before beginning
- Execute timeout procedure (checklist)

Contributions to Adverse Event:

- Began surgery without patient having received prophylactic medication

Why?

Mental Model Flaws:

- They had ordered the immunosuppression and therefore believed the patient had received it.
- They believed the nurses in the CICU were very familiar with the necessary pre-operative medications as they specialize in cardiac patients.

Contextual Factors:

- On the order screen of the EHR, there is no record of whether an order has been acknowledged and carried out. To see if the orders are completed requires leaving the EHR order screen and going to an entirely separate screen, called the electronic medication administration record (eMAR). The eMAR lists the medications and the time they were given to the patient, but there is nothing in the EHR that clearly shows an order was given and not carried out. One has to be looking specifically for it to pick up on this scenario. The surgical team and the circulating nurse had no reason to suspect that the medication had not been given.

- The patients all came from the CICU, where the surgical team knows and trusts the nurses and, therefore, they do not feel the need to check up on their work. The CICU nurses specialize in cardiac patients so the OR team had no reason to think the nurses would be unfamiliar with the pre-operative medications.
- The pre-operative timeout (checklist) was performed, but it contains no question about pre-operative immunosuppression. The only medication-related question is about pre-operative antibiotics. There is a general question asking about any other concerns, which would not be expected to prompt staff to think about pre-operative immunosuppression.

Additional Questions raised:

- *Why does the EHR order screen not have a record of whether an order has been acknowledged and carried out?*
- *Why isn't the timeout checklist specialized for different types of surgery?*

Circulating RN (Operating Room Nurse)

Responsibilities:

- Perform the final check that the patient is ready for surgery.

Contributions to Adverse Event:

- Did not stop surgery from proceeding despite the patient not having received immunosuppression.

Why?

Mental Model Flaws:

- Believed the patient had received immunosuppression.

Contextual Factors:

- On the order screen of the EHR there is no record of whether an order has been acknowledged and carried out.
- Nobody raised a concern during the timeout procedure.
- The pre-operative timeout checklist is a general OR checklist, so it does not explicitly ask about pre-operative immunosuppression.
- The circulating nurse was not familiar with the special needs of cardiac transplant surgery

Additional Questions Raised:

- Why was the circulating nurse not familiar with the special needs of this type of surgery?

Looking at the CAST results, even at this lower level of the control structure, provides a great deal of useful information to improve practices, such as potential changes in the checklist, the EHR, the handoff, the order sets, and training. Considering only the direct actors in the events will, however, limit the opportunity to prevent future accidents. Extending the analysis to the higher-level controllers can potentially have a much greater effect.

Figure 3 shows that the nursing supervisor, the intensive care unit administration, and the operating room administration have broad responsibilities for preventing adverse events involving medication (and other) errors. Understanding why these controls were not effective is an important source of information about how the system can be redesigned to prevent them.

Samost found that several of the 30 adverse events studied using CAST at this hospital had the same or very similar causal factors, although this fact was not noticed by the original RCA teams. If the first events had been thoroughly investigated in terms of determining why the safety controls did not prevent them, instead of focusing on who among the medical staff were to blame, the later events might have been prevented.

Nursing Supervisor

Responsibilities:

- Ensure medication procedures and handoffs are being executed properly
- Identify necessary improvements in procedures

Contributions to Adverse Event:

- Assigned inappropriate nursing staff to CCU
- Did not ensure proper training and procedures were being followed

Why?

Mental Model Flaws:

- Not aware that medication errors were occurring (note the feedback flaw here).

Contextual Factors:

- Budget had been cut and could not always find cardiac qualified CICU nurses.

Questions raised:

- *Had the nursing supervisor complained about the shortage of qualified nurses? What was she told to do about it, if anything? Were there alternatives actions she might have taken? Did the surgical staff know about the shortage?*
- *How was she supposed to be informed about medication errors? Why did she not know about the related medication errors that had occurred previously in the unit?*

Along with the new questions identified so far, the following original questions can be answered at the higher levels of the control structure, i.e.,

- *Why did it take several weeks to discover the missing pre-operative immunosuppressant administration? While it might not have made a difference in this case as the patient was very quickly treated for graft rejection, the answer to the question might have important implications for preventing or mitigating other events involving medication ordered but not administered.*
- *Why did it take several instances of this type of adverse event for clinical staff to learn about it and make recommendations to handle it? While many medication errors are undiscovered, in this case the previous errors had resulted in incident reports being written.*

CICU Management (Administration)

Responsibilities:

- Oversee medication procedures
- Oversee implementation of procedures
- Investigate adverse events involving medication errors
- Ensure safe practices are being followed in the CICU
- Maintain staffing levels and training
- Establish safe medication procedures

Contributions to Adverse Event:

- Did not establish a safe, standardized medication procedure for administering pre-operative immunosuppressant.

Why?

Mental Model Flaws:

- Believed staff knew how to order and administer all medications
- Believed adverse events were being communicated and handled properly.

Contextual Factors:

- All medication orders were placed as intended in this case.
- The new medication procedure was not communicated to this level.
- The hospital did not have an established *management of change* procedure
- Separate management silos for surgery and intensive care complicates communication between the two departments
- Financial pressures led to cutting overtime and placing staff in jobs they were not trained for. They assumed that cuts were not leading to adverse events because problems were never reported to them.

Questions raised:

- *How are adverse events communicated to those responsible for handling them in this hospital and what procedures are used to investigate them?*
- *Were feedback channels non-existent, non-operational, or unreliable?*

OR Administration

Responsibilities:

- Ensure medications are available
- Ensure safe practices in the OR
  - Establish proper medication procedures
  - Oversee timeout procedures to ensure they are used and effective
- Investigate adverse events

Contributions to Adverse Event:

- Did not ensure safe practices in the OR
- Allowed inappropriate timeout procedures for transplants
- Did not identify systemic factors in previous incidents
- Did not thoroughly investigate adverse events (the same thing happened twice before any of the events were investigated)

Why?

Mental Model Flaws:

- Thought timeout was appropriate, procedures were adequate, use of RCA was appropriate
- Believed the staff knew how to order and administer all medications
- Unaware of this incident until several months later

Contextual Factors:

- Previous medication errors were investigated (if they were investigated at all) using standard RCA that focused on determining who was to blame. Most incidents in this unit are never written up as incident reports. (The majority of the 30 incidents were identified through the assistance of a PA doing chart reviews on cases to assist Dr. Samost in her checklist study.)
- Separate management silos for surgery and intensive care complicate communication between the two departments.
- Tradeoffs between cost and safety influenced decisions including staffing levels, personnel training, and equipment inventory. No procedures were provided to assist in determining the impact of such decisions.
- There was a complete turnover of management team. The new team had different specialties and different styles from the previous team. The new team came from a hospital that specialized in cardiac transplants.
- No plan was established to move from rarely doing transplants to one that specialized in them. The new management was unfamiliar with the department and their background and expertise and did not recognize the limitations of this hospital with respect to procedures related to cardiac transplants.

Additional Questions Raised:

- *What are the procedures for determining whether and how an incident is reported and how thoroughly it is investigated*?
- *What kinds of procedures are followed when significant changes are made to determine whether the changes introduce new or increased hazards*?
- *What channels exist for communication between surgery and CICU management*?

Step 4. Identify causal factors related to the overall design of the safety control structure

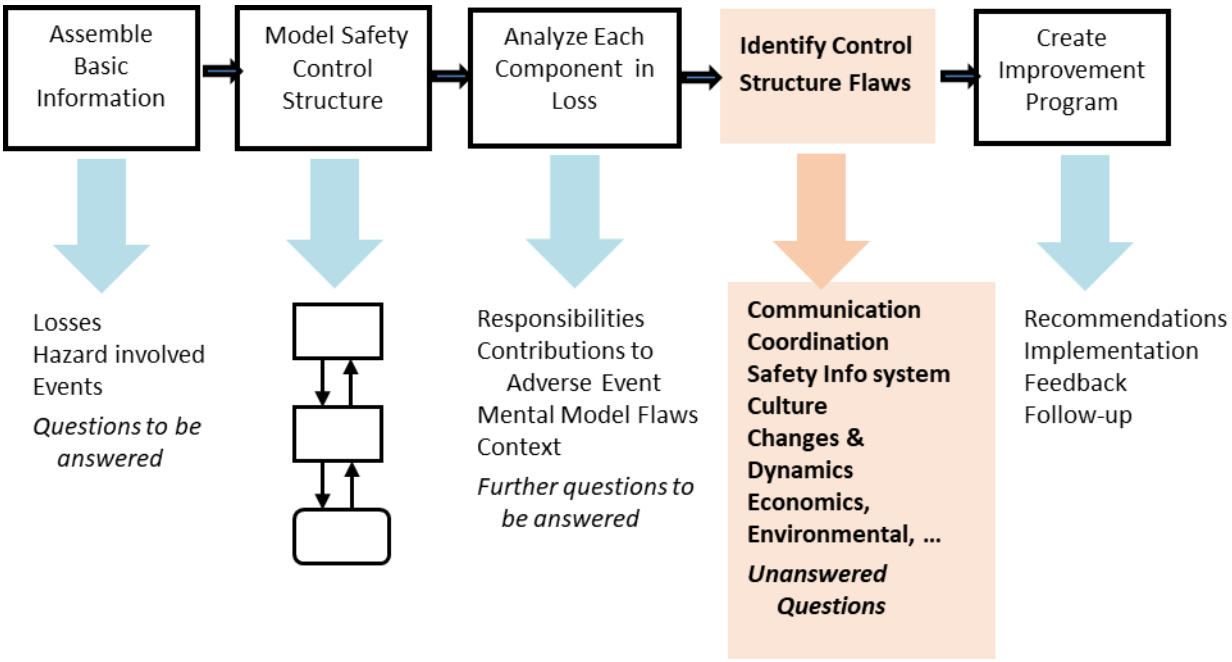

Fully understanding why an accident occurred and making changes to prevent future ones is not complete without considering the systemic (organizational and social) factors that influenced the physical design and the human behavior involved in the loss. Examples of such factors are internal and external economic factors, conditions or practices that lead to fatigue on the part of personnel and poor morale, organizational culture, and so on.

Many of the systemic causal factors involved in this accident and accidents in general cannot be linked by direct arrows to the events and to the behavior of individual components in the control structure (as are used in some accident analysis techniques). Instead, they result from the overall control structure design and factors that impact the behavior of many or all of the individual system components and their interactions. Direct causal factors may be easier to identify than indirect ones, but many times the indirect systemic factors may be more critical to fix in order to prevent a broad range of future accidents, not just the specific circumstances that happened this time.

The important factors to consider will differ with specific adverse events. Examining the contextual factors identified in the previous CAST step will identify some general control structure flaws. In addition, some general categories occur frequently and should be considered when investigating an adverse event. Using any checklist is dangerous in controlling hazards as people tend not to consider anything beyond those on the checklist, surprisingly even safety experts [7]. However, general categories do assist in looking for specific problems. In addition, the individual component causal analyses generated using CAST, particularly in terms of flawed mental models and contextual factors, will help to identify general safety control structure categories that should be investigated. The graphical control structure model can also be useful in identifying many types of systemic factors.

Only some common systemic factors in adverse events are described here: ill-defined responsibility, authority, and accountability; communication and coordination problems; safety culture; inadequate management of change; an incomplete safety information system; and economic and management factors. These were primarily identified during the contextual analysis of the individual system components.

Ill-defined Responsibility, Authority, and Accountability

The control structure model identifies the responsibilities of those involved in the adverse event. Examining these responsibilities should assist in identifying gaps and redundancies.

Confusion may exist about who is doing what, leading to gaps in covering the responsibility. A second cause may be that conflicting control actions are provided. A third potential problem can arise from redundancy. While theoretically redundancy could be considered useful, often it simply ensures that the responsibility is performed completely by nobody. The participants may think that someone else is taking care of it, resulting in nobody doing it.

In this case, both the OR administration and the ICU management had responsibility for handling adverse event reports, but the in-house reporting system was not functioning well nor were reports investigated in a timely manner. The fact that two different management chains were responsible for investigating adverse reports with little communication between them may have contributed to the lack of timely handling of the reports.

This was not a case where the medication error was unknown; it became clear immediately after surgery when the patient began to show signs of rejection that the immunosuppressant might not have been administered, which should have led to asking questions about whether it had been. The eMar (electronic medication administration record) clearly showed it had not. Why did two other similar events have to occur before an investigation ensued? A recommendation to evaluate the adverse event reporting and handling system is certainly well justified.

A basic principle of safety management is that one group needs to be responsible for overall safety, including the adverse event reporting and handling process. This group should be independent and free of conflicting responsibilities. The group need not perform the investigations, but it should ensure that investigations are taking place, that their quality is acceptable, and that recommendations are generated and implemented.

Controlling adverse event reporting and investigation should not be the responsibility of those potentially being investigated. This principle does not mean that those involved should not be part of the investigation process; in fact, they should. But some members should be independent. Also, decisions about whether an investigation should occur and for maintaining reporting channels need to be independent of the direct management involved.

Another common problem occurs when there are controls and controllers for individual system components but nobody is controlling the boundaries and interactions among them. In this case study, there did not seem to be anyone responsible for controlling the handoff between the CICU staff and the surgical team and ensuring it was well defined and operational.

As in any management situation, responsibility, authority, and accountability need to be carefully defined.

Communication and Coordination

Communication flaws between individuals and departments can lead to false assumptions and inconsistent practices, which in turn can create delays and misunderstandings. Looking at the management levels of this hospital's control structure shows that management is divided into autonomous silos operating each division of the hospital. This division makes the boundaries between divisions, such as the line between the CICU and the OR, risky areas for patient care. Differences in protocols, communication styles, assumptions and expectations all increase the risk of information being missed.

In general, a hospital environment is designed to have several different providers working to coordinate patient care [1]. Often, communication related to patient care is very transactional, as each participant (nurse or doctor) is likely only interacting with that case for a limited amount of time and working within a specific set of policies and procedures created to govern the patient. Health

information systems, including EHRs, are used as communication intermediaries among the hospital caregivers to control patient care actions. When there are different doctors and nurses on different shifts, knowledge about the continuity of care for the patient is usually communicated in a computer system.

One can think of the EHR as a shared mental model. A communication process requiring entering information into an information system is highly inefficient in a patient care setting and places an extra burden on each healthcare provider to make certain that the information system is updated promptly to ensure safe patient care. In many cases, the design of the information systems such as EHRs does not support this process.

The rapid pace and constant pressures on the physicians and nurses in the hospital environment can lead to policy and procedure workarounds and multi-tasking that compromise patient safety and contribute to errors. Baker [1] describes the problem with respect to medication errors. In a medication error she investigated using CAST, the hospital requires that medication orders be written and externally verified before administration. In this case, when the physician could not write the order because of duties in the catheterization lab, the care team reverted to verbal medication orders. Even though a verbal confirmation of the order was done, several safety checks in the EHR defined for medication order entry and between care providers were bypassed when the order was relayed verbally between the caregivers. The workplace environment may make enforcing policies and procedures to ensure proper communication difficult.

The hospital administration in this case had developed policies over time to ensure patient safety and quality care. Overall, these policies had been communicated well. What the administrators had *not* communicated to the physicians and nurses sufficiently was how to rapidly prioritize several conflicting policies in a highly dynamic hospital care environment under unusual circumstances.

In general, communication problems can arise in feedback and in coordinating actions. The latter includes coordination and synchronization of mental models.

Feedback

A common contributor to adverse events is missing feedback. Sometimes the feedback channels have never been defined. In other cases, they exist but are not used or do not operate correctly for some reason. In the example of CAST used in this handbook, the nursing and surgical staff did not receive feedback that the immunosuppressant had not been administered even though it had been ordered. Many obvious recommendations could be devised to deal with this feedback problem.

There were also problems in feedback to the higher-level control structures in the hospital. One cannot control a process without feedback about its current state. Were the feedback channels to OR and ICU administration functional? Were there feedback channels for those responsible for the design and use of the EHR and were they aware of previous related adverse events? Was there a feedback channel to inform them? Was anyone responsible for performing a human factors/hazard analysis of the EHR design?

Coordination

When multiple controllers have the same or closely related responsibilities, communication to coordinate activities is important. In this hospital, autonomous silos existed between divisions in the hospital, which led to difficulties in coordinating the activities of different departments.

Potential confusion about shared responsibilities might have occurred in this case, for example, the attending cardiac surgeon, the surgical fellow, and circulating nurse were all responsible to check that the ordered medications had been administered before surgery. Redundancy is a good way to enhance reliability, but can decrease safety. Reliability and safety are two different system properties that can conflict. In this case, redundancy introduced to enhance the reliability of medication administration led

to confusion about who was supposed to administer the immunosuppressant and flawed redundant checks that it had been administered before surgery. All of the redundancy here was undone by a common flawed feedback channel (the EHR). In engineering, this phenomenon is called a *common-cause failure*/*error*. Redundancy can also lead to dangerous complacency.

Other types of communication problems existed here such as in the handoff between the SICU nurse and the surgical team.

Mental model synchronization and update

In the heart transplant example, the people directly involved in the events (the surgeons and nursing staff) all had incorrect process models: The surgeons and circulating nurse thought that immunosuppressant medication had been administered, while the SICU nurse was not aware she needed to give an immunosuppressant and assumed it would be administered in the OR. Feedback and communication channels need to be designed to ensure correct synchronization and updating of mental models related to the hazard.

One important communication and feedback source is the EHR. This patient was very sick and had been admitted to the hospital pre-operatively. Therefore, all pre-operative medications and testing were ordered by the surgical team the night before to be given by the nurses the morning before the procedure. The use of order sets decreases the chance that the surgical team will forget to place important orders. The investigation of this incident and others at the hospital found that the orders were indeed placed, so the control worked as intended. However, the order sets introduce a different source of confusion.

The EHR at this hospital has a poor layout in terms of giving clear orders from the physician to the nursing staff and providing feedback regarding the successful carrying out of those orders. The orders are given a time to be carried out, but they do not explicitly say who is responsible for performing it. In many cases, this omission is not a problem. Common surgeries, such as a bypass, always require pre-operative antibiotics, which are ordered the night before to be given by the anesthesia team in the OR. The nurses in the CCU or SICU know that it is not their responsibility to give these antibiotics despite seeing the order in their order set. However, in the case of a less common surgery like a cardiac transplant, which includes less common orders like immunosuppression, the ambiguity of these order is far more likely to cause confusion about who gives that medication and when.

The other aspect of the EHR that can lead to confusion is the lack of clear feedback regarding the status of the medications being ordered. The order screen shows the medications that were ordered, but not whether they had been administered. To see if the orders were completed, one has to leave the order screen and move to an entirely separate screen in the EHR, the electronic medication administration record (eMAR). The eMAR lists the medications and the time they were given to the patient. There is nothing in the EHR, however, that will flash bright red or be otherwise very obvious to show that an order was given and not carried out. One has to be looking specifically for it to pick up on this scenario.

In the case of common surgeries, the division of labor is clearly known and understood by all parties. However, in the case of a less common surgery, such as cardiac transplantation, which includes less common orders such as immunosuppression, confusion about who is expected to give the medication and when it is given is more likely and, in fact, occurred in this case. A potential recommendation from this CAST investigation is to perform a human factors evaluation of the EHR with respect to medication ordering. Perhaps such an evaluation had in fact been done, and there was a good reason—related to safety or some other healthcare goal—for designing it the way it was. A revisiting of such decisions, however, might be appropriate to determine whether the assumptions underlying the decision were correct.

A second potential source of communication between the CICU nurses and the surgical team about a patient's readiness for surgery occurs in the handoff when the surgical team picks up the patient for transport to the operating room. The handoff is a time when the nursing staff can communicate any concerns and the surgical staff can ask any questions about the patient. However, there is no formal structure in the handoff at this hospital, so important information may not be shared. In this case, the nurses had no concerns as they were unaware that were supposed to provide immunosuppressant medication.

In general, when one controller is responsible for process A and another for process B, nobody may feel responsible for the interface between A and B.

A third opportunity for communication about the patient's readiness for the operation is the pre-operative timeout. The timeout, however, has no question about pre-operative immunosuppression. The only medication-related question explicitly asked is about pre-operative antibiotics along with a general question asking about any other concerns, neither of which could be expected to prompt the staff to think about pre-operative immunosuppression.

A design decision was made when the timeout checklist was adopted to make it generic and the same for every operating room. Designers adapted this checklist from the WHO Surgical Safety Checklist. However, neither the original WHO checklist, nor this adapted version, was ever validated in cardiac surgery. Are there components of this checklist that can be improved to prevent this type of accident in the future? For example, there are many questions on the checklist that do not apply to a cardiac case. Questions regarding the laterality of the surgery and the site marking by the surgeon are superfluous for a cardiac patient. One cannot operate on the wrong heart.

A question asking specifically whether the patient had gotten pre-operative immunosuppression would have been far more helpful in this scenario. However, there is a balance required in writing checklists. Too much detail makes the checklist too long and people will not complete it in a setting with such time pressure, such as an OR. Too little detail and important things get missed, as in this and other adverse events analyzed by Samost.

Creating a long and detailed checklist is not a good solution either. A long checklist is unlikely to be fully completed. In fact, airlines often try to do everything they can to keep the "before takeoff" checklist (the final one before barreling down the runway) as short as possible. Only the most critical or "killer" items are included. Compliance rates are much more likely to be high with short checklists [3].

One recommendation that might result from this CAST causal analysis is to tailor the surgical checklist to make it more specific to cardiac surgery rather than using the same one in every operation. Some general questions, such as the site of the surgery, are not relevant to cardiac surgery, whereas some questions specific to this type of surgery are omitted from the general checklist.

Safety Culture

The culture of an organization is defined by Edgar Shein (known as the "father of organizational culture") as the values and assumptions in the industry or organization used to make safety-related decisions [25]. These values and assumptions provide the foundation for creating organizational rules, policies, and practices. The cultural values and assumptions are established by the highest-level management in the organization or sometimes by the industry as a whole.

Safety culture is the subset of the organizational or industry culture that reflects safety values. Flawed safety culture may include overconfidence and complacency, inappropriate priority assigned to safety, and flawed resolution of conflicting goals. A common belief is that safety conflicts with productivity and profits—in fact, it does not [14].

Features of a good safety culture include openness about safety and safety goals, a willingness to hear bad news, an emphasis on doing what is necessary to decrease adverse events rather than just complying with government regulations or producing a lot of paperwork, belief by employees that

managers can be trusted to hear their concerns about safety and will take appropriate action, belief by managers that all employees are worth listening to and are worthy of respect, and employees feeling safe about reporting their concerns and believing that their voice is valued. Employees need to know they will be supported if they exhibit a reasonable concern for safety in their work and if they give priority to safety in short-term conflicts with other goals such as schedule and cost.

Safety is enhanced when it is a shared responsibility and all employees are considered to be part of the solution and not just the problem. At the same time, responsibility cannot be just placed on the workforce to keep themselves and others safe. Appropriate responsibilities need to be assigned and communicated to individuals at all levels of the organization.

The "Just Culture" concept promotes a process where mistakes or errors do not result in automatic punishment, but rather a process to uncover the source of the error. Attempts to create a just culture in healthcare have focused on eliminating blame but have not always successfully implemented the second necessary piece, which is to uncover the true sources of errors. A just culture is not simply about eliminating blame or evaluating human behavior in a different way. Instead, the focus needs to be on identifying the systemic (system design) flaws leading to the behavior and fixing those flaws or conditions that lead to undesired behavior. CAST provides a way to accomplish this goal.

One particular safety culture problem that exists widely in the healthcare community and contributes to the blame focus is the "bad apple" theory. Blaming workers in accident investigation is based partially on this outdated and erroneous theory [4,5,6]. The bad apple theory rests on the assumption that there are a certain number of "bad apples" or a few people in a system that are responsible for most of the accidents. Safety would not be a problem, in this view, if it were not for a few unreliable humans.

The concept goes back to 1925 when both German and British psychologists were convinced that they would solve the safety problem by identifying and getting rid of the bad apples in an organization. They used statistical analysis over 50 years to determine that there were a cohort of "accident-prone" workers. These were people with personal characteristics that, they claimed, predisposed them to making errors and thus precipitating accidents. The data seemed to imply that a small percentage of people are responsible for a large percentage of accidents. If these people are removed, it was assumed that the system would become much safer.

The bad apple theory existed until WW II, when the complexity of the systems we were creating and that humans had to work within, started to increase dramatically. The theory was finally put to rest in 1951 by two statisticians, Arbous and Kerrich.

The theory did not work because of a major statistical flaw in the argument. For the accident-prone or bad apple theory to work, the risk of error and accidents must be equal across every system. But, of course, it is not. Newer views of accident causation conclude that personal characteristics do not carry as much explanatory power for why accidents occur as context does.

> *When faced with a human error problem you may be tempted to ask 'Why didn't they watch out better? How could they not have noticed?' You think you can solve your human error problem by telling people to be more careful, by reprimanding the miscreants, by issuing a new rule or procedure. They are all expressions of the 'Bad Apple Theory' where you believe your system is basically safe if it were not for those few unreliable people in it. This old view of human error is increasingly outdated and will lead you nowhere.* [4].

The bad apple theory is especially prevalent in medicine, but it permeates almost every industry in the form of the argument that most accidents are caused by human operators. An alternative explanation for the statistics is that humans are usually assigned the blame for accidents, but may not be the real cause.

If one believes that 5% of bad doctors cause most accidents, for example, then simply identifying and getting rid of the doctors who get the most complaints and are involved in adverse events should drastically reduce medical error. Unfortunately, it does not. It may simply get rid of a group of doctors who do the really difficult, tricky work, for example, some oncological cases with a negative prognosis. Certainly, there are people in healthcare (as in all fields) that are not competent and need to be removed. The problem is that they are not the only source of medical error or even a large percentage of it.

In general, if there are system features that are likely to induce human mistakes under some circumstances, such as poor facilities with limited resources and financial and other stresses, then swapping the humans involved, but not changing the system features that are creating the erroneous behavior, will have little impact on the accident rate. We know now that the design of systems has the potential to create various classes of errors or, conversely, to reduce them. Surprisingly, most humans don't realize they are being influenced by their environment.

In the example heart transplant adverse event, individual behavior was influenced by the fact that the EHR had a poor layout and did not provide clear instructions from the physician nor provide clear feedback about the carrying out of orders. The order set, used to decrease the chance of the OR team forgetting to place important orders, as described previously, introduced potential for confusion. For common surgeries, everyone knew their role. But heart transplantation in this hospital was less common and had not been done for a long time, leading to confusion about who was to do what. In fact, the behavior of each of the people involved was reasonable considering the context in which they were working, which is usually the case.

Management of change

Most accidents in all industries are preceded by some type of change, such as, in this case, the change to doing heart transplants in this hospital. Without any change, there is no reason for the process not to behave as safely as it has previously.

In fact, complex systems are always adapting and changing, and therefore simply observing what "went right" in the past does not provide any information about the future [11]. Changes may occur over time and may lag documentation or even recognition of those changes [18].

Changes can be planned or unplanned. Sophisticated organizations use *management of change* (MoC) procedures to identify and evaluate any new or increased hazards that may result from a planned change. The same principles should be used for changes in surgical or other hospital procedures.

In the heart transplant example, the entire OR cardiac leadership team turned over with new leaders coming to the hospital from other academic medical centers. Every medical center operates slightly differently. They have different specialties and different styles. This medical center performs relatively few cardiac transplants every year and standard practices did not account for new requirements. The new leadership came from a hospital that specialized in performing cardiac transplants. Their goal for this hospital was to transform it into a cardiac transplant specialty center. There was no plan in place to help move from a center that rarely performed transplants to one that specialized in them.

Even when there are documented MoC procedures, they are often skipped in practice, resulting in losses that need never have occurred. There must be responsibility assigned for enforcement of MOC procedures and feedback channels to determine whether they are being followed and, if not, why. If the cause of skipping them is in the MOC change procedures themselves—too onerous, too difficult, too time consuming, etc.—then the MOC procedures may need to be changed. This hospital did not seem to have any documented management of change procedures.

Cardiac transplants require a team effort, far more than just bringing in surgeons with appropriate specialized surgical experience. The idea of managing change and ensuring safety during these transitions is not an idea that healthcare uses as much as other industries. Hospitals may have no standard procedures for introducing new surgical techniques or transitioning to different types of care delivery. Most institutions improvise as best they can and solve issues when they arise instead of trying to take a proactive approach to safety during periods of transition.

Planned changes may also not take into account established workarounds that circumvent the problems in the system design that can inhibit acceptable patient care. Changes can undermine those workarounds. A holistic view of the system is needed to prevent this degradation.

The problem becomes more difficult to handle when changes occur that are *not* planned and may be unnoticed until an adverse event occurs, and sometimes not even then. In general, systems tend to migrate to states of higher risk due to pressures to advance various goals [21]. Under pressure, people start to violate their own rules and justify it by arguing that safety will not be affected. When no adverse events occur over time, our perception of risk is often reevaluated downward even though risk has not actually changed at all or may even have increased due to complacency and the changes occurring. At some point, risk gets so high that an adverse event becomes almost inevitable.

The processes in the safety control structure need to interrupt this risk re-evaluation before safety margins are eroded. In addition, there needs to be an alerting function to a responsible person when behavior is contrary to the true level of risk. Complacency can be short circuited by providing information that allows accurate risk assessment by decision makers. Another strategy is to allow flexibility in how safety goals are achieved so that ignoring rules is less frequently necessary.

Identifying when risk may be increasing requires designing and using feedback channels to those responsible for safety.

Leading indicators are one way to identify when risk is increasing in a system. A new approach to identifying leading indicators has been created, called *assumption-based leading safety indicators* [15]. Briefly, the idea is that certain assumptions are made during system development that are used to design safety into a system. When, over time, those assumptions no longer hold, then the organization is likely to migrate to a state of higher risk. Leading indicators, then, can be identified by checking the original safety-related assumptions during operations to make sure that they are still true. The assumptions underlying the design, of course, must be identified and documented for this approach to work.

Managing change is beyond the scope of this handbook, but an important goal in adverse event investigation is to identify how changes, either planned or unplanned, contributed to the events and the limitations in managing these changes within the organization.

Safety Information System

A comprehensive and usable safety information system (SIS) is key to having success in preventing adverse events. Good decision making requires accurate and up-to-date information about the safety of the system. After studying organizations and accidents, Kjellan concluded that an effective safety information system ranked second only to top management concern about safety in discriminating between safe and unsafe companies matched on other variables [12].

The SIS is essentially a “shared process model” and a source to update individual mental models, to learn from events, and to improve the safety management system (SMS). It can provide the information necessary to detect trends, changes, and other precursors to an accident; to evaluate the effectiveness of the safety controls; to compare models and risk assessments with actual behavior; and to learn from events and improve the design of the SMS itself. After major accidents, it is often found that the information to prevent the loss existed but was not used or was not available to those involved. Often,

lots of information is collected only because it is required for government reports and not necessarily because it is considered a necessity for the operation of an effective SMS.

This hospital had been collecting data on incidents for the past three years using an in-house incident reporting system. Additionally, unexpected clinical outcomes were investigated by the clinical team to determine if the outcomes were caused by preventable errors. In this case, the adverse event was not recognized and investigated until several months later, after two similar incidents had occurred.

This fact points to several challenges with adverse event investigation in healthcare [22]. The first is the challenge of even recognizing that an adverse event has happened. Was a patient's death or complication an expected side effect or the result of an adverse event? Cardiac transplant patients can show signs of rejecting the organ and die even in the presence of correct immunosuppressive regimens. In this case, however, the rejection process was recognized right after surgery. It would have been possible to check whether the immunosuppression had been administered. It was not until a pattern of harm was seen over several patients that an investigation ensued.

In addition, what constitutes a reportable incident differs from hospital to hospital and department to department. Implementing standard requirements for reporting and investigation might improve learning from events in healthcare.

Finally, there is the process of investigation. Once an accident has been identified and reported to management, management still needs to perform an investigation that will capture meaningful recommendations for system changes to prevent future incidents. In the cardiac transplant case, as stated, the investigation was delayed. When it was finally discovered after several occurrences where immunosuppressant had been ordered but not given, the recommendation was made to add immunosuppression to the checklist. Samost found that this recommendation had yet to be implemented at the time of her general study of timeouts, suggesting that there is room for improvement also in the way that management responds to adverse events and the recommendations that result from them.

In general, information may be collected by individual organizations or by industries and coordinating groups. Both adverse events and incidents or near misses need to be collected. Examination and understanding of near misses can warn of an impending accident and provide important information about what conditions need to be controlled. Civil aviation and nuclear power are industries where near misses are reported widely. Healthcare has special difficulty in learning about near misses because they may be unknown or very hard to identify. It can be challenging to determine whether a patient's death or complication was an expected side effect or the result of an adverse event or even that an unsafe action has occurred without a noticeable adverse result.

To be most useful, the information must be accurate and timely and it must be disseminated to the appropriate people in a useful form. Three factors are involved: collection, analysis, and dissemination. Many groups in healthcare do collect information, but major hurdles exist in analyzing the data and disseminating it to those who need it.

Data from near-miss reporting shows that it is often filtered. The role of blame again limits the collection of information. The Just Culture concept was devised to try to overcome this problem and to encourage reporting, but those making reports must be convinced that the information will be used for constructive improvements and not as a basis for criticism or disciplinary action. Potential legal liability may come into play.

Aviation has made great use of anonymous monitoring systems to collect information. Relatively recently, automated monitoring systems have been installed on most commercial aircraft, but implementing such systems would be much more difficult in healthcare, both technically and culturally.

All industries have the problem of unconscious filtering and omission of the systemic factors from reporting systems.

Information about the efficacy of the SIS itself needs to be evaluated when investigating adverse events. Were previous similar events identified and properly investigated? What types of causal factors were included in previous investigations? Was the information recorded and properly disseminated to those who could prevent future events? Did the information get to those who could do something with it? Have the same causal factors been identified previously? If so, were recommendations implemented to try to prevent reoccurrences? If they were implemented, why were they not effective? If they were not implemented, why not? Questions such as these need to be raised in the adverse event investigation.

The same and similar causal factors in the cardiac transplant case had occurred previously and been identified but analysis and intervention did not occur until after several instances. An important goal of the investigation once it got started should have been to identify why the previous cases were not investigated or did not lead to attempts to prevent reoccurrences and how to change the hospital processes to react and intervene sooner.

Adverse events are often repeated, and their investigation reports are one of the most important sources for hazard analysis and control. When conducting an investigation, the safety information system and any possible impact on the events needs to be examined. Was the information needed to prevent the loss not collected or stored? Was it lost in the collection or analysis process? Was it not available or not easily retrieved in the daily activities of those involved?

Disseminating information about adverse events is the third part of the equation along with collection and proper analysis of causality. Healthcare is unique in the occurrence of medical equipment providers contractually preventing hospitals from sharing information about the contributions of their equipment to adverse events. At the other extreme, payor organizations (such as Medicare) have been very successful in ensuring that adverse events are thoroughly investigated and solutions put into place to prevent repetition.

Local hospital safety information systems will have limited efficacy compared to national or international ones because trends and problems can be spotted faster and disseminated more widely. But the existing more global systems are often limited in their functionality.

Economics and Resources

Many hospitals are for-profit entities in the U.S. and, even when not, resources are usually limited in some way. Economic concerns increase the focus on cost reduction and personnel efficiency while trying to keep patient outcomes optimal [1]. Healthcare providers are continually being pressured to provide more efficient care more quickly, leading to inevitable errors and subsequent losses. Nurses frequently find themselves needing to care for more patients in a shorter amount of time. All of these pressures lead to shortcuts and workarounds and therefore pushing the boundaries of safe practice to save money. The causal factors related to economics need to be included in adverse event investigations.

Actively re-engineering the hospital workplace environment might alleviate the need for such workarounds [1]. In the end, the costs of an adverse event, not only monetary but the incalculable costs such as provider mental health (the “second victim” phenomenon), usually are greater than the amount saved by not investing in safety. It is usually easier and more cost-effective in the long term to redesign the system structure than to attempt to change people.

Organizational Design and Safety Management System

The design of an effective safety management system is the topic of many books and beyond the scope of this handbook. But the organizational design and safety management system within which the adverse event occurred is clearly an appropriate area of investigation. Why did the features of the existing safety management system not prevent the adverse event?

A common factor in adverse events is bottom-up, decentralized decision making and lack of coordination among decision makers. Each local decision may be "correct" with respect to the mental models of those participating in that particular decision and within the local context. But independent decisions in complex, sociotechnical systems can lead to a loss when decisions and organizational behaviors interact in unsafe ways.

Flawed decisions may not only result from limitations in the boundaries of the system view available to the individual participants. In the case study in this handbook, the decisions made by the physicians and nursing staff were correct with respect to their own mental models and limited view of the overall state of the system and the patient (medications administered). Flawed decisions may also result from the activities of other decision makers within the system as a whole. Individuals cannot be expected to be cognizant of those other decisions and activities when the interfaces and interactions are ill-defined. The high reliability of individual behavior in a complex system does not ensure that the system will be safe.

Additional Systemic Factors

Other general categories of systemic factors should be considered in a CAST analysis. Because the adverse event case study used in this handbook was not investigated until long after it occurred and only as part of an experimental study, not enough was known to determine whether these other factors might have been involved. One important factor is poor risk management in regard to tradeoffs between safety and profits (for profit-making care facilities) or tradeoffs between safety and other factors such as efficiency and resources available.

Step 5: Generate Recommendations and Create an Improvement Program

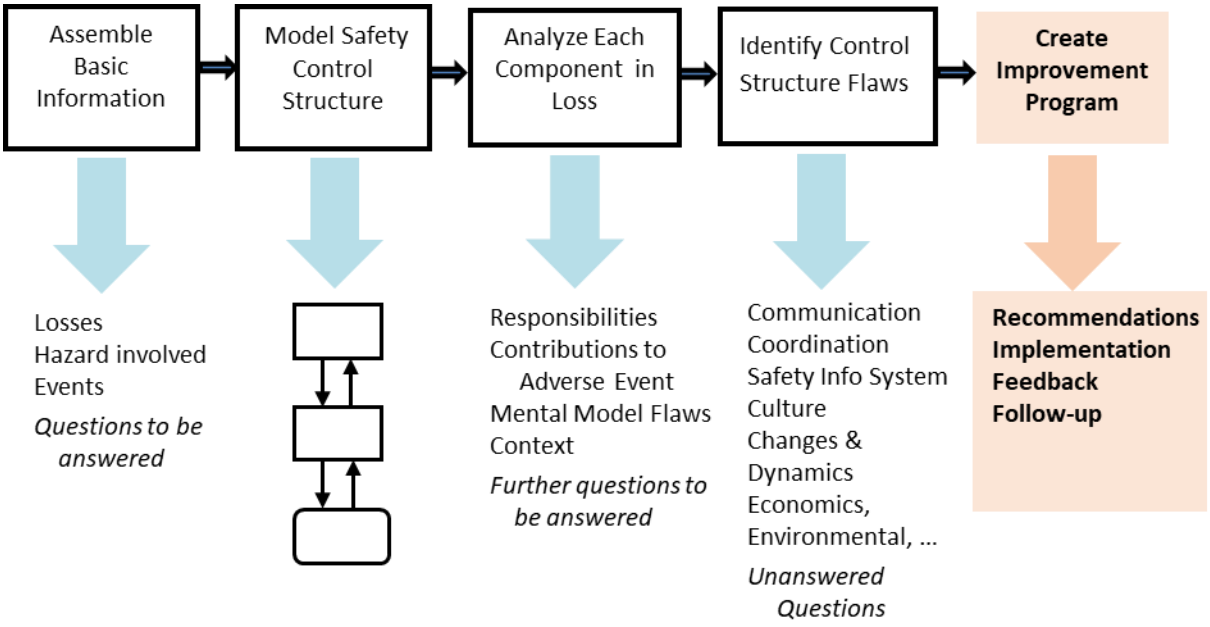

Summarizing the Causal Factors Identified

By the end of the CAST analysis process, there should now be a complete explanation of why the adverse event or incident occurred. In addition, all the questions that were generated should have been answered. If not, further action may be warranted to obtain the answers, such as further analysis and study.

Following this process allows for direct traceability from accident causal factors—process model flaws, contextual factors, systemic factors—to recommendations. This traceability serves two functions: (1) ensuring that all causal factors are addressed and (2) providing a rationale for system changes. This rationale allows analysts and engineers in the future to understand why these particular policies exist and what to consider when changing them.

Using the standard blame approach, the causal factors identified in this example appear to be:

1. The CICU nurse did not give immunosuppressant medication.
2. Nurses did not tell the surgeon in the handoff that it had not been given.
3. The surgeon started surgery without the patient receiving immunosuppressant, despite executing the timeout (checklist) before surgery.

The actual RCA report on the adverse event identified these factors and focused on encouraging the use of checklists and additional policies and procedures to bring the human behavior into compliance, rather than addressing potential safety problems at the system level, which would make compliance a non-issue.

In summary, rather than focusing on who did what wrong, a systems approach to adverse event investigation shifts the focus to why did the events occur and what changes can be made in the system design to prevent similar events in the future. The nurses and physicians as well as the administrators all acted reasonably given the information they had.

In CAST, there is no "root" cause. Or at least there is always the same root cause, i.e., the system design contributed to or did not prevent or control the hazards so as to prevent an adverse event [14]. Instead of assigning blame, the goal is to identify the flaws in the system structure that contributed to the adverse event. The goal is to obtain the information necessary to determine how to redesign the safety control structure (safety management system) to be more effective. In the case study, the nurses and physicians as well as the administrators all acted reasonably given the information they had. Eliminating such causes requires (1) redesigning the system as a whole to eliminate unsafe interactions and to ensure that everyone has the correct information on which to base their actions, and (2) designing a culture based on making tradeoffs in ways that allow the participants to make decisions with a proper emphasis on safety.

If the adverse event investigators focus their attention on the particular individuals involved or even on classes of people such as management, the problems that need to be fixed, i.e., the system design as a whole and the dynamics of the system, may not be identified or get appropriate attention. Focusing on blame detracts from learning. Healthcare professionals work within a high-stress environment. Focusing on blame when something goes wrong does nothing to change the system design features that lead to unsafe behavior.

All human decision-making is based on the person's mental model of the current state and operation of the system being controlled. Determining how to prevent unsafe behavior in the future requires not only identifying flaws in the participants' mental models but also *why* those flaws existed.

Generating Recommendations

Once the other parts of the analysis are completed, generating recommendations should be straightforward. The results of the Samost CAST analysis and the recommendations of the hospital RCA report differ significantly. The CAST findings and recommendations focused on the control structure of the entire medication administration system within cardiac surgery in the hospital and not on individual behaviors in the scenario. The CAST recommendations are framed as procedural and environmental changes that can be taken at the hospital level rather than reinforcing policies for enhanced human compliance.

Examples of some recommendations that might be generated from the CAST analysis include:

- Change the EHR to provide more obvious feedback on both screens if an order has not been carried out or missing doses of medication. More generally, perform a human factors safety analysis of the EHR to find other design features that can contribute to adverse events.
- Evaluate the pre-operative checklist and consider changes to make it more specific to cardiac surgery. A checklist designed for every OR is likely to be less useful than more tailored designs.
- Make the wording more explicit on order sets as to who is responsible for carrying out the orders in addition to when they should be completed. Order sets are useful to prevent forgetting orders, but bad when they cause ambiguity. Consider breaking up the large order set into orders for the SICU team and orders for the OR team.
- Institute a formal handoff procedure between the CICU and the surgical team. Include mention of all the pre-operative medications and labs.
- Implement a formal Management of Change (MoC) process. The MoC should include changes in leadership, ensuring that everyone understands the expectations and assumptions.
- Create a robust incident reporting system to help track problems, especially on the interface between two services (e.g., surgery and intensive care). Make it easy to write up and access the reporting system. Show people that you take their reports seriously and are doing something to make improvements.
- Conduct weekly management meetings between the CICU and the Cardiac Surgery leadership. Promote communication and create policies for interactions between departments.

These recommendations can all be traced back to the contextual factors or process model that they are designed to address.

Note that the CAST recommendations focus on how the hospital administration can improve its safety activities and how the facility's culture can be changed or evolved to avoid losses. The RCA analysis and recommendations omitted entirely those factors, mostly likely because the process does not encourage examining them.

Examining the CAST analysis of all thirty adverse events, Samost found that the majority of the factors related to equipment were specific to the incident. At the other end of the spectrum, however, the same management and contextual factors kept appearing over a wide range of incidents. The new cardiac surgery management as well as financial and time pressures contributed to many of the incidents, ranging from missing immunosuppression to a delay while a piece of equipment was located. Process/policy and management factors contributed to nearly every adverse event, while infrastructure played less of a role. This suggests recommendations targeting these management and process contextual factors will have a broader impact that implementing recommendations that apply more to a specific piece of equipment or changing the infrastructure of the medical facility.

The biggest complaint about CAST we hear is that it generates too many recommendations. This complaint is only justified if the goal of the accident investigation is to make as few recommendations as possible. Accident investigation has had the goal of identifying "root causes" or "probable causes" for so long that it may be a cultural shock to start looking at more general causal factors that generate more recommendations.

One of the objections raised to including a large number of recommendations is that responding to them is overwhelming. This is simply a logistical problem and not one that should be solved by learning less from each accident. There is no reason that recommendations cannot be prioritized according to specified criteria. There is also no implication that all the recommendations must be implemented immediately. Some recommendations will be relatively straightforward to implement immediately while

others may take longer. Some may require such extensive changes that implementing them will take a great deal of effort and resources. Examples of the latter include establishing a new oversight agency or changing regulations and laws. Difficulty of implementation is not an excuse to omit a recommendation from the accident report, but it may be a good reason to categorize it as a longer-term goal rather than an immediate fix.

In addition, taking steps to "jury-rig" short-term solutions should not be an excuse for endlessly delaying comprehensive and effective solutions.

Creating a Learning Process

Sometimes recommendations are made but never implemented. Not only must there be some way to ensure recommendations are followed, there must also be feedback to ensure that they are effective in terms of achieving the goals and strengthening the safety control structure.

Essentially there are three basic requirements for creating a continuous learning system:

1. Assigning responsibility for implementing the recommendations
2. Checking that the recommendations have been implemented
3. Establishing a feedback system to determine whether the recommended changes were effective in strengthening the controls.

The third requirement implies the need to collect evidence about the effectiveness of the recommendations. Such feedback can come from audits and inspections and from the analysis of later incidents to determine whether previous recommendations were successful. Such an activity is a critical component of any safety management system, but it is often omitted. Rather than waiting for another adverse event, positive steps should be taken to ensure that the recommendations have been implemented and are effective.

Subsequent accidents or losses, particularly if they are analyzed using CAST, provide a rich source of information to understand why previous recommendations were not effective and what else is needed. Was the original causal analysis flawed? Were assumptions about the potential effectiveness of particular improvements incorrect? Did other changes occur that thwarted the attempt to strengthen the safety control structure? Did the planned changes result in unforeseen consequences? If the fixes were not effective in removing the causes of adverse events, then an investigation of the process for creating recommendations and responding to them may be warranted. That is, not only must the factors involved in the incident be corrected but also the process that led to inadequate fixes being implemented after previous incidents or adverse events.

The goal here is to ensure that your organization is continually learning and improving its risk management efforts so that the potential for losses is reduced over time.

Training and Implementation

If possible, some participant(s) in the investigation of adverse events should have special training and knowledge. A trained investigation team including experts in human factors, management, organizational culture, and multiple specialties is desirable. In aviation, accidents are investigated by experts with different knowledge. Local participants in the events or processes should also be included.

The Cost of Sophisticated Adverse Event Analysis

One question remains: Does it cost more to use a more sophisticated adverse event investigation process as described in this handbook? The answer is usually no. The CAST process involves asking different questions and interpreting the information obtained in a different way. The cost of collecting the information usually remains the same.

In addition, the alternative to better adverse event analysis is to continue to have preventable incidents, most of which usually dwarfs the cost of more sophisticated analysis of those incidents. Every major accident in all fields has had precursors that might have been used to prevent the major loss. The cost and time required to do a complete causal analysis will usually pale in comparison to the cost of a future loss due to the same factors that were never fixed. Doing a series of superficial investigations that omit the most important causal factors that could contribute to future losses is not cost effective: it wastes both resources and opportunity.

Summary

There is always a temptation to identify individual human errors as the cause of major contributor behind adverse events. To maximize learning, the focus instead has to be on identifying why the people involved behaved the way they did—why their actions made sense given the contextual factors surrounding them at the time. The CAST example in this handbook pointed to general weaknesses in the controls used at this hospital to prevent such events. The results can be used to not only reduce highly similar events but to address a larger class of related events.

Some unique aspects of CAST compared to other RCA processes are

- A shift in focus from blaming individuals and individual errors to considering the role of flaws in the system design;
- Analysis driven by a common model of the system and how the processes work, how components fit together and interact, and emphasizing a larger scope of factors in the investigation;
- Capture of a broader set of causal factors such as managerial decision and organizational flaws
- Consideration of such things as mental models and other human factors aspects and the role of feedback from the system within the healthcare professionals are working.

From a longer-term perspective, we need to design safety control structures to prevent adverse events *before* they happen and not rely on learning after the fact. The type of prospective analysis required is called *hazard analysis* in engineering. Most traditional hazard analysis techniques (such as FMEA or Fault Trees) are based on a component reliability model and do not provide the information necessary to identify the systemic changes required for preventing accidents in complex sociotechnical systems like healthcare [17]. New, more powerful hazard analysis methods, such as STPA (System-Theoretic Process Analysis), are based on the same systems theory causality model as CAST.

References

1. Baker, Elizabeth (2022). *Safety in Hospital Medication Practices: Applying STAMP Processes*, S.M. Thesis, System Design and Management Program, January.
2. Checkland, P. *Systems Thinking, Systems Practice*, New York: John Wiley & Sons, 1981.
3. Degani, A. and E. Weiner (1993). Cockpit Checklists: Concepts, Design, and Use, *Human Factors*, Vol. 35, Issue 2, June.
4. Dekker, S. (2006). *The Field Guide to Understanding Human Error,* Ashgate Publishers, 2006.
5. Dekker, S. (2016). *Just Culture*, 2nd edition. CRC Press, 2016.
6. Dekker, S. and Leveson, N. (2014) he systems approach to medicine: Controversy and misconceptions, *British Medical Journal: Quality and Safety*, 2014 and "The bad apple theory won't work", British Medical Journal: Quality and Safety, Nov.

7. Fischoff, B., P. Slovik, and S. Lichtenstein (1978), "Fault Trees: Sensitivity of Estimated Failure Probabilities to Problem Representation, Journal of Experimental Psychology: Human Perception and Performance, 4(2): 330-334.
8. Grimmett, W. (2020) Early Australian Experience with using CAST to investigate medical events, *STAMP Workshop*, March, http://psas.scripts.mit.edu/home/wp-content/uploads/2020/08/Early-Australian-Experience-with-CAST.pdf
9. Grimmett, W. (2021) Introducing STAMP to a major health organization, STAMP Workshop, June 2021, http://psas.scripts.mit.edu/home/2021-stamp-workshop-program/
10. HHS, New HHS Data Shows Major Strides Made in Patient Safety, Leading to Improved Care and Savings, May 7, 2014, U.S. Dept. of Health and Human Services, https://innovation.cms.gov/files/reports/patient-safety-results.pdf
11. Hollnagel, Erik (2014). *Safety-I and Safety-II*. CRC Press.
12. Kjellan, Urban (1982) An evaluation of Safety Information Systems at Six Medium-Sized and Large Firms, *Journal of Occupational Accidents* 3:273-288.
13. Kohn, L., Corrigan, J., Donaldson, M., (1999). *To Err is Human*. Washington, DC: Institute of Medicine.
14. Leveson, Nancy (2012) *Engineering a Safer World*, Cambridge, MA: MIT Press, 2012.
15. Leveson, Nancy (2015), A Systems Approach to Risk Management through Leading Safety Indicators, *Reliability Engineering and System Safety*, 136(April):17-34.
16. Leveson, Nancy (2019). *CAST Handbook: How to Learn More from Incidents and Accidents*, downloadable from http://psas.scripts.mit.edu/home/materials/
17. Leveson, Nancy (2023). *An Introduction to System Safety Engineering*, Cambridge, MA: MIT Press
18. N. Leveson, Safety III: A Systems Approach to Safety and Resilience, Boston, MA, 2020. [Online]. Available: http://sunnyday.mit.edu/safety-3.pdf
19. Leveson, N., A. Samost, A., Dekker, S., Finkelstein, S, and Raman, J. (2020). A Systems Approach to Analyzing and Preventing Hospital Adverse Events, *Journal of Patient Safety*, Vol. 16, No. 2, pp./ 162-167.
20. O'Neil, M. (2014). *Application of CAST to Hospital Adverse Events*. MIT S.M. Thesis.
21. Rasmussen, Jens (1997). Risk Management in a Dynamic Society: A Modeling Problem. *Safety Science*, 27, no. 2/3: 183-210.
22. Samost, Aubrey. *A Systems Approach to Patient Safety: Preventing and Predicting Medical Accidents Using Systems Theory*, MIT Masters thesis, June 2015
23. Senge, Peter (1990). *The Fifth Discipline: The Art and Practice of the Learning Organizatio*n, New York: Doubleday.
24. Silvis-Cividjian, N. (2022), Applying CAST to Analyze an incident in radiation therapy, STAMP Workshop, http://psas.scripts.mit.edu/home/stamp-workshop-presentations/
25. Shein, Edgar (1985) *Organizational Culture and Leadership*, CA: Jossey-Bass Publishers.
26. Sterman, John (2000) *Business Dynamics: Systems Thinking and Modeling for a Complex World*, New York: McGraw Hill.

27. Thwink.org, https://www.thwink.org/sustain/glossary/EventOrientedThinking.htm

28. Von Bertalanffy, Ludwig (1969) *General Systems Theory: Foundations*, Braziller, NY

29. Wallace, L., Spurgeon, P., & Earll, L. (2006). *Evaluation of the NPSA's 3 Day RCA Programme: Report to the Department of Health Patient Safety Research Programme*. Coventry: Coventry University.

30. Weiner, Norbert (1965) Cybernetics: of the Control and Communication in the Animal and in Engineering, Cambridge, MA: MIT Press.

31. Wong, L.M.K. Enabling Effective Safety Learning in Healthcare: Implementing CAST and Designing the STAMP-Enhanced Learning and Reporting System, Ph.D. Dissertation, MIT, Sept. 2021.

32. Wu, Albert (2000). Medical error: The second victim, *BMJ*, 320(7273):726-727.

33. Wu, A., Lipshutz, A., & Pronovost, P. (2008). Effectiveness and efficiency of root cause analysis in medicine. *Journal of the American Medical Association,299*(6), 685-687.

Appendix: A Brief Introduction to Systems Theory

Much of engineering involves building models of the system being created or studied and analyzing those models. When the system is a physical or natural system, the models often are composed of mathematical equations, primarily using continuous mathematics (e.g., calculus and differential equations). The model or analysis of human or social systems requires a very different type of engineering model and analysis method, as demonstrated in this handbook, based on systems theory or systems thinking.

Systems theory provides the scientific foundation for the study and design of complex systems. It can be traced back about 60 years ago to researchers studying biological systems [28], the designed aspects of social systems [2], and the relationship between man and machines [30]. It quickly spread to management, the social sciences, and system engineering. As one example, Margaret Mead is often credited with introducing systems theory in anthropology. System dynamics [26], mostly used in management, is based on systems theory. Only a very brief introduction to the general principles is provided here.

Systems theory provides the scientific foundation for the study and design of complex systems that have the following basic properties:

Goal-Oriented (Teleological): Engineered (designed) systems, including social systems, are not just a set of connected components that interact with one another but have an *overall purpose or goal*.[3] The highest-level system purpose—providing healthcare, in this case—is achieved through the operation of the separate components and by the operation of the system as a whole. The components are designed with a subgoal or purpose in mind that is part of the overall goal. For example, a goal of the healthcare diagnostic laboratory system, which is itself a subsystem within the overall healthcare system, is to provide accurate and timely laboratory test information for decisions made in the larger healthcare system. The successful achievement of the subgoals by the components is necessary to achieve the overall system purpose.

[3] In contrast, *complexity theory*, created a few years after system theory, is most applicable to systems that do not necessarily have a purpose, such as weather.

Interdependent and Interconnected:

The world and the systems in it are interconnected and interdependent. Science and engineering have traditionally handled complexity by decomposing complex systems into components, analyzing the components separately for some property, and then combining the results. For example, problems in taking samples, data transfer between components, diagnostic problems within individual laboratories, and so on are solved in isolation from each other. An assumption is made that this combining process provides an accurate result for the system as a whole.

While useful to some extent, this focus on individual components of a complex system can lead to missing important problems that occur in the interactions among the components, such as the sharing of electronic data with other laboratories or healthcare facilities or communication in the transfer between the CCU and the operating team.

The connections are usually not simple in a complex system because the components mutually interact with each other, with one impacting the behavior of the other(s). For example, a laboratory test order may be initiated by a medical provider within an electronic health record (EHR), a uniquely coded message is sent to a laboratory information system (LIS), which in turn sends a uniquely coded message to an in vitro diagnostic (IVD) device, where the test is performed by a lab technician. Once the test is complete, the lab results are returned to the LIS and EHR using a series of system interfaces and ultimately, to the provider, shared with the patient, and sent to the billing system. While each step may be performed perfectly, serious problems can arise in the interfaces between these components. For example, the test requested may be recorded using a different coding system than that used by the laboratory and/or the EHR and LIS (Laboratory Information System].

Interdependencies are common in the designed parts of the overall healthcare system. Such interdependencies can lead to undesired effects when the design of one component is changed in isolation from the others. This phenomenon is called the *Law of Unintended Consequences*. As a simple example, changing data formats in one part of the system may have unintended consequences throughout the entire system.

Holistic

If we want to fix something or intentionally change the behavior of a complex system, we must first understand the system as a whole or we are very likely not to achieve our goals. In the worst case, we increase the number or type of adverse events.

In addition, some system properties, such as safety, *emerge*, that is, are created through the operation of the system as a whole. Diagnostic laboratory data safety emerges from the interaction and properties of multiple system components such as the way that samples are taken from patients, the mode of the transfer of required information to the laboratory facility and back, the calibration of instruments in the laboratory, the knowledge of laboratory workers, the consistent interpretation of data among the various groups involved, the implementation of accepted data standards, and so on.

A common way of expressing this idea is by saying that *the whole is greater than the sum of the parts*, an idea that has been attributed to Aristotle. For example, individual laboratory tests (e.g., serum glucose, hemoglobin AIC, oral glucose tolerance test) provide clinicians and patients with important information, but the combined results of a series of laboratory tests are required to diagnose diabetes. Laboratory data safety may depend on the way that diagnosis is performed in different laboratories, the expectations of those receiving the data, timing, adverse event reporting and handling, technological changes affecting different parts of the system, and so on.

The concept of emergence means that properties of complex systems may “emerge” when the parts operate together that may not be visible when looking at the separate components in isolation.

Understanding a system property, such as safety, requires looking at all the components as an integrated system, that is, as a whole.

Contextual

All behavior is affected by the context in which it occurs. We cannot understand, predict, or change the behavior of something without looking at the context in which that thing (or person) is operating. For example, the process of healthcare personnel taking a sample involves not only the ability of the individual involved but also features of their environment such as the equipment available, distraction, time stresses, patient compliance (e.g., "needle phobia"), and so on. Leveson has suggested that human error is a sign that the system design (the context for the activities in the system) needs to be changed [16].

Human behavior is not only driven by the context in which it occurs, but also indirectly by our mental models of that context. Our mental models impose a structure that allow us to deal with a "messy" world. For example, physicians believe that the data they received has been processed in specific ways and will act on that belief. Perception is also affected by expectation, that is, we often interpret what we see through the lens of what we expect to see. If we expect to see the results from a diagnostic lab expressed in micrograms, we might not notice that the units reported differ from what we expected, even if the new units are noted somewhere on the lab results report.

Dynamically Complex and Adaptive

In dynamically complex systems, such as healthcare, cause and effect are not related in a simple way. Understanding and changing such systems is difficult because they are continually changing and adapting to the current conditions, both within the system and in its environment. As an example, a healthcare payor may try to reduce the number of potentially dangerous incidents of a particular type by creating financial incentives for hospitals that have a low number of such events. Hospital administrators, in response, may create incentives for employees to reduce those types of incidents.

Unfortunately, the result of such incentives may not be what the payors expected: Instead of reducing the incidents, the incentives may lead only to reduced reporting of them. As a result, adverse events may not decrease and may even increase. Or the attempts to reduce that particular type of incident may lead to increases in other types of incidents, perhaps leading to even worse adverse events. In general, attempts to reduce adverse events may not have the intended result because the system reacts and adapts in unexpected ways to those reduction attempts.

The solution is to impose constraints in system design and operation to control the dynamics that prevent the system goals from being achieved and the attempted fixes to be effective.

Non-Linear

Causality is sometimes simplified to assist in understanding and preventing adverse events. The most common simplification is to assume causality is linear. Linear causality means that each event is the cause of an event that directly preceded it. A common conception of causality in safety is to think of it as holes in Swiss cheese, with the holes lining up in a linear fashion. While this model can be imposed on any set of events preceding an adverse event, it omits important information such as the reasons *why* the events occurred—which are usually much more complicated than just the existence of a single preceding event or the failure of a protection device. Linear causal models therefore do not provide enough information to effectively prevent large categories of adverse events.

In addition, causality can be circular. For example, doing something successfully (without an adverse event occurring) can lead to complacency and an assumption that the process is and will always be safe.

Doing A leads to success, which leads to doing A the same way again and again, reinforcing the appearance of the safety of doing A. This circular loop may continue (with complacency increasing as nothing bad happens) until some feature of or change in the system or its environment is encountered such that A leads to an adverse event. Another example of circular causality is the one mentioned previously where financial incentives to reduce incidents and therefore adverse events leads to the same number of incidents or maybe an increased number.

Senge, a leader in introducing a systems approach to management wrote: *Reality is made up of circles, but we see straight lines* [23, p. 73)].

In event-oriented thinking, everything is explained by straight lines, that is, causal chains of events. The root causes are the events starting the chains of causes and effect, such as A and B in the following [27]:

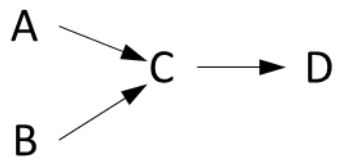

In systems thinking, alternatively, a system's behavior emerges from the structure of its feedback loops. Root causes are not the individual nodes—if it makes sense even to think of root causes. Instead, they are the forces emerging from the feedback loops.

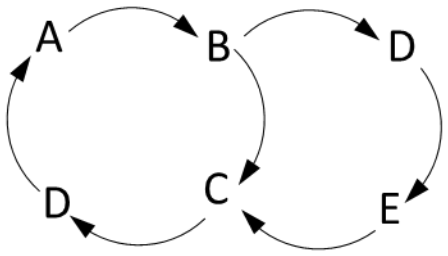

Figure 1 shows an example of a simple feedback loop in healthcare where a physician is treating a patient.

Summary

In summary, to deal effectively with complex systems, our understanding of causality—why the events occurred—has to use models that include non-linear (non-sequential) behavior and identification of *why* the events occurred. Non-linear causality may include feedback and other types of communication between components and events. In general, goal-seeking behavior includes feedback and monitoring of information about the state of the system and the components in it.

Instead of thinking of accidents as a system component (usually a human) "failure" problem, a systems theoretic approach to safety considers accidents as a control problem: An adverse event occurs because the overall system design does not have effective controls to prevent the event.

Human behavior is assumed to be impacted by the design of the system in which that behavior occurs so the most effective way to change human behavior is to redesign the system. The behavior of front-line personnel may be affected by the training they receive; the resources they are provided with, including equipment such as EHRs or the design of operational controls and displays on equipment; incentive structures, etc. Healthcare personnel may provide incorrect inputs to a machine because of

flaws in the design of the displays and controls, difficulty in obtaining the information necessary to respond safely, poor design of alerts and alarms, incorrectly designed incentive structures, and so on.

The solution then is not tp identify a human (or humans) as the cause of the adverse event and telling them to behave differently, but to change the system designs to reduce or eliminate human “failures” or errors. The goal is not to find something to blame—which limits the causes and solutions that can be identified— but to determine how to redesign the system as a whole to eliminate or reduce the occurrence of adverse events. Blame is a distraction as it usually focuses on identifying who to punish rather than how to redesign the system to eliminate the adverse events.