CHAPTER

7

ROOT-CAUSE ANALYSIS

Farrokh Alemi, Jee Vang, and Kathryn Laskey

Root-cause and failure-mode analyses are commonly performed in hospitals to understand factors that contribute to errors and mistakes. Despite the effort that healthcare professionals put into creating these analyses, few models of root causes are validated or used to predict future occurrences of adverse events. This chapter shows how the assumptions and conclusions of a root-cause analysis can be verified against observed data. This chapter builds on Chapter 3 and Chapter 4.

Root-cause analysis, according to the Joint Commission on Accreditation of Healthcare Organizations (JCAHO) (2005) is a "process for identifying the basic or causal factors that underlie variation in performance, including the occurrence or possible occurrence of a sentinel event." *Sentinel events* include medication errors, patient suicide, procedure complications, wrong-site surgery, treatment delay, restraint death, assault or rape, transfusion death, and infant abduction. Direct causes bring about the sentinel event without any other intervening event, and most direct causes are physically proximate to the sentinel event. However, the direct causes are themselves caused by root causes. Because of accreditation requirements and renewed interest in patient safety, many hospitals and clinics are actively conducting root-cause analyses.

When a sentinel event occurs, most employees are focused on the direct causes that have led to the event. For example, many will claim that the cause of medication error is a failure to check the label against the patient's armband. But this is just the direct cause. To get to the real reason, one should ask *why* the clinician did not check the label against the armband. The purpose of a root-cause analysis is to go beyond the direct and somewhat apparent causes to figure out the underlying reasons for the event (i.e., the root causes). The objective is to force one to think harder about the source of the problem. It is possible that the label was not checked against the armband because the label was missing. Furthermore, it is also possible that the label was missing because the computer was not printing the labels. Then, the direct cause is the failure to check the label against the armband and the root cause is computer malfunction. Exhorting employees

This book has a companion web site that features narrated presentations, animated examples, PowerPoint slides, online tools, web links, additional readings, and examples of students' work. To access this chapter's learning tools, go to ache.org/DecisionAnalysis and select Chapter 7.

to check the armband against the label is a waste of time if there is no label to check in the first place. A focus on direct causes may prevent a sentinel event for a while, but sooner or later the root cause will lead to a sentinel event. Inattention to root causes promotes palliative solutions that do not work in the long run. The value of root-cause analyses lies in identifying the true, underlying causes. An investigation that does not do this at best is a waste of time and resources and at worst can exacerbate the problems it was intended to fix. But how can one know if the speculations about the root causes of an event are correct?

One way to check the accuracy of a root-cause analysis is to examine the time to the next sentinel event. Unfortunately, because sentinel events are rare, one has to wait a long time to see them occur again, even if no changes were made. An alternative needs to be found to check the accuracy and consistency of a root-cause analysis without having to wait for the next sentinel incident.

Many who conduct root-cause analyses become overconfident about the accuracy of their own insights. No matter how poorly an analysis is carried out, because there is no easy way of proving a person wrong, people persist in their own fallacies. Some people are even incredulous about the possibility that their imagined causal influences could be wrong. They insist on the correctness of their insights because those insights seem obvious. However, a complex problem that has led to a sentinel event and that has been left unaddressed by hundreds of smart people for years is unlikely to have an obvious solution. After all, if the solution was so obvious, why was it not adopted earlier? The search for obvious solutions contradicts the elusiveness of correcting for sentinel events. If a sound and reliable method existed for checking the accuracy and consistency of a root-cause analysis, then employees might correct their misperceptions and not be so overconfident.

Simple methods for checking the accuracy of root-cause analyses have not been available to date (Boxwala et al. 2004). This chapter suggests a method for doing so. As before, clinicians propose a set of causes. But now several additional steps are taken. First, probabilities are used to quantify the relationship between causes and effects. Then, the laws of probability and causal diagrams are examined to see if the suggested causes are consistent with the clinician's other beliefs and with existing objective data. Through a cycle of testing model assumptions and conclusions against observed data, one improves the accuracy of the analysis and gains new insights into the causes of the sentinel event.

Bayesian Networks

Bayesian causal networks can be used to validate root-cause analyses. A *Bayesian causal network* is a mathematical model of causes and effects. It consists of a set of nodes, typically pictured as ovals, connected by directed arcs. Each node represents a mutually exclusive and collectively exhaustive set of possible events. For example, Figure 7.1 shows a Bayesian network with two nodes. The node labeled "armband legible?" has three possible values, of which exactly one must occur and two cannot coincide. These values are "no armband," "poor legibility," and "good legibility." The other node, labeled "armband checked?" has two possible values: "sure" and "not sure." A node with two possible values is called a *binary node*. Binary nodes are common in root-cause analyses.

A Bayesian network is a cyclical directed graph, meaning that you cannot start from a node and follow the arcs to arrive back to where you started. In a Bayesian network, the relationships among any three nodes can be described as having one of the following three structures: serial, diverging, or converging. Each of these graph structures can be verified through tests of conditional independence and are further explained through the examples below.

The relationship between the armband being legible and its being checked in Figure 7.1 is a direct causal relationship. Bayesian networks can also represent indirect causal relationships through the concept of conditional independence, as shown in Figure 7.2. Figure 7.2 illustrates a *serial* graph structure, in which the sentinel event is independent of the root cause given the known value for the direct cause. In this example, the root cause labeled "understaffing" is an indirect cause of the sentinel event; there is no direct arc from this root cause to the sentinel event. This means that the action of the root cause on the sentinel event is indirect, operating through an intermediate cause. The direct cause of a medication error is a fatigued nurse. The root cause, understaffing, is conditionally independent of the sentinel event given the intermediate cause. This means that if you intervene in any given instance to relieve a fatigued nurse, you can

FIGURE 7.1
A Bayesian Causal Model with a Local Probability Table

	Armband checked?	
Armband legible	Done	Not done
No armband	0%	100%
Poor legibility	80%	20%
Good legibility	99%	1%

FIGURE 7.2
Serial Example of Direct and Root Causes of Medication Error

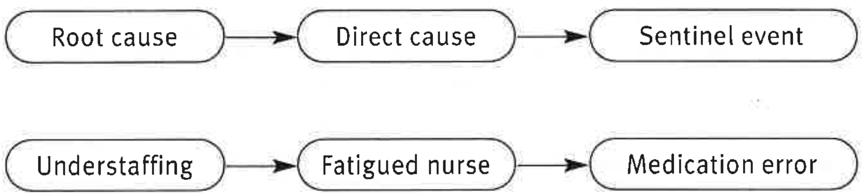

break the link from the root cause to the sentinel event, thus reducing the probability of the sentinel event to its nominal level. However, this solution is a palliative one and will not produce a long-term solution unless the root cause is addressed.

Another type of conditional independence occurs when a cause gives rise independently to two different effects, as depicted in Figure 7.3. This type of graph structure is known as *diverging*. In this example, high blood pressure and diabetes are conditionally independent given the value of weight gain, but they are correlated because of the influence of the common cause. That is, the two effects typically either occur together (when the common cause is present) or are both absent (when the common cause is absent). This type of conditional independence relationship is quite useful for diagnosing the presence of root causes that can lead to multiple independent effects that each influence different sentinel events. For example, understaffing might lead to several different intermediate causes, each of which could be a precursor of different sentinel events. If several of these precursor events were to be observed, one could infer that the understaffing problem was sufficiently severe to affect patient care. Proactive remediation could then be initiated before serious adverse medical outcomes occur.

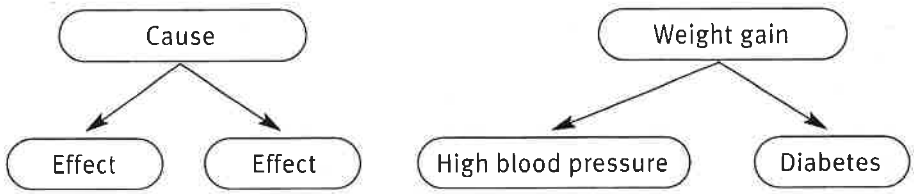

FIGURE 7.3
Conditional Independence Is Assumed in a Diverging Structure

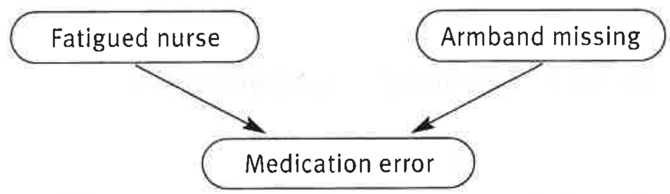

FIGURE 7.4
Two Causes Converging into a Common Effect

Figures 7.2 and 7.3 illustrate serial and diverging causal structures, respectively. As you have seen, a serial structure represents the action of an indirect causal relationship, and a diverging structure represents multiple independent effects of a single cause. In both of these cases, the two terminal nodes are conditionally independent of each other given the middle node. A different kind of causal structure, *converging*, is shown in Figure 7.4. A converging structure occurs when two different causes can produce a single effect, as when either a fatigued nurse or a missing armband can cause a medication error. Notice that in this case, the terminal nodes are not conditionally independent given the middle node. For example, if the sentinel event is known to occur, and you learn that the armband was present, this will increase the probability that the nurse was unacceptably fatigued. Likewise, if you find that the armband was missing, this will reduce the likelihood that the problem was caused by fatigue.

Data can be used, if available, to validate the graph structure of a Bayesian causal network. As noted above, when a connection is serial or diverging, the terminal nodes are conditionally independent given the intermediate node. In general, a node in a Bayesian network is conditionally independent of all its nondescendents given its parents. This general condition implies a set of correlations that should be equal to zero if the causal assumptions correct. Although it is tedious to verify all of these relationships by hand, it is straightforward to automate the verification process, and computer programs have been written to accomplish the task.

Given a causal graph, one can read off the assumed conditional independencies. Conditional independencies are identified by examining serial or diverging graphs in causal models so that removing the condition would sever the directional flow from the cause to the effect. Often, a complicated root-cause analysis can be broken into smaller components containing serial and diverging structures. If these structures are observed, and if removing the condition in these structures would sever the link between the other two nodes, then a conditional dependency has been identified. Careful examination of conditionally independent relationships is an important element of specifying and validating a Bayesian network for root-cause analyses.

Validation of Conditional Independence

Once conditional independencies have been identified, the assumptions can be verified by examining data or by querying experts. If data are available, the correlations in a serial structure between the root cause and the sentinel event should equal the correlation between the root cause and the direct cause times the correlation between the direct cause and the sentinel event:

$$R_{\text{root cause, sentinel event}} = R_{\text{root cause, direct cause}} \times R_{\text{direct cause, sentinel event}},$$

where

- $R_{\text{root cause, sentinel event}}$ is the correlation between the root cause and the sentinel event,
- $R_{\text{root cause, direct cause}}$ is the correlation between the root cause and the direct cause, and
- $R_{\text{direct cause, sentinel event}}$ is the correlation between the direct cause and the sentinel event.

In a diverging structure, a similar relationship should hold. In particular, correlation between the two effects should be equal to the multiplication of the correlation between the cause and each effect:

$$R_{\text{effect1, effect2}} = R_{\text{cause, effect1}} \times R_{\text{cause, effect2}},$$

where

- $R_{\text{effect1, effect2}}$ is the correlation between the two effects,
- $R_{\text{cause, effect1}}$ is the correlation between the cause and the first effect, and
- $R_{\text{cause, effect2}}$ is the correlation between the cause and the second effect.

If data are not available, the analyst can ask the investigative team to verify assumptions of conditional independence based on their intuitions. For example, in the serial structure in Figure 7.2, if you know that the nurse was fatigued, would information about staffing add much to your estimate of the probability of medication error? If the answer is no, then the assumption of conditional independence has been verified. Another way to ask the same questions is, does understaffing affect medication errors only by creating a fatigued nurse? In this method, the exclusivity of the mechanism of change is checked. Still another way of verifying conditional independence is by asking for estimates of various probabilities:

Question: What do you think is the probability of medication error when the nurse is fatigued?
Answer: It is higher than when the nurse is not fatigued but still relatively low.

Question: What do you think is the probability of medication error when the nurse is fatigued and working in an understaffed unit.
Answer: Well, I think understaffing leads to a fatigued nurse, but you are not asking about that, are you?

Question: No, I want to know about the probability of medication error in these circumstances.
Answer: I would say it is similar to the probability of medication error among fatigued nurses.

If conditional independence is violated, then the serial or diverging structures in the graph are incorrect. If these conditions are met, then the causal graph is correct.

Let's look at slightly more complicated sets of causes. Figure 7.5 shows four proposed causes for medication error: understaffing, fatigued nurse, vague communications, and similar medication bottles. Two root causes (understaffing and vague communications) are shown to precede the direct cause of a fatigued nurse. Removing the node labeled "fatigued nurse" would stop the flow from these two root causes to the medication

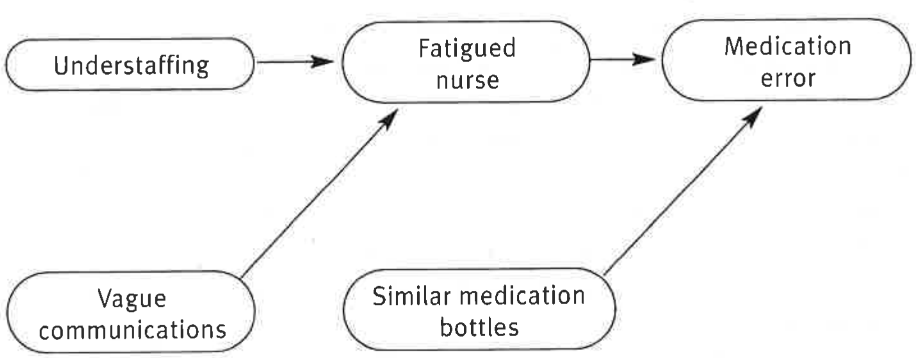

FIGURE 7.5 Four Possible Causes of Medication Error and Their Relationships

error. Therefore, a conditional independence is assumed. This assumption can be verified either through data or through experts' judgments. Assume that if you know that the nurse is fatigued, understaffing adds no additional information to the probability of medication error; therefore, this independence is verified. But suppose that even when the nurse is not fatigued, vague communications may lead to medication errors. Therefore, the assumption of the conditional independence of vague communications and medication error is not met.

Because the assumptions of the model are not correct, the causal network needs to be modified. Further exploration may indicate that vague communications, similar medication bottles, and a fatigued nurse directly affect medication errors. This example shows how verifying conditional independence could help revise root-cause analyses.

Predictions from Root Causes

The causal model behind root-cause analyses can be used to predict the probability of a sentinel event, and this probability can then be compared to the intuitions of the investigative team. The probability of the sentinel event can be calculated from each of the direct causes, and the probability of direct causes can be calculated from their root causes:

$$P(\text{Sentinel} \mid \text{event} \mid \text{Various causes}) = P(\text{Sentinel event} \mid \text{Direct causes}) \times P(\text{Direct causes} \mid \text{Root causes}) \times P(\text{Root causes}).$$

To calculate the probability of sentinel event S given a set of different unobserved (C_U) and observed causes (C_i), you can use the following formula:

$$P(S \mid C_1, C_2, \ldots, C_n) = \sum_{C_U} P(S \mid C_1, C_2, \ldots, C_n) + P(C_{U_1}) + P(C_{U_2}) + \ldots + P(C_{U_N}).$$

The above formula requires careful tracking of numerous probabilities. Because these calculations are tedious, investigative teams can use widely available software, such as Netica, to simplify the calculations. An example can demonstrate how such calculations are made using this software. Suppose Figure 7.6 shows root causes for wrong-site surgery in a hospital. First, note that the root causes listed are poor physician training and understaffing as it contributes to a fatigued nurse. These are the root causes because they are independent of the sentinel event given the various direct causes. The direct causes listed are the nurse marking the patient wrong, the surgeon not following the markings, and the patient providing

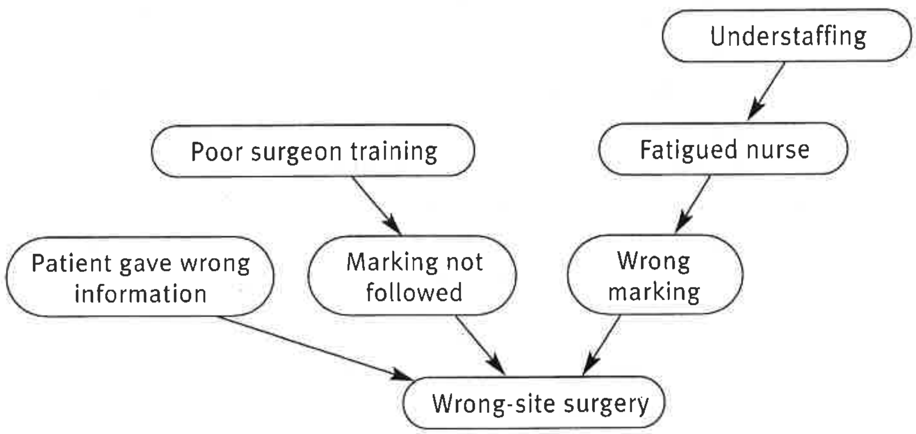

FIGURE 7.6 Root Causes for Wrong-Site Surgery

wrong information. These are direct causes because an arc connects them to the sentinel event.

Given the root-cause analysis in Figure 7.5, the next step is to estimate the probability of the various causes and effects. These probabilities are obtained by asking the expert to assess the conditional probabilities implied in the graph (Ludke, Stauss, and Gustafson 1977; Spizzichino 2001). Each node is conditioned on its direct causes. For example, to estimate the probability of having a fatigued nurse, the investigators need to ask the expert the following two questions:

1. In 100 occasions in which a unit is understaffed, what is the frequency of finding a fatigued nurse?
2. In 100 occasions in which a unit is not understaffed, what is the frequency of finding a fatigued nurse?

Obviously, estimates of probabilities from experts are subjective and therefore may be unreliable. But if experts are provided with tools (e.g., calculators, paper, pencils), brief training in the concept of conditional probabilities, and available objective data (e.g., JCAHO's reports on prevalence of various causes), and if experts are allowed to discuss their different estimates, then experts' estimates are usually accurate and reliable. These probabilities may not be accurate to the last digit, but can provide for a test of consistency. Suppose that through interviewing experts or through analyzing hospital data, the investigative team has estimated the following probabilities:

$$P(\text{Understaffing}) = .40$$
$$P(\text{Patient provided wrong information}) = .05$$
$$P(\text{Poor surgeon training}) = .12$$

P(Fatigued nurse | Understaffing) = .30
P(Fatigued nurse | No understaffing) = .05
P(Nurse marked patient incorrectly|Fatigued nurse) = .17
P(Nurse marked patient incorrectly|Not fatigued nurse) = 0.01
P(Surgeon did not follow markings|Poor training) = 0.10
P(Surgeon did not follow markings|Good training) = 0.01
P(Wrong-site surgery | Patient gave wrong information,
Nurse marked patient incorrectly and Surgeon did not follow markings)
as given as in Table 7.1

Using these estimates, you can use Netica software to calculate the probability of wrong-site surgeries when no information about any causes is present as 0.06 (see Figure 7.7 to see these calculations with Netica software). Does this seem reasonable to the investigative team? If the probability is significantly higher than what the investigative team expected, then perhaps important constraints that prevent wrong-site surgeries have been missed. If it is too low, then an important cause or mechanism by which wrong-site surgeries occur might have been missed. If the probability is in the ballpark but not exactly what was expected, then perhaps the estimated probabilities might be wrong. In any case, when there is no correspondence between the probability of the sentinel event and the investigative team's intuition, it is time to rethink the analysis and its parameters.

Other probabilities can also be calculated and compared to the experts' intuitions. Suppose on a particular unit on a particular day, you find the nurse was fatigued but the clinician was well-trained and the patient provided accurate information. Given the above estimates and the root cause in Figure 7.6, the probability of wrong-site surgery on this day is

TABLE 7.1 Estimated Probabilities of Wrong-Site Surgery Given Various Conditions

Conditions			*Probability of Wrong-Site Surgery Given Conditions*
Patient Provided Wrong Information	*Surgeon Did Not Follow Markings*	*Nurse Marked Patient Incorrectly*	
True	True	True	0.75
True	True	False	0.75
True	False	True	0.70
True	False	False	0.60
False	True	True	0.75
False	True	False	0.70
False	False	True	0.30
False	False	False	0.01

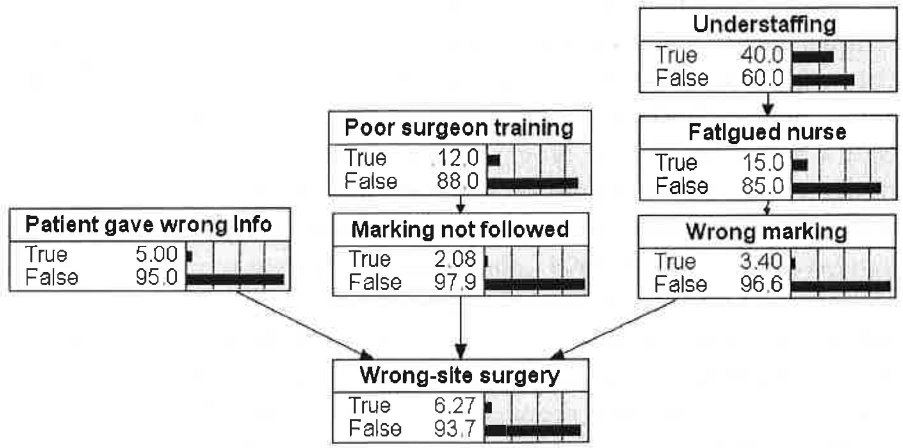

FIGURE 7.7 Application of Netica Software to Root-Cause Analysis from Figure 7.6

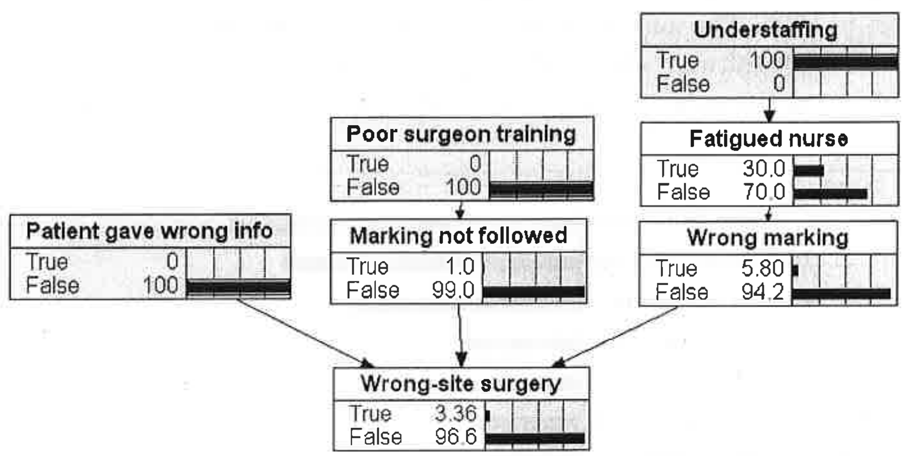

FIGURE 7.8 Probability of Wrong-Site Surgery

calculated as 0.03 using Netica (see Figure 7.8). If this corresponds to the investigative team's expectation, then the analysis is consistent and one can proceed. If not, one needs to examine why not and look for adjustments that would fit the model predictions to experienced rates.

Reverse Predictions

The Bayesian network can also be used to calculate the probability of observing a cause given that an effect has occurred. This is the reverse of how most people think about causes and effects. Most people start with a cause and want to predict the probability of the effect. Bayesian probability models allow you to do the reverse. One can start with known sentinel events and ask about the prevalence of a particular cause among them. Because causes are not as rare as sentinel events, this procedure allows one to check on the adequacy of the analysis without having to wait a long time for the reoccurrence of the sentinel event. To make things easier, JCAHO publishes the prevalence of categories of causes among sentinel events (see http://www.jcipatientsafety.org). Despite limitations (Boxwala et al. 2004), these data can be used to examine the consistency of a root-cause analysis done in one organization against the industry patterns roughly reported through JCAHO's voluntary system. A large discrepancy between observed prevalence of causes among sentinel events and assumed prevalence of causes in the investigative team's model suggest errors in assignments of probabilities as well as possible missed causes or constraints.

Netica software can calculate the prevalence of understaffing in the model of wrong-site surgeries. First, the probability of wrong-site surgery is set to 100%. The software then reports the prevalence of understaffing.

The software calculated that understaffing was present in 44 percent of wrong-site surgeries (see Figure 7.9). But is this a reasonable estimate?

FIGURE 7.9
Root Causes of Sentinel Events, 1995–2004

SOURCE: Joint Commission on Accreditation of Healthcare Organizations. 2005. "Sentinel Event Statistics." [Online information; retrieved 6/16/05]. http://www.jcaho.org/accredited+ organizations/ambulatory+care/sentinel+events/root+causes+of+sentinel+event.htm.

In contrast, JCAHO reports staffing levels to be a cause of sentinel event in fewer than 20 percent of surgeries (see Figure 7.10). Obviously, there are many reasons for a healthcare organization to differ from other aggregate data reported by JCAHO. But JCAHO's data can be used as a rough benchmark. Because the two probabilities differ considerably, these differences suggest the need to rethink the analysis.

Overview of Proposed Method for Root-Cause Analyses

Sentinel events can be reduced if healthcare organization create a blame-free environment, conduct a root-cause analysis, and take concrete actions in response to the analysis. Conduct a verifiable root-cause analysis by completing the following steps:

1. Before a sentinel event occurs, an investigative team is organized. The team should include a facilitator and a team leader. The facilitator's responsibility is to organize tasks, serve as staff to the team, and conduct team meetings in an efficient and effective method (see Chapter 6) for details). The facilitator should be trained in probability models. The leader's responsibility is to make sure that the investigation is carried out thoroughly and to provide content expertise.
2. When a sentinel event is reported, the employees closest to the incident are asked to record facts (not accusations) about the event, including what happened, who was present, where the event

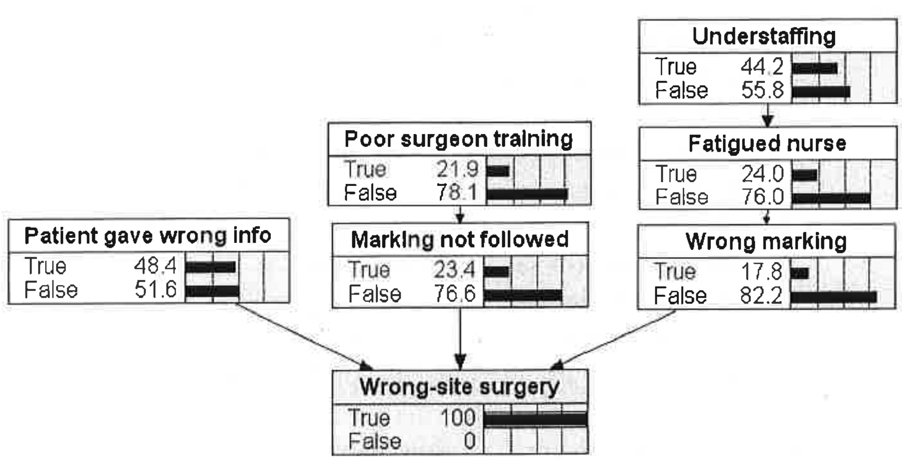

FIGURE 7.10 Prevalence of Understaffing Among Wrong-Site Surgery

occurred, when it occurred, and what the time sequence of the events that preceded the sentinel event was.

3. The investigative team meets and brainstorms the following: (1) potential causes for the incidence and (2) key constraints that would have prevented the incidence if they had been in place. Two steps are taken to make sure the listing is comprehensive. First, the framing bias is reduced by using alternative prompts. Because constraints can be thought of as reverse causes, the team should be asked to list both the constraints and causes. Furthermore, because the team is focused on conditions that led to the sentinel event, they should also be asked to examine conditions that prevented sentinel events on other occasions.
4. The facilitator interviews the investigative team or uses existing data to assign a probability to each cause and a conditional probability for each effect following the cause.
5. The facilitator checks the accuracy of the causal model and asks the investigative team to revise their model. The following steps allows one to check the accuracy or consistency of the causal model:
   a. The facilitator uses the model to predict the probability of the sentinel event. If this probability is several magnitudes higher than historical patterns or investigative team's intuitions, the facilitator seeks additional constraints that would reduce the probability of the sentinel event. If the probability is lower than historical experience or the investigative team's intuitions, the team is asked to describe additional mechanisms and causes that may lead to the sentinel event.
   b. The facilitator uses the model to calculate the prevalence of the causes among sentinel events. These data are checked against the investigative team's intuitions as well as against observed industry rates published by JCAHO.
   c. The facilitator checks that claimed root causes are conditionally independent from the sentinel event. If a root cause is directly linked to the sentinel event, the investigative team is asked to redefine the direct cause to be specific to the mechanism used by the root cause to affect the sentinel event. If few root causes have been specified, the investigative team is asked to rethink the reasons why the direct causes occur.
   d. The facilitator checks the marginal probabilities against objective data. If the probabilities do not match, the facilitator should use the objective probabilities whenever available.
6. The findings are documented. A flowchart shows the nodes for the root causes, direct causes, and sentinel events connecting to each

other with arrows. Arrows are drawn from root causes to direct causes and from direct causes to sentinel events.

Summary

Investigative teams often rely on their own intuitions for listing the root causes of a sentinel event. They rarely check the validity of their analysis. Bayesian networks can be applied to root-cause analyses to test the validity or consistency of the analyses. Real analysis should be a careful examination of facts and not a cover for wishful speculation. By creating a Bayesian network and estimating the probabilities of various events, one can scrutinize assumptions made in a root-cause analysis. In particular, one can check if important root causes have been missed, if the analysis is focused on root causes or direct causes, if the frequency of the sentinel event corresponds to expectations and experienced rates, if the prevalence of the causes of sentinel events corresponds to known rates, and if the assumptions of dependence or independence are wrong. These are not exact ways of checking the accuracy of the analysis, but these methods allow one to check the intuition of investigative teams and help them think through the implication of their analysis.

Review What You Know

1. When A causes B, B causes C, and there are no other causal relationships, what implication do these relationships have for the conditional probabilities?
2. What are the steps in conducting a root-cause analysis?
3. How can you validate the root-cause analysis? List specific ways that assumptions in root-cause analysis can be verified.
4. If a root-cause analysis of wrong-site surgery exceeds by several folds the observed frequency of wrong-site surgery, what implication does this have for the analysis?
5. What are serial, diverging, and converging structures, and which ones imply conditional independence?
6. What is meant by reverse prediction, and why is that more useful than directly predicting a rare accident?
7. In the root-cause analysis of wrong-site surgery, what is the probability of finding that the patient was responsible? If in the past you have reviewed 100 wrong-site surgeries and found that 5 percent of them were because of patient misinformation, what is the implication of this finding for the root-cause analysis?

Rapid-Analysis Exercises

1. Interview a colleague at work to analyze root causes of an adverse outcome (not necessarily a sentinel event). Make sure that you list at least three direct causes or constraints and that you include the categories suggested by JCAHO. Draw a flowchart.
2. Indicate the direct and root causes of the sentinel event in your model.
3. Give an example question that can check the conditional independence assumption associated with root causes. Make sure the question is not awkward.
4. Verify all assumptions of conditional independence in your model by interacting with your expert. Show what assumptions were checked and what assumptions were violated.
5. Estimate marginal and conditional probabilities by interviewing your expert.
6. Use Netica to estimate the probability of the sentinel event.
7. Use Netica to calculate the probability of sentinel event in at least three different scenarios (i.e., combination of causes occurring or not occurring).
8. Ask your expert if the various estimates in questions 6 through 7 are within your expert's expectations.
9. Calculate the prevalence of root causes for the sentinel event in your analysis. Compare these data to JCAHO's reports on prevalence of causes of sentinel events. Report the difference between your model assumptions and JCAHO's data.
10. Suggest how you would change the causal model to better accommodate your expert's insights. Show how your root-cause analysis changed as a consequence of the data you examined.
11. Bring your work to class.

Audio/Visual Chapter Aids

To help you understand the concepts of root-cause analysis, visit this book's companion web site at ache.org/DecisionAnalysis, go to Chapter 7, and view the audio/visual chapter aids.

References

Boxwala, A. A., M. Dierks, M. Keenan, S. Jackson, R. Hanscom, D. W. Bates, and L. Sato. 2004. "Organization and Representation of Patient Safety Data: Current Status and Issues Around Generalizability and Scalability." *Journal of the American Medical Informatics Association* 11 (6): 468–78.

Joint Commission on Accreditation of Healthcare Organizations (JCAHO). 2005. "Glossary of Terms." [Online information; retrieved 06/16/05.] http://www.jcaho.org/accredited+organizations/sentinel+event/glossary.htm.

Ludke, R. L., F. F. Stauss, and D. H. Gustafson. 1977. "Comparison of Five Methods for Estimating Subjective Probability Distributions. *Organizational Behavior and Human Performance* 19 (1): 162–79.

Spizzichino, F. 2001. *Subjective Probability Models for Lifetimes.* Boca Raton, FL: Chapman and Hall.