ORGANISATIONAL MATTERS

Assessing risk: the role of probabilistic risk assessment (PRA) in patient safety improvement

J Wreathall, C Nemeth

Qual Saf Health Care 2004;13:206–212. doi: 10.1136/qshc.2003.006056

Morbidity and mortality due to “medical errors” compel better understanding of health care as a system. Probabilistic risk assessment (PRA) has been used to assess the designs of high hazard, low risk systems such as commercial nuclear power plants and chemical manufacturing plants and is now being studied for its potential in the improvement of patient safety. PRA examines events that contribute to adverse outcomes through the use of event tree analysis and determines the likelihood of event occurrence through fault tree analysis. It complements tools already in use in patient safety such as failure modes and effects analyses (FMEAs) and root cause analyses (RCAs). PRA improves on RCA by taking account of the more complex causal interrelationships that are typical in health care. It also enables the analyst to examine potential solution effectiveness by direct graphical representations. However, PRA simplifies real world complexity by forcing binary conditions on events, and it lacks adequate probability data (although recent developments help to overcome these limitations). Its reliance on expert assessment calls for deep domain knowledge which has to come from research performed at the “sharp end” of acute care.

and their consequences, if only on a relative scale—for example, events are frequent v unlikely. Even though there are many types of FMEAs, the analyses tend to focus on hardware and software failures. RCA relies on investigator experience to identify the possible contributors to adverse events, yet it is often performed by staff members who are not familiar with clinical issues. The RCA method classically has no clear rule as to when to stop because a “deeper” cause can always be found. RCA searches are often within a single level instead of across the levels of an organization, look for a single cause or set of causes instead of what are usually multiple contributing causes, and can be blind to events that intervene across organizational boundaries and extended time periods. Both FMEA and RCA methods are gradually being extended to include more complexity and completeness.

PRA is used systematically to identify and review all of the factors that can contribute to an event, including equipment failure, human erroneous actions, departments or units involved, and their interactions. It is performed to understand the causes that contribute to a class of undesirable outcomes and determines how to reduce, eliminate, or improve barriers to them. The approach is frequently used in technology driven industries such as chemical manufacturing, offshore drilling and production facilities, and aviation. However, PRA is best known as a safety assessment tool in the commercial nuclear power industry. The information produced by PRA provides a basis for resource allocation decisions and evaluation of performance goals in terms of safety related criteria.

This paper discusses PRA tools and process and examines its strengths, limitations, and relevance to patient safety. It also describes applications in other fields and identifies issues that are of interest to healthcare applications. Further background information on the evolution of thought about “medical error” and “patient safety” is available from Nemeth.⁵

PRA MODELS AND MODELING

The PRA process begins with the identification of a bad event. “Bad” refers to an event that is of concern to safety but does not necessarily involve real harm. The bad event that is selected is important, as all the subsequent searches for

Morbidity and mortality resulting from what are frequently referred to as “medical errors” compel a better understanding of health care as a system. Human factors methods¹ are designed to understand complex systems. One such method is probabilistic risk assessment (PRA) (also known as probabilistic safety assessment (PSA) or quantitative risk assessment (QSA)). PRA is a “top down” analytical process that can be used to identify the cause, consequence, and frequency of adverse outcomes in a system.

Efforts to improve patient safety already use two analytical processes: failure modes and effects analysis (FMEA) and root cause analysis (RCA). FMEA uses a table format to identify system components, to identify the ways that different elements in a system can fail, and to estimate how the failures might affect the system. Recent studies^{2–4} have described the use of FMEA in healthcare applications. Many FMEAs provide approximate failure probabilities

See end of article for authors' affiliations

Correspondence to: J Wreathall, John Wreathall & Co Inc, 4157 MacDuff Way, Dublin, OH 43016-9578, USA; john@wreathall.com

Accepted for publication 29 November 2003

causes will focus on it. A bad event could involve serious harm to a patient—for example, removal of a healthy organ or limb—or death. It could also be an adverse drug event in which the harm is transient or minimal. It may even be a near miss in which the wrong drug is dispensed but is detected by the nurse before it is delivered.

The purpose of a study strongly influences the selection of the bad event. PRA studies are performed too often without an explicit purpose other than some vague goal such as “measuring safety” or “identifying improvements”. Such studies are frequently of limited use because they omit actions and behaviors that are critical in some later use of the study. Once a bad event has been chosen, the analyst identifies the potential causes of the event and how they are related. Most PRAs use two complementary graphical tools to do this: event tree analysis and fault tree analysis (FTA).

Event tree analysis

The event tree is a logical structure in the form of a tree branch that maps out the different pathways by which the bad event can come about. All of the paths that cause an adverse outcome must be included and analysts routinely rely on the experience of subject matter experts to know which events to include. The tree structure enables the analyst to order events (usually chronologically), to separate clusters of events from each other, and to show whether or not events are important.

The branching structure shows how an initiating event that starts a sequence at the left side of the tree may lead to the bad event that is shown at the far right side. Events or options that depend on other events are shown to the right of those events on which they depend. Figure 1 shows an everyday example—the problem of being late for class. We will take as the bad event “being late again” (while this may seem a trivial application it is convenient in scale and requires no specific clinical knowledge for a real medical application. The analytical process is identical, however). The particular sequence of events we will consider starts with the subject waking up late and being time pressured to get to class.

The subject has three ways to get there. The normal way is by driving his/her own car via a freeway that is subject to periodic overcrowding and delays while driving. The first alternative would be to use public transport (say, a local subway or commuter train) and the second is to call a colleague to ask for a ride.

Figure 1 shows an event tree including alternatives and the different things that could lead to the student being late again. The alternative outcomes are shown in the right hand column (“Late again”) as either “yes” or “no” for the outcome of the path that terminates to the left of the outcome answer. Trace back from the outcomes towards the

left hand side of the tree along the horizontal paths. There are a series of vertical branches labelled “Y” (for yes) and “N” (for no) that are connected to earlier paths. Each of the vertical branches represents the response (yes or no) to a question that appears at the top of the tree. Tracing back from the first “no” under “Late again” we come to the first branch labelled Y/N: “Freeway clear?” The up branch represents “yes” and indicates that the freeway on this particular morning was clear. The student was not therefore held up by traffic on the freeway and arrived on time. The down branch “no” means that the freeway was not clear and the student was late. This branch is attached to the earlier path and represents the condition that the car did start. The up branch corresponding to the question “Car starts?” indicates “yes”. Because the car did start, there is no need to consider the backup alternatives of the subway or the colleague. To keep things simple we ignore other failure modes such as the car having a flat tyre, being involved in an accident, or any of the other things that seem to happen when time is of the essence.

What are the possible outcomes if the car does not start? Work from left to right, starting on the lower “No” branch associated with the question “Car starts?”. The next question is “Train/subway available?” The “yes” path goes straight to the outcome of not being late again. Notice that the questions in the event tree sound very simple. However, in order to satisfy our analysis there are several subway or train possibilities that need to be considered. Is the train sufficiently frequent and conveniently located to get the student to class on time? Is the day being analysed a holiday with reduced service? Has there been an accident or breakdown on the line in question? When we discuss fault trees in the next section we show that these kinds of questions are addressed to answer “Train/subway available?”

If the answer is “no” then we are left with the colleague option, and whether he is available and willing to give the subject a ride in time. If not, he will be late. If yes, whether the freeway is clear must be considered. Heavy freeway traffic can still cause a late arrival.

Fault tree analysis (FTA)

A fault tree is an extension of the event tree method. Park⁶ considers FTA “a method of system reliability/safety analysis that .. shows a logical description of the cumulative effects of faults within the system”. Like the event tree, events in a fault tree are arranged to show how they are related. Event trees are portrayed in a logic structure that branches from left to right and uses only OR gate logic. In contrast, a fault tree is organized in a “top to bottom” hierarchy and uses both AND and OR gate logic. The fault tree diagram (fig 2) adds logic diagram symbols to the tree structure.

The diagram represents cause and effect relations among events that culminate in a “top event”. Logic symbols at each intersection (or gate) indicate what is required to occur for its condition to be satisfied. The AND gate requires multiple events to occur at the same time—that is, the output condition exists when *all* the inputs exist. Thus, in fig 2 the condition “no gas” exists when both conditions “no gas in tank” and “no gas in spare can” are true. The OR gate is satisfied when *any* one event occurs. Referring to fig 2, the condition “no backup electrics” (electrical power) exists when either “no jumper cables” or “no second battery” exists. The output condition is also satisfied when more than one input exists in an OR gate. Bahr⁷ observes that the more AND gates a tree contains, the more fault tolerant (and safer) a system typically is. A proliferation of OR gates depicts a failure prone situation.

FTA can be used to build a model to predict the likelihood of each branch for which the analyst has no direct

Figure 1 Example event tree structure.

experience. For our commuting student, fig 2 refines events into probabilities and dependent conditions. A subject matter expert can estimate the likelihood of certain combinations in order to figure AND and OR gate probabilities. An AND gate condition provides a simple example. Events that are assigned to each AND gate will affect the likelihood that both will occur and thereby meet the gate's logical requirement. The probability of the output being "true" (P_o) for two inputs, A and B, with probabilities of being true (P_a and P_b , respectively) is given by:

$$(P_o = P_a \times P_b)$$

If we know that the probability on any given day that the gas tank will be empty is 0.01 (that is, 1%), and the probability that the spare can is empty is 0.3 (that is, 30%), then the probability that the car will not start because of no gas is:

$$P_o = 0.01 \times 0.3 = 0.003 \text{ (or 0.3\%)}$$

The mathematical formulation is a little more complex with an OR gate. Using the same terms as above (P_o , P_a and P_b) for the output and the two inputs, the probability of the output being "true" is calculated by:

$$(1 - P_o) = (1 - P_a) \times (1 - P_b),$$

$$\text{or } P_o = (1 - (1 - P_a) \times (1 - P_b))$$

For the backup electrics we assess that the likelihood of not having or being able to locate the jumper cables is 0.1. The likelihood of not having a spare battery or a helpful neighbor not being home is also 0.1. The likelihood of not being able to use the backup system either due to no cables *or* no second battery is:

$$P_o = (1 - (1 - 0.1) \times (1 - 0.1))$$

$$\text{or } (1 - (0.9 \times 0.9)) = 0.19$$

Figure 2 shows how the process is extended to include the other contributing factors to the car not starting when needed. The probability of not starting on any routine day is 0.004, or approximately 1 in 10 years. This probability would be used as an input for the probability of failure for the event "Car starts?" in fig 1. Other fault tree models would be created for the other events in the event tree.

Figure 3 shows how circumstances change when the occasion of misplacing car keys is added. Not finding the keys is an immediate cause of the car not starting. It is added to the top event, where we estimate that on any given morning there is a 1% chance of not finding them. The change increases the overall probability of the car not starting to 0.014, or approximately 3–4 times per year (a 30–40-fold increase over the analysis which neglected the keys not being found).

This brings us to an important point that novice PRA analysts often overlook. Human behavior varies, and that variety is often linked to adverse events. Leaving human performance out of the PRA model will cause the analyst to miss some dominant influences. This will produce results that are very different from the real world. Modeling human performance and reliability is a specialized area within PRA and is beyond the scope of this article. More information on human reliability analysis (HRA) and different methods

Figure 2 Example fault tree.

that have been used to perform it are available in the literature.^{8,9}

Fault trees can be used alone to calculate specific events of concern such as an error in a specific procedure. However, they are not very well suited to represent conditional events. Stephenson¹⁰ describes event tree analysis as a variant of FTA that can be used to "explore both success and failure alternatives at each level". While event trees are meant to show "the path by which we got here", fault trees are not

Figure 3 Example fault tree (extended).

conditioned by what has happened before. It has therefore been said that fault trees have no “memory”.

Even though they can be calculated precisely, PRA results are not exact because, in most cases, the input data have inherent uncertainties and the analysis results are based on explicit assumptions. Assumptions about the number of people working in a team could affect the probabilities of human erroneous actions. Assumptions about the type of equipment in use may affect the rates of errors (in, for example, an infusion pump interface design that influences practitioner performance). Because of this, the prediction that PRA produces is best used to compare proposed solutions on a common basis.

USES OF PRA IN OTHER INDUSTRIES

The early applications of PRA techniques (primarily fault tree models) assessed the reliability of the Minuteman missile system during the design stage, before operational testing could be performed. Designers built fault trees and used failure probabilities for the components that had been derived from testing or from experience in other applications. By using this approach, designers were able to identify the most likely sources of system failures and to make changes in the design stage in order to improve overall system reliability.

Perhaps the most extensive use of PRA has been in the commercial nuclear power industry. The US Atomic Energy Commission, the precursor to the current US Nuclear Regulatory Commission, initiated the reactor safety study in the early 1970s. The reactor safety study was an extensive analysis that was commissioned to estimate the frequencies of accidents that could lead to uncontrolled releases of radioactive materials from reactors.¹¹ Many of the techniques that had been developed for the Minuteman program were also used in the study. The additional need to consider time-sequence dependencies for accident development led to the development of event trees. Data for many of the component failures modelled in the fault trees and event trees were developed from maintenance records.

Work has continued in the nuclear industry to improve the PRA modelling processes, especially in the area of human performance. Human performance analyses in the reactor safety study were principally concerned with fairly simple discrete tasks such as misreading an indication, selecting a wrong switch, or skipping a step in a written procedure. However, the nuclear plant accident in March 1979 at Three Mile Island, Harrisburg, PA showed that its operators misunderstood the conditions in the reactor.¹² This misunderstanding was so fundamental that the operators systematically and purposefully took exactly the wrong actions which led to the accident. Since that event, much work has been performed to develop models of human performance to assess the likelihood of such misunderstandings—see, for example, the US Nuclear Regulatory Commission.¹³

Several different domains such as aviation and space operations already use PRA. The US Federal Aviation Administration allows the use of PRA as part of its demonstration of the acceptability of aircraft designs. The method is used to show that the likelihood of certain types of failures of aircraft systems that would cause a crash is “extremely improbable”.^{14*} More explicit failure frequency requirements are being developed in Europe¹⁵ in response to the standards that have been set by the Joint Aviation Authorities. PRA has become a standard tool to assess the safety of hazardous industries in Europe and elsewhere. For example, PRA analyses are required as part of the formal

*Taken to be less frequent than one failure in a billion flight hours of operation.

safety case for offshore platforms that are regulated by the UK Health & Safety Executive.¹⁶ It is also finding use in new applications. As train controls start to use new computer based systems, PRA is being used to assess railroad reliability and safety—for example, the US Federal Railroad Administration has proposed PRA as one means to assess the safety of new train control systems.^{14–17}

Many safety regulations specify three requirements: (1) that single human or equipment failures should not result in an unacceptable accident; (2) that those who are responsible will determine whether existing or proposed barriers are appropriate given the levels of risk; and (3) that changes to improve safety will be evaluated according to whether they can be implemented economically and efficiently. PRA cannot be used alone to answer all of these conditions. However, as part of a larger safety assessment program, PRA can be used to search for potential combinations of events that can lead to failure.

Experience in other industries suggests that PRA can be most effective when certain conditions exist. Firstly, it must be possible to describe the functional interconnections between different entities such as people and equipment, and to describe how failures in one entity can affect safety. Secondly, there must be some kind of consistency among the processes that are under study. Thirdly, it should be possible to describe failures within the context of a task or subtask—for example, it is not practical to model failures in artistic endeavors in which there are almost infinite routes to success. While medicine may be described as “the art of healing”, the use of standardized protocols and the move towards evidence-based practices removes much of the idiosyncrasy from many procedures. Unlike fine art, medical practice requires some kind of specifiable outcome by which success or failure can be judged. The outcome might be the occurrence or avoidance of injuries, fatalities, or economic losses of some specified magnitude.

STRENGTHS AND LIMITATIONS OF PRA

As with any method, PRA has its own strengths as well as limitations that the analyst should understand.

Strengths

PRA offers the healthcare analyst a number of benefits including element integration, prospective analysis, change evaluation, and direct graphical representation.

Element integration

PRA can account for all major and minor elements in the causes of events. These include combinations of human actions, hardware/software faults, procedural mistakes, and circumstantial factors. Few other approaches to safety assessment allow such “real life” combinations.

Prospective analysis

PRA can be used to anticipate and remedy potential adverse events without having to wait for them to happen to practitioners or patients.

Change evaluation

PRA allows “what if” studies to examine the effectiveness of changes to the system and to look for the most efficient and effective solutions. For example, the method can be used to consider the potential effects of changing the type or the design of equipment that is used in a healthcare setting.

Direct graphical representation

Event and fault trees explicitly describe event relationships through easy to understand diagrams. These graphical tools make it possible for diverse groups to interact and to develop,

or at least discuss, shared perspectives on the causes of bad events.

Limitations

The analyst must also be aware of limitations in PRA that include potential for naïve analyses, reduction, unavailable probability data, reliance on expert estimation, presumption of binary states, and tunnel vision.

Naïve analyses

PRA tools are so simple and easy to apply that an inexperienced analyst can come to inaccurate conclusions. Checking whether a valve is open or closed is a common industrial task that illustrates this point. Some believe that assigning more people to check its status will drive the likelihood of the valve being out of the correct state to almost zero. Experience tells us, though, that the more people who are assigned to check something, the more likely it is that people will shirk the task. The mind set is “I know Charlie checked it and he’s good, so I don’t need to waste my time checking it”. The notion is that everyone will check it. The result is that no one checks it.

When events occur that are not reflected in the PRA models, the results either lead to ineffective solutions or the method itself can be discredited. Relying on experienced analysts and peer group reviews to make sure that events are included can minimize this possibility.

Reduction

Simplifying human experience into tree structure statements excludes the full richness of human performance and associated problems, even though Reason, Rasmussen and others^{9 18 19} have developed methods to evaluate the causes and effects of different kinds of human errors. For example, heroic or inspired actions of people are often ignored. Only their departures from an ideal performance are modeled as failures.

Unavailable probability data

Systems may be in the early stages of development or existing systems may need changes that have not been made. As a result, data may not be available to estimate event likelihood. However, there are ways to extrapolate existing data to new situations. In some situations filtering and scaling can be used to adjust existing data to estimate how proposed changes may affect performance.²⁰ In other situations expert judgment can be used to estimate probabilities provided that care is taken to avoid known biases and limitations.²¹

Reliance on expert estimation

The more complex the circumstance, the fewer data are available to support accurate probability estimates. This requires the analyst to rely on the judgment of those with experience to estimate the likelihoods of occurrence. This can be done by asking the expert to perform the tasks that are being modeled. It can also be done by asking the expert to observe the tasks, which can be less subject to bias. Hollnagel²² cautions that knowledge of human behavior must go beyond what is observable in order to grasp what causes erroneous acts.

Vulnerable to bias

The method relies on the analyst’s integrity and awareness. It is possible to “game” the method by assigning certain events in ways that would indicate a more favorable outcome. Modeling poorly defined (“gray”) areas in studies can be biased because the analyst unwittingly prefers a particular solution. For example, it may be assumed that certain faults can be easily detected and corrected. Those faults could then

be excluded from the analysis without any actual test to support such a choice.

Presumption of binary states

All events in fault trees and event trees represent some kind of binary state. Either a failure happens or it does not. Not all matters can be reduced to such simple decisions, particularly human performance. The complexity of human performance begs further exploration beyond such black and white statements. For instance, when is a human erroneous act a failure? Is it a failure if the practitioner detects it? If another member of the care team detects it? Do subconscious slips count if no harm is done? What is the standard by which performance is judged? Is it perfection? Is it minimally adequate care?

Tunnel vision

The analyst who focuses on hardware failures may inadvertently omit human performance from fault trees and event trees. This limited vision is one of the foremost reasons why PRA predictions fail to match actual experience. In another instance of limited vision, changes can be evaluated too narrowly. For example, Ford’s decision not to reinforce Pinto fuel tanks at a cost of \$6.65/car does not appear to have taken larger issues such as the perceived cost of harm into account. The decision eventually resulted in significant litigation costs and public disapproval.²³

ROLE OF PRA IN HEALTH CARE

In high hazard, low risk systems such as chemical plants, the potential for harm is great but the incidence of accidents is low. PRA has traditionally been used to assess the designs of these systems by predicting the types and likelihood of major accidents that might occur. More recently it has been used in commercial nuclear power plants, aviation, and chemical manufacturing plants²⁴ to explore how things went wrong, what alternative outcomes were possible, and whether new or changed defences and barriers may be more effective. These applications have been driven by a need for models that can predict the types and frequencies of major accidents in order to assess whether the design ensures that they are adequate.

Health care, particularly acute care, combines high hazard with high risk. A small number of studies using risk assessment tools have already been performed for events. One study explored their use in anesthesia²⁵ while another involved radiation brachytherapy.²⁶ Even so, the broader use of PRA requires the analyst to understand how health care differs from other high hazard applications. Rather than the continuous operation of single design systems (as in a chemical plant), health care routinely involves aggregations of equipment that are assembled and interchanged quickly. Furthermore, healthcare processes are far more diverse than traditional applications in which PRA has become an accepted safety management tool.

- Personnel: a wide range of doctors, nurses, pharmacists, technicians, clerical staff, contract employees, manufacturers, and pharmaceutical equipment and consumables vendors perform in important roles that affect safety.
- Procedures: treatments are specific and particular to a patient or patient’s condition and they compel the use of specific qualifications, procedures, skills, tools, and equipment. There are no opportunities to “set it and forget it”.
- Facilities: each healthcare location—from major trauma centres to ambulatory clinics, to assisted living facilities and hospices—has its own unique set of safety concerns.
- Criticality: patient health trends, diagnoses, anatomies, compliance, and response to treatment all affect outcomes.

Key messages

- Human factors methods help analysts to understand health care as a system.
- Probabilistic risk assessment helps the analyst to understand large systems as a whole, rather than isolated parts.
- It can be used to identify the causes, potential severity, and frequency of adverse events.
- Probabilistic risk assessment leads the analyst to portions of a system that may have safety related issues and indicates where to allocate resources for improvement.
- The analyst can use the probabilistic risk assessment model to explore the potential risk of changes before they are implemented.

Every patient is vulnerable, having arrived in the system as a result of a need for treatment.

“Risk” in PRA focuses on influences that lead to bad events, making it more specialized than “risk management” in health care. Risk management combines quality management and the management of patient safety in the interest of limiting malpractice liability.²⁷ Even with this difference, there is potential for PRA to benefit health care. Both PRA and health care risk management seek to understand and control the potential exposure of patients and staff to hazards. PRA can be used to improve patient safety efforts in four ways: formality, logic, structure, and prioritization.

- **Formality:** PRA provides a formal systematic way to identify and represent the factors that contribute to adverse events and near misses.
- **Logic:** PRA provides a logical basis to explore interventions and barriers that could reduce iatrogenic harm. PRA qualitative tools (the fault trees and event trees) can be used to identify all potential changes, while PRA quantitative tools can be used to measure change effectiveness and then to rank potential changes accordingly.
- **Structure:** PRA provides a framework for data reporting requirements. The structure provided by PRA will define the scope, conditions, and other aspects of adverse event data gathering, making it purposeful.
- **Prioritization:** PRA provides a way to focus on the most vulnerable aspects of the system and to protect them from failure.

PRA represents a logical extension to tools that are already in use within the patient safety community, specifically FMEA, RCA, and data gathering efforts. For example, the information created by FMEA and RCA can be used to build fault trees and event trees. However, several changes are needed in the use of PRA before it can be applied to healthcare analysis. The reliance of PRA on expert assessment calls for deep domain knowledge. Such knowledge necessarily comes from research into the deep structure of work that is performed at the “sharp end” of acute care. It includes three essential elements:

- Data on sharp end practice based on first hand observation.
- Analyses and findings that characterize sharp end practice drawn from basic research into healthcare operations.

- Insightful representations of sharp end practice and the influence of management policy and procedures at the system level.

Without such research, PRA models will lack the substance that they require to be valid and run the risk of producing naïve analyses that appear useful but lack rigor or validity.

CONCLUSIONS

PRA is a proven method for identifying and evaluating risk in high hazard applications that has the potential to improve patient safety efforts in health care. Its formal and structured procedures make it a promising way to identify and assess potential adverse events. PRA can build on FMEA, RCA, and data collection efforts that are already in place in healthcare settings. By integrating these methods, PRA can make them even more useful.

In order to use PRA effectively, further work must be done in healthcare research. Analysts must initially develop a deeper understanding of health care as a system. That understanding can then be used to correctly account for events to populate the event and fault trees, assign well considered links among events, and estimate the likelihood of occurrence with a reasonable degree of confidence.

Authors' affiliations

J Wreathall, John Wreathall & Co Inc, Dublin, OH 43016-9578, USA

C Nemeth, Cognitive Technologies Laboratory, University of Chicago, Chicago, IL 60637, USA

REFERENCES

- 1 **Nemeth C.** *Human factors methods for design*. London: Taylor & Francis/CRC Press, 2004.
- 2 **Feldman SE**, Douglas DW. Medical accidents in hospital care: applications of failure analysis to hospital quality appraisal. *Jt Comm J Qual Improv* 1997;23:567–80.
- 3 **DeRosier J**, Stalhandske E, Bagian JP, et al. Using health care failure modes and effects analysis: the VA National Center for Patient Safety's Prospective Risk Analysis System. *Jt Comm J Qual Improve* 2002;28:248–67.
- 4 **Cohen MR**, Senders J, Davis NM. Failure modes and effects analysis: a novel approach to avoiding dangerous medication errors and accidents. *Hosp Pharmacy* 1994;29:319–24.
- 5 **Nemeth C.** *The master schedule: how cognitive artifacts affect distributed cognition in acute care*. *Diss Abs Int* 2003;64:3990.
- 6 **Park K.** Human error. In: Salvendy G, ed. *Handbook of human factors and ergonomics*, 2nd ed. New York: John Wiley & Sons, 1997:150–71.
- 7 **Bahr NJ.** *System safety engineering and risk assessment*. London: Taylor and Francis, 1997.
- 8 **Reason J.** *Human error*. New York: Cambridge University Press, 1990.
- 9 **Gertman DI**, Blackman HS. *Human reliability and safety analysis data handbook*. New York: John Wiley & Sons, 1994.
- 10 **Stephenson J.** *System safety* 2000. New York: Van Nostrand Reinhold, 1991.
- 11 **Nuclear Regulatory Commission (NRC).** *Reactor Safety Study – an assessment of accident risks in US commercial nuclear power plants*. Report No WASH-1400 (NUREG 75/014). Rockville, MD: US Nuclear Regulatory Commission, 1975.
- 12 **Kemeny J.** *The need for change. Report of the President's Commission on the accident at Three Mile Island*. New York: Pergamon Press, 1979.
- 13 **Nuclear Regulatory Commission (NRC).** *Technical basis and implementation guidelines for a technique for human event analysis (ATHEANA)*. Report No NUREG-1624, Rev. 1. Rockville, MD: US Nuclear Regulatory Commission, 2000.
- 14 **Federal Aviation Administration.** *Equipment, systems and installations* 14 CFR 25.1309. Washington, DC: Department of Transportation, Federal Aviation Administration, 2003.
- 15 **Joint Aviation Authorities.** *Equipment, systems and installations. Joint Aviation Regulation 25.1305*. Hoofddorp, The Netherlands: Joint Aviation Authorities, 2003.
- 16 **Health & Safety Executive.** *Assessment principles for offshore safety cases*. London: Health & Safety Executive, 1998.
- 17 **Department of Transportation.** *Standards for development and use of processor-based signal and train control systems*. Proposed rule (Federal Railroad Administration). 49 CFR Part 209. Federal Register 2001;66:155.
- 18 **Rasmussen J.** Role of error in organizing behavior. *Ergonomics* 1990;33:1185–90.
- 19 **Reason J.** *Managing the risks of organizational accidents*. Brookfield, VT: Ashgate Publishing Company, 1998.
- 20 **Wreathall J**, Bley DC, Roth E, et al. Using an integrated process of data and modeling in HRA. *Reliability Engineering System Safety* 2004 (in press).

- 21 **Budnitz RJ**, Apostolakis G, Boore DM, *et al.* *Recommendations for probabilistic seismic hazard analysis: guidance on uncertainty and use of experts*. Report No NUREG/CR-6372. Rockville, MD: US Nuclear Regulatory Commission, 1997.
- 22 **Hollnagel E**. *Cognitive reliability and error analysis method (CREAM)*. New York: Elsevier Science, 1998.
- 23 **Associated Press**. Ford rejected \$6.65 safety change on Pinto, witness says. *Boston Globe* 8 February 1980.
- 24 **Battelle**. *Guidelines for hazard evaluation procedures*. New York: American Institute of Chemical Engineers, 1985.
- 25 **Paté-Cornell ME**, Lakats LM, Murphy DM, *et al.* Anaesthesia patient risk: a quantitative approach to organizational factors and risk assessment. *Risk Anal* 1997;17:511–23.
- 26 **Ostrom LT**, Leahy TJ, Novack SD. *Summary of 1991–1992 misadministration event investigations*. Report No NUREG/CR-6088. Idaho Falls, ID: Idaho National Engineering Laboratory, 1994.
- 27 **Dawson MC**, Munro AP, Appleby KJ, *et al.* Risk management and medical errors. In: Rosenthal MM, Sutcliffe KM, eds. *Medical errors: What do we know? What do we do?* San Francisco: Jossey-Bass, 2002.