

26

Causal Risk Analysis

Farrokh Alemi
Georgetown University

Manaf Zargoush
ESSEC Business School

Introduction	26-1
History of Probabilistic and Causal Analysis.....	26-2
What Is a Cause?.....	26-3
Modeling Multiple Causes: Causal Diagrams.....	26-5
Modeling Multiple Causes: Probability Networks	26-5
Modeling Multiple Causes: Simultaneous Equations	26-7
Causal versus Noncausal Analysis.....	26-8
References.....	26-9

Introduction

The purpose of this chapter is to model risks of adverse events using causal probability models. To infer a cause and an effect, an experiment is needed where the cause is introduced and withdrawn and the effect is examined. Most risk analysis is not based on experimental data and uses observational data from naturally occurring variation in occurrences of causes and effects. For example, to infer that an environmental pollutant has led to an adverse event, experimental data are needed. The pollutant must be introduced in randomly chosen subjects or circumstances and the effect monitored. Because putting people at increased risk of adverse events is unethical, experimental data is seldom used in risk analysis. Instead, the analyst relies on observational data. But causal inferences from observational data are suspect. The observed increased risk of adverse events may be due to another variable not studied. It may, for example, show that an individual exposed to a pollutant has increased risk for breast cancer. The patient's breast cancer might have been caused by her smoking habit and not by the external pollutants. In analysis of observational data, it is difficult to control for alternative explanations of the adverse event. In recent years, progress has been made in making causal inferences from observational data. The purpose of this chapter is to review this progress and provide a blueprint of how causal risk analysis should be done.

Risk analysis is risk analysis, meaning that the methods of risk analysis are the same even when applied to many different application areas. In predicting the risk of adverse hospital events (e.g., wrong side surgery), risk of security breaches, risk of privacy violations, health risk of pollutants, risk of excess mortality from medications, and many other non-healthcare applications (e.g., risk of space disaster, risk of flood, risk of nuclear accidents) the methods are the same. All of these applications share a common problem structure. First, there is an adverse event of interest, for example, mortality, cardiac event, operating room fire, terrorist action, or nuclear accident (Figure 26.1).

AQ1

Second, there are a set of causes often referred to as risk factors (e.g., smoking, slippery floors, new medication). Each of these causes could lead to the adverse event. Third, the impact of the causes (risk factors) may be moderated by various circumstances, which we refer to as moderating factors. Some factors may enhance the impact of the causes and others may prevent it. For example, in environmental risk analysis, the moderating factors are the way the pollutant spreads in the environment and is absorbed




FIGURE 26.1 Common structures of risk analysis problems.

by the population. Analysts need to model how the dispersion and absorption processes mitigate the impact of the pollutant. In patient safety, patient falls may be caused by slippery floors. Some patients wake up at night and would like to go to bathroom and fall in the process. The impact of slippery floors may be mitigated by lowering the height of the patient beds at nights and thereby enabling patients more control over their movements. For still another example, new medications may cause unanticipated excess mortality. In analysis of impact of medication on excess cardiac events, the impact may be mitigated by the patient's genetic predispositions. While the application of risk analysis varies, the structure of the problem at hand is similar and the methods of analysis are the same.

To illustrate the methods of causal risk analysis, we show how it can be applied to the determination of risk of excess cardiac events caused by the use of rofecoxib (commonly known as Vioxx). Merck & Co. discovered and marketed Vioxx for treatment of osteoarthritis, acute pain conditions, and dysmenorrhea. The Food and Drug Administration (FDA) approved the use of Vioxx on May 20, 1999. It quickly gained a large market share, particularly among physicians treating patients with arthritis. A series of small studies of the long-term impact of Vioxx raised concerns that it might increase cardiac events. Two large-scale prospective studies were undertaken (Bombardier et al. 2000, Bresalier et al. 2005), one of which was stopped before the end of the study because preliminary analysis indicated that experimental patients were at increased risk of cardiac events. Finally, a very large study of patients in Kaiser Permanente involving 2.3 million person-years of follow up indicated that Vioxx increased the risk of cardiac events (Graham et al. 2005). This is the study that we will use throughout this chapter because it made causal inferences from observational data within electronic health records. On September 30, 2004, Merck voluntarily withdrew Vioxx from the market. Vioxx was the largest drug to be withdrawn from the market. In the year before withdrawal, Merck had sales revenue of US\$2.5 billion from Vioxx (Reuters 2006). At the time of withdrawal, Merck and Co. also disclosed that their earlier unpublished studies had shown an increased risk of cardiac events, but that findings were not significant. Merck and Co. was criticized for failing to disclose data on adverse events of Vioxx to FDA. Congressional hearing led to interest in development of methods of monitoring impact of medications after FDA approvals, the so-called post-release monitoring. Congress passed a law empowering FDA to create the sentinel database for monitoring impact of medications. Later, the American Recovery Act funded the widespread use of electronic health records, which increased the availability of observational data on impact of medications. Thus, it made it easier to analyze the impact of medications from observational data within electronic health records. We selected the Vioxx case as an example for demonstrating causal risk analysis because it is a model for how risk of medications can be analyzed and because it demonstrates the issues and difficulty associated with causal inferences from observation data.

History of Probabilistic and Causal Analysis

Probabilistic risk analysis started in 1967 in the space industry, where, except for a short period of time, it has been the standard approach for analysis of risks associated with shuttle flights (Colglazier and Weatherwax 1986, Bell and Esch 1989, Planning Research Corporation 1989, Science Applications International Corporation 1989, Pate-Cornell and Fischbeck 1994, Hoffman et al. 1998, Cooke and Jager 1998). In nuclear safety, probabilistic risk analysis has been used to assess reactor safety (Environmental Protection Agency 1976, Union of Concerned Scientists 1977, Kemeny 1979, Rogovin

and Frampton 1980, Kaplan and Garrick 1981). Probabilistic risk analysis has been applied to cyber terrorism (Taylor et al. 2002), earthquake predictions (Chang et al. 2000), floods and coastal designs (Voortman et al. 2002, Kaczmarek 2003, Mai and Zimmermann 2003), environmental pollution (Slob and Pieters 1998, Moore et al. 1999), waste disposal (Garrick and Kaplan 1999, Ewing et al. 2004), and environmental health (Cohen 2003, Sadiq et al. 2003). Applications to healthcare problems have been less common (Marx and Slonim 2003, Wreathall and Nemeth 2004, Alemi 2007). Probabilistic risk analysis has been applied to allergic reaction to computational tomography contrast agent (Abduljawad 2007), medication errors in a hospital (Hover and O'Donnel 2007), falls in an assisted living home (Song 2007), and infant mortality (Yang and Kitsantas 2007). In health care, many probabilistic risk analysis studies are reported under other headings, for example as statistical analysis of adverse effects of new medications.

AQ2

In recent decades, probabilistic analysis has evolved into causal analysis. The early causal analysis focused on path analysis and simultaneous equations (Wright 1921). Judea Pearl used probabilistic networks for causal analysis (Pearl 1986, 1988, Rebane and Pearl 1987). More recent works have highlighted how causal graphs, probability networks, and simultaneous equations are interrelated and can be used interchangeably for causal analysis (Pearl 2000).

The application of causal analysis to risk analysis is rare. Cox reports one of the earliest applications of causal risk analysis to the understanding of the impact of environmental pollutants (Cox 2001).

What Is a Cause?

For more than a millennium, scientists have been clarifying what can legitimately be considered a cause of an event. Aristotle, for example, speaks of classes of causes more than a millennium ago. Furthermore, many psychologists believe that humans think through causal reasoning (Sloman 2005). We say that spark causes forest fires, smoking causes lung cancer, or Vioxx leads to excess cardiac events. Despite widespread use of the term, despite a long history of active research in this area, a great deal of confusion remains about what is a cause. When we ask people to list the causes of excess cardiac events, we were surprised by the answers they gave. Much of what they said could not be considered a cause at all. Some listed the goals of the healthcare system, e.g., one person said that she wanted to reduce excess cardiac events in order to improve quality of care. Improving quality of care maybe a reason why we reduce excess cardiac events but it is not a cause of it.

When asked what causes cardiac events, some responded “being male.” While it is true that men have more cardiac events, it is hard to imagine the Y-chromosome plays an active role in creating cardiac events. Being male has a probabilistic but not a causal relationship to cardiac events.

So many people make errors in listing causes of cardiac events that it occurred to us that we need to distinguish a cause from other concepts such as goals, reason, and motivation. We need to define a cause so that independent observers can judge if claims of causality are plausible.

At the simplest possible level, we can imagine a world in which there is one cause and one effect (Figure 26.2). The cause and the effect are shown in two separate nodes and an arrow shows the direction of the influence. Given this simple cause and effect diagram, the question is what we need to show that supports our claim that one event causes the other. For example, what evidence is necessary to claim that the medication Vioxx led to excess cardiac events.




FIGURE 26.2 A simple causal diagram.

First, to clarify what is a cause, we need to be clear about what is an event. Events are physical or mental changes that have a start and an end. In this sense, taking the Vioxx medication is considered an event. This event causes another event called “cardiac event.” Causes link events to each other. It is easy to talk of events causing other events because an event marks a change. It is hard to say the same about non-events, e.g., physical object. Physical objects do not typically cause changes because they seldom change themselves. It is hard to imagine that hospital walls cause cardiac events (even though some hospitals have more cardiac events than others) because these walls do not change and therefore cannot be considered events. The very definition of events and the association between cause and change leads to the obvious statement that causes are temporary events. They are absent at one point in time and present at another point in time. The statement that being male is a cause of cardiac events does not sound reasonable to us because it is always present, even when there are no cardiac events. On the other hand, the statement that smoking is a cause of cardiac events could be reasonable because a person might change his smoking habit and see the consequence of it in terms of lower cardiac events. A legitimate candidate for causes of cardiac events needs to be a temporary event associated with cardiac events.

Second, we expect causes and effects to be associated with each other. The relationship between a cause and effect might be measured in terms of a correlation or the conditional probability of observing the effect given the presence of the cause. It is important to point out that the relationship between a cause and an effect is often not deterministic. A cause does not always lead to the same effect. A previous heart attack increases the risk of cardiac events but it is not true that everyone who has a history of heart attacks will have additional cardiac events.

Third, causes describe a mechanism where an event leads to another (Susser 1991, Morabia 2005). A series of events can cause each other. The process continues until the effect emerges. The original event can be called the root cause. The cause immediately preceding the effect is called direct cause. The more the mechanism from the cause to the effect is clear, the stronger the claim that one has found a cause and not merely two associated events. We are more likely to believe that Vioxx leads to excess cardiac events if it is clear how the active ingredient in Vioxx creates the heart attack.

Fourth, another point that should be obvious is that causes must precede the effect (in Figure 26.2, the cause node should precede the effect node and the direction of the arrow should be from cause to the effect). One cannot talk of Vioxx causing a cardiac event, if the heart attack preceded taking the medication. Nor can one talk of cardiac events causing the taking of Vioxx. Even if aspirin is taken as a preventive measure, still it is the aspirin that is the cause of the risk reduction. One cannot claim that the medication was taken because of a future event but the medication led to the reduced risk of a future event.

Fifth, and perhaps most important, if one claims that an events causes the other, then the counterfactual should also be true. A counterfactual refers to the claim that for individuals that the effect has occurred, the effect would not have happened if the cause was not present. This is perhaps key criterion that separates a causal event from other types of events. For our example, where we claim that Vioxx leads to excess cardiac events we also need to show that for patients who had a heart attack, they would have not had the heart attack if it were not for the Vioxx. Here is a scenario under which a counterfactual does not hold. Suppose that the intervention occurs only among patients who are severely ill. For these patients, removing the cause will not lead to improvements. It is likely that they would die anyway. The intervention was a shot in the dark, it could help but it could not make the situation worse. For example, suppose that we give Vioxx to patients who have such severe arthritis pain that they cannot function well. These patients tend to be much older and also at increased risk of cardiac events in any case. If we see excess cardiac events among patients who took Vioxx compared to patients who took an alternative medication, it is now difficult to attribute the excess cardiac events to Vioxx. In fact, in these circumstances, the counterfactual claim is no longer valid. It is not true that if we had not given the medication these patients would not have had the cardiac event anyway. Blaming Vioxx for death of these patients is akin to blaming the fireman for the fire. It blames a treatment of last resort for outcomes that would have occurred anyway. Thus, the verification and testing of counterfactual claims is central to causal risk analysis.




FIGURE 26.3 A causal diagram showing side effects of Vioxx medication.

Modeling Multiple Causes: Causal Diagrams

Most events have multiple causes, and one task of the risk analyst is to sort out which of the many causes is the central reason for the occurrence of the adverse event. At the same time, while one cause is too few, too many causes could make the analysis much more difficult. It is possible to claim that everything has an effect on something else, and in the end—despite several degrees of separation—everything is a cause. This type of thinking is not very helpful. It is important to limit the analysis of causes to events that have a large impact and are under the control of the decision maker. From a practical perspective, a select few causes are likely to explain most of the effect. Being struck by lightning can lead to cardiac events, but the event is so rare as not to be relevant in most causal analysis.

Simple cause–effect relations are known from intuition or can be verified using the five criteria we presented earlier, but more complex situations require modeling to track the many interactions among the causes of various events. There are three ways to do so: causal diagram, probabilistic network, and simultaneous equations.

A causal graph shows events as nodes and the cause and effect relationship as a directed link between the nodes. For example, the graph in Figure 26.3 shows that patient's arthritis pain leads to the prescription of Vioxx, which in turn leads to both reduced pain and the side effect of excess cardiac events.

In a causal diagram, every arrow specifies a cause and effect relationship. Missing arrows also tell a great deal. The absence of an arrow shows the lack of a direct cause and effect. As we will see in a later section, each missing arrow shows an assumption of conditional independence. For example, in Figure 26.2, there is no direct link between arthritis pain and cardiac events. This suggests that, if it were not for the Vioxx, arthritis pain and cardiac events were independent events.

Modeling Multiple Causes: Probability Networks

A probability network goes a step beyond causal diagrams by assigning a specific probability to each direct cause and by using the calculus of probabilities to measure the impact of a change somewhere




FIGURE 26.4 A possible cyclic causal diagram.

in the causal diagram on the entire network of nodes. As before, the nodes represent events, arrows indicate causality (now measured in terms of conditional probabilities), and missing arrows encode conditional independencies between the events. A Bayesian probability network is assumed to be a directed and acyclic graph. By “directed,” we mean that any two events that are related to each other have a direction of influence on each other. By “acyclic,” we mean that it is not possible to start at any node on the graph and return to the same point by going through other nodes in the analysis. In this sense, Bayesian probability networks cannot model causes that feed into themselves. For example, poor care leads to lower market share (see Figure 26.4). Lower market share leads to lower volume of services, which in turn leads to less practice and thus more poor care.

Strictly speaking, probability network are not causal networks unless we restrict the variables listed as causes and consequences.




FIGURE 26.5 Possible causal relations among three events.

A key component of probability networks is how multiple causes interact. One way to understand this is to examine three possible ways that three nodes can be represented in a causal relationship. Between any three events, three possible relationships are possible: serial, diverging, or converging structures.

A serial structure (right side of Figure 26.5) identifies the root and the direct cause of the adverse event. Figure 26.3 also contains an example of a serial structure. In this example, the root cause “arthritis pain” is an indirect cause of cardiac events. There is no direct arc from the root cause to the adverse event. This means that the impact of the root cause on the adverse event is indirect, operating through an intermediate cause. That is, the direct cause of cardiac event is the medication Vioxx. If one removes the Vioxx medication, then there is no longer any relationship between arthritis and cardiac events. A serial graph structure can be identified by examining conditional independence. In a serial structure, the root cause (left node) is conditionally independent of the adverse event (right node) given the direct cause (middle node). In Figure 26.3, the serial drawing implies that among patients who received Vioxx, arthritis pain and cardiac events are independent from each other. This assumption can be tested in the data to verify that the sequence assumed in the causal model is accurate.

A diverging structure (left side of Figure 26.5) depicts a situation where one cause leads to multiple effects. Whether a common cause is leading to multiple effects can also be detected by examining conditional independence of the effects. As can be expected, the common cause leads to the fact that the multiple effects are correlated and dependent on each other. Furthermore, these effects are conditionally independent of each other given the common cause. Figure 26.3 also shows an example where Vioxx is leading to two effects: the intended reduction of pain and the unintended cardiac event. In this example, reduction of pain and cardiac events are independent from each other among patients who took Vioxx.

Finally, middle of Figure 26.5 shows a converging causal structure. In these structures, two causes lead to the same effect. In a converging or common effect, the causes are independent from each other except when conditioned on the effect. Another and perhaps more intuitive way of saying the same thing is that when the effect has been observed and we know that one cause is not present, then our estimate that the alternative cause is present is increased. Figure 26.6 gives an example where patients are assumed to die through two separate and independent causes: severity of cardiac illnesses and severity of other diseases (e.g., cancer). If we know that the patient has died and we know that there was no cardiac event, then we would think it very likely that the patient has died from the other cause, i.e., from other diseases such as cancer. Likewise, if we know that the patient has cancer and the patient has died, then it is likely that the patient did not die from cardiac events.




FIGURE 26.6 Two different causes of mortality.

Once a causal probability network has been constructed, it can be used to predict the occurrences of the adverse event. The way causal networks simplify the calculation of probability of an event is through the use of conditional independence among the events. The probability of the adverse events is assumed to be a function of its direct causes and nothing else. In this fashion, the probability of each node can be written as a function of few causes. The probability of the adverse event, S , can be predicted from the presence of various causes. Even if a cause has not yet occurred but has a chance of occurring (unobserved causes), this information can be used in the predictions. The probability of the adverse event given a set of different unobserved (C_U) and observed causes (C_i) can be calculated using the following formula:

$$P(S | C_1, C_2, \dots, C_n) = \sum_{C_U} P(S | C_1, C_2, \dots, C_n) P(C_{U1}) P(C_{U2}) \dots P(C_{UN})$$

Each of the direct causes in the above formula can be predicted by a different set of events leading to these causes. In this fashion, a large complicated causal diagram could be easily distilled into a set of probabilistic relationships, which can then be used to estimate the probability of the adverse event.

Modeling Multiple Causes: Simultaneous Equations

A third approach to causal analysis is to use simultaneous equations. In this approach, the effect is the dependent variable in the equation and the direct causes are the independent variables. For each node, a different equation is written showing the relationship of the events that directly lead to it. For example, we may start with the adverse event and write an equation that predicts the adverse event from its direct causes. Then, an equation is written to predict the direct causes of the adverse event from events that lead to these causes. The process is continued until every node (except nodes without a direct cause) has an associated equation. Simultaneous solutions of these equations allows one to predict how changes in one event affect the frequency of the adverse event.

Figure 26.7 shows a causal diagram for how Vioxx leads to excess cardiac events. Figure 26.7 assumes that patients who have arthritis pain are prescribed Vioxx, in order to reduce their pain. Some patients develop cardiac events. These patients may also have cardiac events because of their underlying unrelated illness (shown as a direct link between patient's severity of cardiac illness and occurrence of cardiac events). In Figure 26.7, cardiac events are assumed to cause mortality. In addition, mortality may be caused by other noncardiac events (e.g., cancer). Given the structure of Figure 26.7, a number of equations can be written.

We start with equations relevant for predicting mortality (note that in all of the following equations the English letters correspond to the letters within Figure 26.7 and all of the Greek letters correspond to constants that can be estimated from the data):

$$M = \alpha + \beta C + \gamma O$$




FIGURE 26.7 A model of relationship between Vioxx and cardiac events.

This equation says that mortality is a function of patients with cardiac, C , and other, O , diseases. Next, we write the equations for predicting cardiac events:

$$C = \delta + \epsilon V + \theta S$$

This equation states that cardiac events are a function of the patient's severity of cardiac illness and the use of Vioxx. This is the core relationship of interest. The equation, in essence, shows what percent of cardiac events can be attributed to the use of Vioxx. Next, we write the equations for frequency of observing other disease:

$$O = \vartheta + \mu Z$$

This equation states what might be obvious, that frequency of other diseases is a function of severity of other illnesses in the patient. Finally, an equation links the use of Vioxx to the level of pain among patients with arthritis:

$$V = \pi + \rho P$$

Simultaneous solution of the above five equations allows one to detect whether Vioxx has led to excess mortality.

Causal versus Noncausal Analysis

Causal risk analysis starts with a hypothesized model of the relationships anticipated in the data. For us this is Figure 26.7. Within this model, we want to estimate how much of excess mortality can be attributed to the use of Vioxx. The first step is to verify that the causal relationships depicted in the model meet the criteria discussed earlier (i.e., causes must change over time, there should be association between cause and effect, causes must precede effects, and the counterfactual statements should be true). All of these conditions are easily verified except the counterfactual claim. To verify the counterfactual claim, we need to show that patients who had a cardiac event would not have had the event if it were not for the Vioxx. We can predict what might have happened to these patients by examining the severity of their illness. This is simply the situation in Figure 26.6. The predictions of what might have happened is done through the analysis of simultaneous equations described in the previous section (through use of path analysis) or through use of Bayesian probability networks.

What was actually done in the study by Graham and colleagues was different. They classified patients in 10 cardiac risk categories through use of patients' diagnoses, previous medication use, and unexplained mortality. They then conducted a matched case study where patients within the same cardiac risk category were compared to each other based on whether the patient used Vioxx or alternative medications. The study showed that Vioxx led to increased risk of cardiac events compared to the alternative medication.

The two approaches may lead to different conclusions. One reason for the difference is that the causal analysis takes into account risk not related to cardiac events. It does so because elderly patients (most of the patients who received high-dose Vioxx were elderly) present with multiple illnesses. Ignoring mortality through other causes besides cardiac risks would ignore a major risk factor for mortality among these patients. Even though mortality is not a cause of cardiac events, information about patients' mortality from other diseases changes the probability that they had cardiac events (in the context of earlier discussion this is a converging structure with two causes of a common effect).

The accuracy of both approaches relies on the ability to predict cardiac events from patient's severity of illness. Unfortunately, Graham and colleagues do not publish the accuracy of their severity index. They do mention that the highest and the lowest severity scores differed by 12 fold. But this does not

tell us what percentage of variance in cardiac events are explained by the severity index. In previous studies of myocardial infarction, several commercial severity indices did not correctly predict more than 76% of cases (Alemi et al. 1990). Since in severity-adjusted outcome studies, the adverse outcomes not explained by the patient's severity is attributed to quality of care, this suggests that 24% of cases may have received poor care. This is an unreasonably high percentage that is not supported by medical record reviews of quality. Therefore, it is reasonable to expect that the severity index may not have been sufficiently accurate to correctly test for the counterfactual claim.

We do not expect that the objections we have raised should change the faith of conclusions arrived at about Vioxx. This issue has been settled through a number of experimental studies. The study by Graham and colleague's was based on observational data, from which it is difficult to make causal inferences. The procedure described here was intended to improve the methods of Graham and colleagues. In particular, we hope that, in the future, investigators will use causal risk analysis to make more informed decisions from observational data.

Like the Vioxx study, risk analysis, in general, can benefit from use of causal models. The analyst and the policymakers make causal interpretation of the risk analysis whether causal assumptions are explicitly discussed in the report. Since these studies are interpreted in causal terms, it is important to use causal methods for risk analysis.

References

- Abduljawad, Y. 2007. Root cause analysis of allergic reaction to computational tomography (CT) contrast agent. *Qual. Manag. Health Care* 16(2). AQ3
- Alemi, F. 2007. Probabilistic risk analysis is practical. *Qual. Manag. Health Care* October–December; 16(4):300–310.
- Alemi, F., J. Rice, and R. Hankins. 1990. Predicting in-hospital survival of myocardial infarction. A comparative study of various severity measures. *Med. Care* September;28(9):762–775.
- Bell, T. E. and K. Esch. 1989. The space shuttle: A case study of subjective engineering. *IEEE Spectr.* June;42–46. AQ4
- Bombardier, C., L. Laine, A. Reicin, D. Shapiro, R. Burgos-Vargas, B. Davis, R. Day, M. B. Ferraz, C. J. Hawkey, M. C. Hochberg, T. K. Kvien, and T. J. Schnitzer; VIGOR Study Group. 2000. Comparison of upper gastrointestinal toxicity of rofecoxib and naproxen in patients with rheumatoid arthritis. VIGOR Study Group. *N. Engl. J. Med.* November;343(21):1520–1528, 2 page following 1528.
- Bresalier, R. S., R. S. Sandler, H. Quan, J. A. Bolognese, M. Stat, K. Horgan, C. Lines, R. Riddell, D. Morton, A. Lanas, M. A. Konstam, and J. A. Baron. 2005. Cardiovascular events associated with rofecoxib in a colorectal adenoma chemoprevention trial. *N. Engl. J. Med.* 352(11):1092–1102.
- Chang, S. E., M. Shinozuka, and J. E. Moore. 2000. Probabilistic earthquake scenarios: Extending risk analysis methodologies to spatially distributed systems. *Earthq. Spectr.* 16(3):557–572.
- Cohen, B. L. 2003. Probabilistic risk analysis for a high-level radioactive waste repository. *Risk Anal.* October;23(5):909–915.
- Colglazier, E. W. and R. K. Weatherwax. 1986. Failure estimates for the space shuttle. In: Abstracts for Society for Risk Analysis Annual Meeting, November 9–12, 1986, Boston MA, p. 80.
- Cooke, R. and E. Jager. 1998. A probabilistic model for the failure frequency of underground gas pipelines. *Risk Anal.* August;18(4):511–527.
- Cox, L. A. 2001. *Risk Analysis: Foundations, Models and Methods. Series: International Series in Operations Research & Management Science.* Boston, MA: Kluwer.
- Environmental Protection Agency. 1976. Reactor safety study oversight hearings before the subcommittee on energy and the environment of the committee on interior and insular affairs. In *House of Representatives, 94th Congress, 2nd Session, Serial No. 84-61, Washington DC, June 11.*

- Ewing, R. C., C. S. Palenik, and L. F. Konikow. 2004. Comment on "Probabilistic risk analysis for a high-level radioactive waste repository" by B. L. Cohen in *Risk Analysis*, volume 23, 909–915. *Risk Anal.* December;24(6):1417–1419.
- Garrick, B. J. and S. Kaplan. 1999. A decision theory perspective on the disposal of high-level radioactive waste. *Risk Anal.* October;19(5):903–913.
- Graham, D. J., D. Campen, R. Hui, M. Spence, C. Cheetham, G. Levy, S. Shoor, and W. A. Ray. 2005. Risk of acute myocardial infarction and sudden cardiac death in patients treated with cyclo-oxygenase 2 selective and non-selective non-steroidal anti-inflammatory drugs: Nested case-control study. *Lancet* February 5–11;365(9458):475–481.
- Hoffman, C. R, R. Pugh, and F. M. Safie. 1998. Methods and techniques for risk prediction of space shuttle upgrades. In *AIAA Conference*, April 20–23, 1998, Long Beach, CA.
- AQ5 Hover, C. and L. O'Donnel. 2007. Causes of medication error. *Qual. Manag. Health Care* 16(2).
- Kaczmarek, Z. 2003. The impact of climate variability on flood risk in Poland. *Risk Anal.* June;23(3):559–566.
- Kaplan, S. and B. Garrick. 1981. On the quantitative definition of risk. *Risk Anal.* 1:11–27.
- Kemeny, J. 1979. Report of the President's Commission on the Accident at Three Mile Island, Washington, DC.
- Mai, S. and C. Zimmermann. 2003. Risk analysis-tool for integrated coastal planning. In *Proc. of the 6th Int. Conf. on Coastal and Port Engineering*, Colombo, Sri Lanka.
- Marx, D. A. and A. D. Slonim. 2003. Assessing patient safety risk before the injury occurs: An introduction to sociotechnical probabilistic risk modelling in health care. *Qual. Saf. Health Care* December;12(Suppl 2):ii33–ii38.
- Moore, D. R. J., B. E. Sample, G. W. Suter, B. R. Parkhurst, and T. R. Scott. 1999. A probabilistic risk assessment of the effects of methylmercury and PCBs on mink and Kingfishers along East Fork Poplar Creek, Oak Ridge, Tennessee, USA. *Environ. Toxicol. Chem.* 18(12):2941–2953.
- Morabia, A. 2005. Epidemiological causality. *Hist. Philos. Life Sci.* 27(3–4):365–379.
- Pate-Cornell, M. E. and P. S. Fischbeck. 1994. Risk management for tiles of the space shuttle. *Interfaces* 24(1):64–86.
- Pearl, J. 1986. Fusion, propagation, and structuring in belief networks. *Artif. Intell.* 29(3):241–288.
- Pearl, J. 1988. *Probabilistic Reasoning in Intelligent Systems*. San Mateo, CA: Morgan Kaufmann.
- Pearl, J. 2000. *Causality: Models of Reasoning and Inference*. Cambridge, U.K.: Cambridge University Press.
- Planning Research Corporation. 1989. Independent assessment of shuttle accident scenario probabilities for Galileo mission and comparison with NSTS program assessment.
- Rebane, G. and J. Pearl. 1987. The recovery of causal poly-trees from statistical data. In *Proceedings, 3rd Workshop on Uncertainty in AI*, Seattle, WA, pp. 222–228.
- Reuters 2006. Merck sees slightly higher 2007 earnings. *New York Times*, p. A1. December 7, 2006. <http://www.nytimes.com/2006/12/07/business/07drug.html?ex=1323147600&en=19d27b5814f1c1e8&ei=5088&partner=rssnyt&emc=rss>
- Rogovin, M. and G. T. Frampton. 1980. Three Mile Island, a Report to the Commissioners and to the Public, Government Printing Office, Washington, DC.
- Sadiq, R., T. Husain, B. Veitch, and N. Bose. 2003. Distribution of arsenic and copper in sediment pore water: An ecological risk assessment case study for offshore drilling waste discharges. *Risk Anal.* December;23(6):1309–1321.
- Science Applications International Corporation. 1995. *Probabilistic Risk Assessment of the Space Shuttle*. Washington, DC: Center for Aerospace Information.
- Slob, W. and M. N. Pieters. 1998. A probabilistic approach for deriving acceptable human intake limits and human health risks from toxicological studies: General framework. *Risk Anal.* December;18(6):787–798.
- Sloman, S. 2005. *Causal Models: How People Think about the World and Its Alternatives*. Oxford, U.K.: Oxford University Press.

- Song, L. 2007. Tutorial on analysis of causes of falls within an assisted living setting. *Qual. Manag. Health Care* 16(2).
- Susser, M. 1991. What is a cause and how do we know one? A grammar for pragmatic epidemiology. *Am. J. Epidemiol.* April 1;133(7):635–648.
- Taylor, C., A. Krings, and J. Alves-Foss. 2002. Risk analysis and probabilistic survivability assessment (RAPSA): An assessment approach for power substation hardening. In *Proc. ACM Workshop on Scientific Aspects of Cyber Terrorism*, Washington, DC.
- Union of Concerned Scientists. 1977. *The Risk of Nuclear Power Reactors: A Review of the NRC Reactor Study*, WASH-1400. Cambridge, MA: Union of Concerned Scientists.
- Voortman, H. G., P. van Gelder, and J. K. Vrijlin. 2002. Risk-based design of large-scale flood defense systems. In *28th International Conference on Coastal Engineering*, Cardiff, Wales.
- Wreathall, J. and C. Nemeth. 2004. Assessing risk: The role of probabilistic risk assessment (PRA) in patient safety improvement. *Qual. Saf. Health Care* June;13(3):206–212.
- Wright, S. 1921. Correlation and Causation. *J. Agri. Res.* 20(7):557–585. http://www.ssc.wisc.edu/soc/class/soc952/Wright/Wright_Correlation%20and%20Causation.pdf
- Yang, L. and Y. Kitsantas. 2007. Monitoring causes of infant mortality. *Qual. Manag. Health Care* 16(2).

AUTHOR QUERIES

- [AQ1] Please check the inserted citation for Figure 26.1.
- [AQ2] Reference citation Zimmermann (2003) has been changed to Mai and Zimmermann (2003). Please check.
- [AQ3] Please provide page number for reference Abduljawad (2007).
- [AQ4] Please provide volume number for reference Bell and Esch (1989).
- [AQ5] Please update the references Hover and Donnel (2007), Song (2007), Yang and Kitsantas (2007).

