CHAPTER SEVEN

MISTAKING ERROR

David D. Woods, PhD and Richard I. Cook, MD

Introduction

Throughout the brief history of the patient safety movement (Hatlie, 1996), stakeholder groups have asked the authors for definitions and taxonomies of human error grounded in science. The questions are of the same form: Each group feels that their progress on safety depends on having a firm definition of human error. Each group seems to believe that this definition will enable creation of a scorecard that will allow them to gauge where their organization stands in terms of being "safe."

Each group's search for the definition, first in the medical literature and then in the general scientific literature, becomes mired in complexity and terms of reference. As definitions appear in the medical literature about patient safety, they seem too specific to particular areas of health care or too vague if broad enough to cover health care in general. The definitions offered are often the product of committee consensus processes and thus too ad hoc to have scientific standing. The definitions offered involve arbitrary and subjective methods of assigning events to categories (e.g., "definitions" that propose to include as errors those events that "could have led to harm"). The resulting counts and extrapolations seem open to endless reassessment and debate (Leape et al., 1991; Kohn et al., 1999; Brennan, 2000; McDonald et al., 2000; Leape, 2000). The definition of error becomes more elusive.

Some within the stakeholder groups turn to social and behavioral science researchers who have grappled with the "human error problem" that has dogged progress on safety in other high-risk industries. Nuclear power, aviation, manufacturing, and the military have invested heavily in basic and applied research on human error over the past 20 years. Although some of this research–and some outspoken researchers–rely on human error being a discrete, well circumscribed, static entity, progress on safety in these industries has come, in large part, from abandoning efforts to attack error.

Driven partly by spectacular failures–the reactor failure at Three Mile Island in 1979 is the archetype–researchers have developed new means for looking into how systems fail and how people in their various roles contributed to *both* success and failure. The research, collectively known as the "New Look" (Rasmussen, 1986; Woods et al., 1994; Reason, 1997), drew on many different disciplines including cognitive science, organizational theory, and cognitive engineering, but it has consistently focused on empirical studies of people at work. The results of these efforts challenge the conventional "folk" assumptions about the relationship between "error" and failure.

To researchers, beginning with the question "What is error?" misleads stakeholders into a thicket of difficulties where answers seem always just around the corner but never actually come into view. The efforts to answer this seemingly simple question–efforts that inevitably become entangled with social factors and lose sight of the research base–actually block progress on safety (Cook, Woods, Miller, 1998). The New Look offers an alternative to "error" as the target of efforts to improve safety (Woods and Cook, 2002).

The need to redirect effort away from error seems counterintuitive when health care is so obviously confronted by a "human error problem." But this is precisely the situation that confronted researchers at the beginning of the development of the New Look (Hollnagel, 1983). At that time, the folk model of accident causation was firmly in place among researchers and error seemed a plausible target for work on safety. It was only after a long period of empirical research on human performance and accidents that it became apparent that answering the question of what is error was neither the first step nor a useful step, but only a dead end.

How are we to respond to those who seek definitions and taxonomies in their efforts to improve safety? Faced with an equally difficult problem in 1897, Frank Church replied to a letter from an 8-year-old girl named Virginia who asked if there really was a Santa Claus. As a way of summarizing what the New Look has revealed about error and safety, we propose our own version of both Virginia's letter and Church's response.

Dear David and Richard,

Our professional group just sponsored a consensus meeting on patient safety and medical error. We were all shocked by the Institute of Medicine reports. As highly motivated professionals committed to serve our patients, we decided to take proactive steps to improve patient safety in our area of medicine.

In our consensus meeting we developed an agenda for progress and strategic action. At the top of our list, we put the need for agreement on definitions of medical errors. Until we reach agreement on this we will be unable to create the error-tracking programs we need to measure the size of the problem and to evaluate the effectiveness of new interventions. However, despite much struggle, we have been unable to achieve consensus on a set of acceptable and workable definitions.

We are soliciting input from different experts on human error outside health care to help us develop these definitions. *What is human error?*

Thank you for your help.

Virginia

Here is our reply:

Dear Virginia,

You are not the first to ask us about "human error." When researchers began studying the role of human performance in system failure, they asked the same questions, thinking that they would find the answers if only they worked hard enough and long enough.[1]

There were good reasons for them to look for the answer. Accidents, sometimes very big ones, seemed to involve human error, and the costs of these accidents were such that getting a clear idea of what human error is seemed essential for progress on safety. The researchers did not have much success however. Instead, they found a different sort of way of looking at accidents and failures, a way we now call the New Look. This is a way of understanding how systems fail and how people in their various roles contribute to both success and failure.[2]

To answer the question as you have posed it would require us to accept a set of conventional assumptions about error. We now

[1]Cook (1999) traces how individuals and organizations recapitulate the learning steps that gave rise to the current understanding of the relationship between safety, systems, and human performance.

[2]The core of the New Look results can be seen in the work of Jens Rasmussen (1986; 1990a; 1990b; 1994; 1999); Erik Hollnagel (1983; 1993; 1998; 1999; Hollnagel and Amalberti, 2001); and Woods et al. (1994). Parallel work on high reliability organizations can be found in Rochlin (1999) and Weick et al. (1999).

understand that these assumptions are incorrect. The focus on defining and counting error has distracted other industries from productive work on safety. The research results can help you find more useful questions to ask.

First off, we want you to understand that the term "error" is used inconsistently in our everyday conversations about safety and accidents.[3] There are at least three ways that the term is used:

- Sense #1. Error as the *cause* of failure: "This event was due to human error." The assumption is that error is some basic category or type of human behavior that precedes and generates a failure. It leads to variations on the myth that safety is protecting the system and stakeholders from erratic, unreliable people.
- Sense #2. Error as the *failure itself,* i.e., the consequences that flow from an event: "The transplant mixup was an error." In this sense the term "error" simply asserts that the outcome was bad and produced negative consequences (e.g., injury to a patient).
- Sense #3. Error as a *process,* or more precisely, *departure from the "good" process.* Here, the sense of error is of deviation from a standard, that is a model of what is good practice, but the difficulty is there are different models of what is the process that should be followed: e.g., what standard is applicable, how standards should be described, and what it means when deviations from the standards do not result in bad outcomes. Depending on the model adopted, very different views of error result.

While you might think that it would always be clear from the context which of these senses people mean when they talk about error, in practice the senses are often confused with each other. Even worse, people sometimes slip from one sense to another without being aware that they are doing so.

Of course, what people are interested in is "error" in the second sense of bad outcomes and how to prevent them. You can see this yourself by this simple thought experiment. Imagine that we have managed to eliminate all accidents from health care but that many errors remain (third sense). Would anyone be interested in error? Probably not. Now suppose instead that we somehow eliminated all the *errors* from health care but that accidents continued to occur. Would anyone be interested in error? Again, probably not.

The motivation to explore error comes from accidents. The research that forms the New Look has two main sources: (1) the base of behavioral science about how individuals and groups cope with complexity and conflict in real world settings and (2) a growing set

[3]The work of Jens Rasmussen and Erik Hollnagel have led the way. See in particular: Rasmussen (1999), Hollnagel (1983; 1993; 1998), and Hollnagel and Amalberti (2001). The description of the first 3 senses is drawn from Hollnagel's work (1993).

of empirical studies on accidents, near-accidents, and real-world work, sometimes called "technical work" studies. The research view of "error" derived from the New Look has produced five conclusions pertaining to "error":

Conclusion #1: *Defining error-as-cause (Sense #1) blocks learning by hiding the lawful factors that affect human and system performance.*

The critical observation that gave rise to the New Look was that errors were *heterogeneous* and not directly comparable events that could be counted and tabulated. The standard way we say this today is that the label error should be the starting point of study and investigation, not the ending point.

It is tempting to stop the analysis of an adverse event when we encounter a person in the chain of events. Continuing the analysis *through* individuals requires workable model cognition of individuals and of coordinated activity between individuals. It turns out to be quite hard to decide where to halt the causal analysis of a surprising event. Although there are theoretical issues involved in this *stopping rule problem*[4] the decision about when to stop most often reflects our roles as stakeholders and as participants in the system. We stop when we think we have a good enough understanding, and this understanding is, not surprisingly, when we have identified human error as the source of the failure.

The idea of error-as-cause also fails because it trivializes expert human performance. Error-as-cause leaves us with human performance divided in two: acts that are errors and acts that are non-errors. But this distinction evaporates in the face of any serious look at human performance.[5] What we find is that the sources of successful operation of systems under one set of conditions can be what we label errors after failure occurs. Jens Rasmussen likes to quote Mach (1905) on this point:

"Knowledge and error flow from the same mental sources, only success can tell one from the other."

Instead of finding error and non-error, when we look deeply into human work we find that the behaviors there closely match the incentives, opportunities, and demands that are present in the workplace. Rather than being a distinct class of behavior, we find the natural laws that influence human cognition and performance are *always at work,* sometimes producing good outcomes and sometimes producing bad ones. Trying to separate error from non-error makes it to harder to see these factors.

[4]Rasmussen (1990b) is the standard statement of the difficulty. For an approach to deal with the problem see Rasmussen (1994). For examples of the difficulty in health care see Cook et al., *A Tale of Two Stories* (1999).

[5]The best introduction to the complexity of expert performance is found in Gary Klein's book *Sources of Power: How People Make Decisions* (MIT Press, Cambridge MA, 1998).

Conclusion #2: *Defining error-as-consequences (Sense #2) is redundant and confusing.*

Much of the time in health care, the word "error" is used to refer to harm—generally *preventable* harm—to patients. This sort of definition is almost a tautology: it simply involves renaming preventable harm as error. But there are a host of assumptions that are packed into "preventable" and these are almost never made explicit. We are not interested in harm itself but, rather, *how harm comes to be.* The idea that something is preventable incorporates a complete (albeit fuzzy) model of how accidents happen, what factors contribute to them, and what sorts of countermeasures would be productive. But closer examination of "preventable" events shows that their preventability is largely a matter of wishing that things were other than they were.

To use "error" as a synonym for harm gives the appearance of progress where there is none. It would be better if we simply were clear in our use of language and referred to these cases in terms of the kind of harm or patient injuries. Confounding the label error with harm simply adds a huge amount of noise to the communication and learning process.

Conclusion #3: *Defining error-as-deviation from a model of "good" process (Sense #3) collides with the problem of multiple standards.*

The critical aspect of error-as-process-deviation is deciding how to determine what constitutes a deviation. Some have proposed normative models, e.g., Bayes Theorem, but these are rarely applicable to complex settings like health care domains, and efforts to use this approach to assess human performance are misleading.[6]

Some have argued that strict compliance with standard operating practices and procedures can be used to define deviation. In other fields, however, it was quickly discovered that standard operating practices capture only a few elements of work and often prescribe practices that cannot actually be sustained in work worlds. In transportation systems, for example, where striking may be illegal, labor strife has sometimes led workers to adopt a "work-to-rule" strategy. By working exactly to rule, workers can readily make complex systems stop working. Attempts to make complete, exhaustive policies that apply to all cases creates or exacerbates double binds or to make it easy to attribute adverse events to "human error" and stop.[7]

[6]Humans are not Bayesian machines and their success in the world is not the result of Bayesian statistics so comparing their performance to Bayesian models necessarily misrepresents the nature of work, cf. Gigerenzer F., Simple Heuristics: Things that make us smart (Oxford University Press, 1999) and Klein's Sources of Power.

[7]Lucy Suchman's (1987) book illustrates how expert practice is more than just following standard policies and procedures; *Behind Human Error* captures the double binds that can arise; and Woods and Shattuck (2000) summarize the basic tradeoff and forms of failure from this point of view.

Choosing among the many candidates for a standard changes what is seen as an error in fundamental ways. Using finer or coarser grain standards can give you a very wide range of error rates. In other words, by varying the standard seen as relevant, one can estimate hugely divergent "error" rates. Some of the "standards" used in specific applications have been changed because too many errors were occurring or to prove that a new program was working. To describe something as a "standard" when it is capable of being changed in this way suggests that there is little that is standard about "standards."

This slipperiness in what counts as a deviation can lead to a complete inversion of standardizing on good process: rather than describing what it is that people need to do to accomplish work successfully, we find ourselves relying on bad outcomes to specify what it is that we want workers not to do. Although often couched in positive language, policies and procedures are often written and revised in just this way after accidents. Unfortunately, hindsight bias plays a major role in such activities.

Working toward meaningful standards as a means for assessing performance and defining error as deviations might be a long-term goal but it is fraught with hazard. To make standards work requires not only clear statements about how to accomplish work but clear guidance about how conflicts are to be handled. Specifying standards for performance for only *part of the work* to be done creates double binds that undermine expert performance creating conditions for failure. To use standards as a basis for evaluating performance deviations requires the continuous evaluation of performance against the standard rather than (as is often the case) simply after bad outcomes become apparent. One practical test of this is whether or not deviations from standards are actually detected and treated in the same way independent of the actual outcome.

To limit the damage from the multiple standards problem, all must carry forward in any tabulation the standard used to define deviations. *This is absolutely essential!* Saying some behavior was an error-as-process-deviation has no meaning without also specifying the standard used to define the deviation.

There are three things to remember about the multiple standards problem. First, the standard chosen is a kind of *model* of what it means to practice before outcome is known. A scientific analysis of human performance makes those models explicit and debatable. Without that background, any count is arbitrary.

Second, a judgment of error is not a piece of data which then can be tabulated with other like data; instead it is the end result of an *analysis*. Its interpretation rests on others being able to decompose and critique that analysis. The base data is the *story* of the particular episode—how multiple factors came together to produce that outcome. Effective systems of inquiry about safety begin with and continually refer back to these base stories of failure and of success in the learning process.

Third, being explicit about the standard used is also essential to be able to critique, contrast, and combine results across events, studies, or settings. When these standards are dropped or hidden in the belief that error is an objective thing in the world, communication and learning collapse.

In the final analysis, the science has shown that "error" is an example of an *essentially contestable* concept. In fact, any benefit to the search for error only comes from the chronic struggle to define how different standards capture and fail to capture our current sense of what is expertise and our current model of the factors that make the difference between success and failure.

Gradually, the research activities that comprise the New Look have led to recognition of a fourth sense of "error."[8]

Conclusion #4: *Labeling an act as "error" marks the end of the social and psychological process of causal attribution.*

If you really want a definition of error, Virginia, we suggest you use this one.

Taken together, the research on how people actually apply the term "error" shows that "error" is a piece of data about reactions to failure that serves as a placeholder for a set of *socially derived beliefs* about how things happen. As stakeholders, our judgments after the fact about causality are used to explain surprising events. *Thus, in practice, the study of error is the nothing more or less than study of the psychology and sociology of causal attribution.* There are many regularities and biases—e.g., the hindsight bias—that determine how people judge causality. The heterogeneity and complexity of real-world work make these regularities and biases especially important: because the field of possible contributors includes so many items, biases may play an especially important role in determining which factors are deemed relevant.

These results are deeply unsettling for stakeholders because they tell us that the use of the term "error" is less revealing about the performance of workers than it is about ourselves as evaluators. As researchers, advocates, managers, and regulators, we are at least as vulnerable to failure, susceptible to biases and oversimplifications, and prone to err as *those other people.* Fallibility has no bounds in a universe of multiple pressures, uncertainty, and finite resources.

Error is not a fixed category of scientific analysis. It is not an objective, stable state of the world. Instead, it arises from the interaction between the world and the people who create, run, and benefit (or suffer) from human systems for human purposes—a relationship between hazards in the world and our knowledge, our perceptions, and even our dread of the *potential* paths toward and forms of failure.

---

[8]This sense of error was first articulated in Woods et al. (1994) and Cook and Woods (1996).

What is the consequence of error being the result of processes of attribution?

Although you did not say so explicitly, the question of what is error was predicated on the notion that we can and should treat error as an objective property of the world and that we can search for errors, tabulate them, count them. This searching and counting is futile.

The relationship between error and safety is mirage-like. We find ourselves in a desert, seeing safety glimmering somewhere in the far distance. To begin the journey, we feel we must gauge the distance to our goal in units of "error." This presumption about the location of safety is illusory. Efforts to measure the distance to it are little more than measuring our distance from a mirage. The belief that estimates of this number are a necessary or even useful method of beginning an effort to improve safety are predicated on the apparent location of the mirage.

The psychology of causal attribution, however, tells us that it is our beliefs and misconceptions about failure and error that have combined to make the mirage appear where it does. The New Look research tells us that progress toward safety has more to do with the metaphorical sand underneath our feet than it does with the tantalizing image off in the distance. When we look closely, we see how health care workers are struggling to anticipate forms of/paths toward failure, actively adapting to create and sustain failure-sensitive strategies, and working to maintain margins in the face of pressures to do more and do it quickly. Looking closely under our feet we see:

1. How workers and organizations are continually revising their approach to work in an effort to remain sensitive to the possibility for failure;
2. How we and the workers are necessarily only partially aware of the current potential for failure;
3. How change is creating new paths to failure and new demands on workers and how revising their understanding of these paths is an important aspect of work on safety;
4. How the strategies for coping with these potential paths can be either strong and resilient or weak and mistaken;
5. How dependent the culture of safety is on remaining dynamically engaged in new assessments and avoiding stale, narrow, or static representations of risk and hazard;
6. How overconfident nearly everyone is that they have already anticipated the types and mechanisms of failure, and how overconfident nearly everyone is that the strategies they have devised are effective and will remain so;
7. How missing the side effects of change is the most common form of failure for organizations and individuals; and
8. How continual effort after success in a world of changing pressures and hazards is fundamental to create safety.

In the final analysis, safety is not a commodity to be tabulated, it is a chronic value "under our feet" that infuses all aspects of practice. *People create safety under resource and performance pressure* at all levels of the sociotechnical system. They continually learn and adapt their activities in response to information about failure. Progress on safety ultimately comes from helping workers and managers create safety.

The folk models about human error are pervasive and seem self-evident. Challenging those folk models calls into question not only closely held beliefs, but also policy and investment decisions. The dissonance between belief and results is uncomfortable at best. Yet, Virginia, the need to make real progress on safety leaves us no choice but to point out how easy it is for all of us to fall back into fallacies and myths about human error. We must make the contrast very stark:

> The misconceptions and controversies about error and safety are rooted in the collision of two mutually exclusive paradigms or world views. One view is that erratic people degrade an otherwise safe system. Thus, work on safety is protecting the system (us as managers, regulators, and consumers) from unreliable people. This is a Ptolemaeic world view (the sun goes around the earth). To defend this world view in the face of the data on human performance and how complex systems fail takes ever greater effort (more and more epicycles by analogy).
>
> The other paradigm or world view is that *people create safety under resource and performance pressure* at all levels of the sociotechnical system by learning and adapting to information about how we all can contribute to failure (this is the basic lesson from New Look research about human performance, success and failure). This is a Copernican world view (the earth goes around the sun). Progress comes from helping people create safety. This is what the science says, despite how odd it sounds: *help people cope with complexity under pressure to achieve success.*
>
> We can blame and punish under whatever labels are in fashion but that will not change the natural laws that govern human performance, nor will it make the sun go round the earth. The paradigm shift demanded if real progress is to be made on safety it is, not surprisingly, extraordinarily difficult. We have windows of opportunity for improving safety, but only if all of us are up to the sacrifices involved in building, extending, and deepening the ways we can help people create safety.
>
> So are people sinners or are they saints? An old debate, but neither view leads anywhere near to improving safety. Making safety begins with recognizing the paradox that, simultaneously, we are both the source of success and of failure. How could it be otherwise?—as we create, operate, and modify human systems for human purposes.[9]

[9]From Woods (2000b).

Virginia, you wanted a simple, pragmatic answer and did not expect to walk into this onslaught of complexities. But accidents arise from the complexities of the domain, not from its apparent simplicity. The attraction of error as a target for work on safety is illusory. No progress or success is possible if we remain trapped in a Ptolemaeic search for erratic other people or if we try to straddle the two paradigms. Either you are working in the Ptolemaeic paradigm, or you throw it off *completely* and move onto the fascinating and productive challenges of deepening the Copernican paradigm of observing, modeling, and enhancing how we cope with complexity and create safety under pressure.

Virginia, adopt the new paradigm and begin by looking for ways to understand the changing vulnerabilities and pathways that expose patients to risks of injury as a result of care. Investigate how people cope with complexity—usually successfully. Search out the sources of resilience that allow them to produce success when failure threatens.[10] In combination, these efforts will allow you to *create foresight,* to recognize, anticipate, and defend against paths to failure that arise as health care organizations and technology change, and to do so *before any patient is injured.*

David Woods and Richard Cook

Summary of the New Look Research Findings

Doing things safely, in the course of meeting other goals, is and has always been part of operational practice. As people in their different roles are aware of potential paths to failure, they develop failure-sensitive strategies to forestall these possibilities. Failures occurred against this background when multiple contributors—each necessary but only jointly sufficient—combine. Work processes do not choose failure but *drift toward it* as production pressures and change erode the defenses that normally keep failure at a distance. This drift is the result of systematic, predictable organizational factors at work, not simply erratic individuals. To understand how failure sometimes happens one must first understand how success is obtained—how people learn and adapt to create safety in a world fraught with hazards, tradeoffs, and multiple goals (Cook et al., 2000).

[10]Many resources are beginning to emerge on the new paradigm and how it can lead to a new set of methods and means to improve safety. A sampler includes Hollnagel (1999) and Svenson (2001) exploring barrier analysis; Carthy et al. (2000) and Reason (2001) looking at the resilience in surgical services; Woods and Shattuck (2000) in analyzing an accident; Weick et al. (1999) in characterizing how high reliability organizations show high resilience.

It is clear that high levels of performance are achievable. For example, researchers have studied organizations that have been remarkably successful in managing potentially hazardous technical operations, and the empirical results match the New Look (Rochlin, 1999). Achieving such high levels of performance does not flow from rooting out error, but rather through anticipating and planning for unexpected events and future surprises. Past success is not a reason for confidence, instead, continued investment in anticipating the changing potential for failure is energized by the deeply held understanding that our knowledge base is fragile in the face of the hazards inherent in work and the changes omnipresent in the environment.

The theme that leaps out from the New Look results is that failure represents *breakdowns in adaptations* directed at coping with complexity. Success relates to organizations, groups, and individuals who produce resilient systems that recognize and adapt to change and surprise. The measure of success for groups and organizations is the ability to "create foresight–anticipate the changing shape of iatrogenic risk, *before* patients are injured" (Woods, 2000).

References

Amalberti, R. (2001). The paradoxes of almost totally safe transportation systems. *Safety Science*, 37, 109–126.

Brennan, T. A. The Institute of Medicine Report on Medical Error–Could It Do Harm? *N Engl J Med* 2000;342:1123–1125.

Carthey, J., de Leval, M. R., and Reason, J. (2000). Understanding Excellence in Complex, Dynamic Medical Systems. In Proceedings of the 44th Annual Meeting of the Human Factors and Ergonomics Society/IEA2000, July, 2000.

Cook, R. I. (1999). Two Years Before the Mast: Learning How to Learn About Patient Safety. In Hendee, W. (ed.), Proceedings of *Enhancing Patient Safety and Reducing Errors in Health Care.* National Patient Safety Foundation, Chicago IL (held at Annenberg Center for Health Sciences, Rancho Mirage, CA, Nov. 8–10, 1998).

Cook, R. I., Woods, D. D., and Miller, C. (1998). *A Tale of Two Stories: Contrasting Views on Patient Safety.* National Patient Safety Foundation, Chicago IL, April 1998 (www.npsf.org/exec/report.html).

Cook, R. I., Render, M. L., and Woods, D. D. (2000). Gaps in the continuity of care and progress on patient safety. *British Medical Journal,* 320, 791–794, March 18, 2000.

Hatlie, M. *Examining Errors in Health Care: Developing a Prevention, Education and Research Agenda, October 13–15, 1996* (Annenberg Center for Health Sciences at Eisenhower, Rancho Mirage, CA, 1996).

Hollnagel, E. (1983). *Human Error.* Position paper for NATO Conference on Human Error, Bellagio, Italy (available at http://www.ida.liu.se/~eriho/Publications O.html).

Hollnagel, E. (1993). *Human Reliability Analysis: Context and Control.* London: Academic Press.

Hollnagel, E. (1998). *Cognitive reliability and error analysis method–CREAM.* Oxford: Elsevier Science.

Hollnagel, E. (1999). Accidents and barriers. In J.-M. Hoc, P. Millot, E. Hollnagel and P. C. Cacciabue (Eds.), Proceedings of Lez Valenciennes, 28, 175–182 (Presses Universitaires de Valenciennes).

Hollnagel, E., and Amalberti, R. (2001). The Emperor's New Clothes or Whatever Happened to "Human Error"? Proceedings of the Fourth International Workshop on Human Error, Safety, and Systems Development (HESSD-01), Linköping, Sweden, June 11–12, 2001.

Klein, G. A. (1998). *Sources of Power: How People Make Decisions.* Cambridge, MA: MIT Press.

Kohn, L. T., Corrigan, J. M., and Donaldson, M., eds. (1998). *To Err is Human: Building a Safer Health System.* Washington D.C.: Institute of Medicine.

Leape, L. L. Institute of Medicine medical error figures are not exaggerated. *JAMA,* 2000;284:95–97.

Leape, L. L., Brennan, T. A., and Laird, N. M. Incidence of adverse events and negligence in hospitalized patients: Results of the Harvard Medical Practice Study II. *N Engl J Med,* 1991;324:377–384.

Mach, E. (1905). *Knowledge and Error.* Dordrecht: Reidel Publishing Company.

McDonald, C. J., Weiner, M., and Hui, S. L. Deaths due to medical errors are exaggerated in Institute of Medicine report. *JAMA,* 2000;284:93–95.

Rasmussen, J. (1986). *Information Processing and Human-Machine Interaction.* New York: North Holland.

Rasmussen, J. (1990a). The role of error in organizing behavior. *Ergonomics,* 33:1185–1199.

Rasmussen, J. (1990b). Human Error and the Problem of Causality in Analysis of Accidents. *Phil Trans R Soc Lond.* B 327, 449–462.

Rasmussen, J. (1994). Risk Management, Adaptation, and Design for Safety. In B. Brehmer and N.-E. Sahlin (Eds.) *Future Risks and Risk Management.* Dordrecht: Kluwer Academic.

Rasmussen, J. (1999). The concept of human error: Is it useful for the design of safe systems in health care? In C. Vincent and B. deMoll (Eds.), *Risk and Safety in Medicine.* London: Elsevier.

Reason, J. (1997). *Managing the Risks of Organizational Accidents.* Brookfield, VT: Ashgate Publishing.

Reason, J. (2001). *Assessing the Resilience of Health Care Systems to the Risk of Patient Mishaps.*

Rochlin, G. I. (1999). Safe operation as a social construct. *Ergonomics,* 42(11):1549–1560.

Suchman, L. (1987). *Plans and situated actions. The problem of human machine communication.* Cambridge: Cambridge University Press.

Svenson, O. (2001). Accident and incident analysis based on the accident evolution and barrier function model. *Cognition, Technology, and Work,* 3, 42–52.

Weick, K. E., Sutcliffe, K. M., and Obstfeld, D. (1999). Organizing for high reliability: Processes of collective mindfulness. *Res Organizational Behavior,* 21:23–81.

Woods, D. D. (2000a). Behind Human Error: Human Factors Research to Improve Patient Safety. *National Summit on Medical Errors and Patient Safety Research,* Quality Interagency Coordination Task Force and Agency for Healthcare Research and Quality, September 11, 2000.

Woods, D. D. (2000b). Patient Safety and Human Factors Opportunities. *Human Factors and Ergonomics Society Bulletin,* 43(5).

Woods, D. D., and Cook, R. I. (2001). From Counting Failures to Anticipating Risks: Possible Futures for Patient Safety. In L. Zipperer and S. Cushman, Editors, *Lessons on Patient Safety,* National Patient Safety Foundation.

Woods, D. D., and Cook, R. I. (2002). Nine Steps to Move Forward from Error. *Cognition, Technology, and Work,* 4(2):137–144.

Woods, D. D., and Shattuck, L. G. (2000). Distance supervision–local action given the potential for surprise. *Cognition, Technology, and Work,* 2, 86–96.

Woods, D. D., Johannesen, L., Cook, R. I., and Sarter, N. B. (1994). *Behind Human Error: Cognitive Systems, Computers and Hindsight.* Human Systems Ergonomic Information and Analysis Center, Wright Patterson Air Force Base, Dayton OH (available at http://iac.dtic.mil/hsiac/productBEHIND.html).