

Nine Steps to Move Forward from Error

D. D. Woods¹ and R. I. Cook²

¹*Institute for Ergonomics, Ohio State University, Columbus, Ohio, USA;* ²*Department of Anesthesia and Critical Care, University of Chicago, Chicago, Illinois, USA*

Abstract: Following celebrated failures stakeholders begin to ask questions about how to improve the systems and processes they operate, manage or depend on. In this process it is easy to become stuck on the label ‘human error’ as if it were an explanation for what happened and as if such a diagnosis specified steps to improve. To guide stakeholders when celebrated failure or other developments create windows of opportunity for change and investment, this paper draws on generalizations from the research base about how complex systems fail and about how people contribute to safety and risk to provide a set of Nine Steps forward for constructive responses. The Nine Steps forward are described and explained in the form of series of maxims and corollaries that summarize general patterns about error and expertise, complexity and learning.

Keywords: Error; Failure in complex systems; Patient safety

INTRODUCTION

Dramatic and celebrated failures are dreadful events that lead stakeholders to question basic assumptions about how the system in question works and sometimes breaks down. As each of these systems is under pressure to achieve new levels of performance and utilise costly resources more efficiently, it is very difficult for these stakeholders in high-risk industries to make substantial investments to improve safety. In this context, common beliefs and fallacies about human performance and about how systems fail undermine the ability to move forward.

On the other hand, over the years researchers on human performance, human–computer cooperation, teamwork, and organisational dynamics have turned their attention to high-risk systems studying how they fail and often succeed. While there are many attempts to summarise these research findings, stakeholders have a difficult time acting on these lessons, especially as they conflict with conventional views, require difficult trade-offs and demand sacrifices on other practical dimensions.

In this paper we use generalisations from the research base about how complex systems fail and how people contribute to safety as a guide for stakeholders when celebrated failure or other developments create windows of opportunity for change and investment. Nine steps forward are described and explained in the form of series of maxims and corollaries that summarise general patterns about error and expertise, complexity and learning. These ‘nine steps’

define one checklist for constructive responses when windows of opportunity to improve safety arise:

1. Pursue second stories beneath the surface to discover multiple contributors.
2. Escape the hindsight bias.
3. Understand work as performed at the sharp end of the system.
4. Search for systemic vulnerabilities.
5. Study how practice creates safety.
6. Search for underlying patterns.
7. Examine how change will produce new vulnerabilities and paths to failure.
8. Use new technology to support and enhance human expertise.
9. Tame complexity through new forms of feedback.

1. PURSUE SECOND STORIES

When an issue breaks with safety at the centre, it has been and will be told as a ‘first story’. First stories, biased by knowledge of outcome, are overly simplified accounts of the apparent ‘cause’ of the undesired outcome. The hindsight bias narrows and distorts our view of practice after-the-fact. As a result:

- there is premature closure on the set of contributors that lead to failure;

- the pressures and dilemmas that drive human performance are masked; and
- how people and organisations work to overcome hazards and make safety is obscured.

Stripped of all the context, first stories are appealing because they are easy to tell and locate the important 'cause' of failure in practitioners closest to the outcome. First stories appear in the press and usually drive the public, legal, and regulatory reactions to failure. Unfortunately, first stories simplify the dilemmas, complexities, and difficulties practitioners face and hide the multiple contributors and deeper patterns. The distorted view leads to proposals for 'solutions' that are weak or even counter-productive and blocks the ability of organisations to learn and improve.

For example, this pattern has been repeated over the last few years as the patient safety movement in health care has emerged. Each new celebrated failure produces general apprehension and calls for action. The first stories convince us that there are basic gaps in safety. They cause us to ask questions like: 'How big is this safety problem?' 'Why didn't someone notice it before?' and 'Who is responsible for this state of affairs?'

The calls to action based on first stories have followed a regular pattern:

- demands for increasing the general awareness of the issue among the public, media, regulators and practitioners ('we need a conference ...');
- calls for others to try harder or be more careful ('those people should be more vigilant about ...');
- insistence that real progress on safety can be made easily if some local limitation is overcome ('we can do a better job if only ...');
- calls for more extensive, more detailed, more frequent and more complete reporting of problems ('we need mandatory incident reporting systems with penalties for failure to report');
- calls for more technology to guard against erratic people ('we need computer order entry, bar coding, electronic medical records, etc.').

Actually, first stories represent a kind of reaction to failure that attributes the cause of accidents to narrow proximal factors, usually 'human error'. They appear to be attractive explanations for failure, but they lead to sterile responses that limit learning and improvement (blame and punishment; e.g., 'we need to make it so costly for people that they will have to ...').

When we observe this process begin to play out over an issue or celebrated event, the constructive response is very simple. To make progress on safety requires going beyond first stories to discover what lies behind the term 'human error' (Cook et al 1998). At the broadest level, our role is to

help others develop the deeper 'second story'. This is the most basic lesson from past research on how complex systems fail. When one pursues second stories the system in question looks very different and one can begin to see how the system moves toward, but is usually blocked from, accidents. Through these deeper insights learning occurs and the process of improvement begins.

1. The Second Story Maxim

Progress on safety begins with uncovering 'second stories'.

The remaining steps specify how to extract the second stories and how they can lead to safety improvement.

2. ESCAPE FROM HINDSIGHT BIAS

The first story after celebrated accidents tells us nothing about the factors that influence human performance before the fact. Rather the first story represents how we, with knowledge of outcome and as stakeholders, react to failures. Reactions to failure are driven by the consequences of failure for victims and other stakeholders and by the costs associated with changes made to satisfy stakeholders that the threats represented by the failure are under sufficient control. This is a social and political process about how we attribute 'cause' for dreadful and surprising breakdowns in systems that we depend on (Woods et al 1994; Schon 1995).

Knowledge of outcome distorts our view of the nature of practice. We simplify the dilemmas, complexities and difficulties practitioners face and how they usually cope with these factors to produce success. The distorted view leads people to propose 'solutions' that actually can be counterproductive

- (a) if they degrade the flow of information that supports learning about systemic vulnerabilities; and
- (b) if they create new complexities to plague practice.

Research-based approaches fundamentally use various techniques to escape from hindsight bias. This is a crucial prerequisite for learning to occur.

3. UNDERSTAND THE WORK PERFORMED AT THE SHARP END OF THE SYSTEM

When we start to pursue the 'second story', our attention is directed to people working at the sharp end of a system such as health care. The substance of the second story resides at the sharp end of the system as organisational, economic, human and technological factors play out to create outcomes. Sharp end practitioners who work in this setting face a variety of difficulties, complexities, dilemmas and trade-offs and are called on to achieve

multiple, often conflicting, goals. Safety is created here at the sharp end as practitioners interact with the hazardous processes inherent in the field of activity in the face of the multiple demands and using the available tools and resources.

To follow second stories, one looks more broadly than a single case to understand how practitioners at the sharp end function – the nature of technical work as experienced by the practitioner in context. This is seen in research as a *practice-centred view of technical work in context* (Barley and Orr 1997).

Ultimately, all efforts to improve safety will be translated into new demands, constraints, tools or resources that appear at the sharp end. Improving safety depends on investing in resources that support practitioners in meeting the demands and overcoming the inherent hazards in that setting.

II. The Technical Work in Context Maxim

Progress on safety depends on understanding how practitioners cope with the complexities of technical work.

When we shift our focus to technical work in context, we begin to ask how people usually succeed. Ironically, understanding the sources of failure begins with understanding how practitioners coordinate activities in ways that help them cope with the different kinds of complexities they experience. Interestingly, the fundamental insight 20 years ago that launched the New Look behind the label human error was to see human performance at work as human adaptations directed to cope with complexity (Rasmussen 1986).

One way that some researchers have summarised the results that lead to Maxim II is that:

'The potential cost of misunderstanding technical work' is the risk of setting policies whose actual effects are 'not only unintended but sometimes so skewed that they exacerbate the problems they seek to resolve'. 'Efforts to reduce error misfire when they are predicated on a fundamental misunderstanding of the primary sources of failures in the field of practice [systemic vulnerabilities] and on misconceptions of what practitioners actually do.' (Barley and Orr 1997, p. 18; emphasis added)

Three corollaries to the Technical Work in Context Maxim can help focus efforts to understand technical work as it effects the potential for failure:

Corollary IIA. Look for Sources of Success

To understand failure, understand success in the face of complexities.

Failures occur in situations that usually produce successful outcomes. In most cases, the system produces success despite opportunities to fail. To understand failure requires understanding how practitioners usually achieve success in the face of demands, difficulties, pressures and dilemmas.

Indeed, it is clear that success and failure flow from the same sources (Rasmussen 1985).

Corollary IIB. Look for Difficult Problems

To understand failure, look at what makes problems difficult.

Understanding failure and success begins with understanding what makes problems difficult. Cook et al (1998) illustrated the value of this approach in their tutorial for health care, 'The tale of two stories'. They used three uncelebrated second stories from health care to show progress depended on investigations that identified the factors that made certain situations more difficult to handle and then explored the individual and team strategies used to handle these situations. As the researchers began to understand what made certain kinds of problems difficult, how expert strategies were tailored to these demands and how other strategies were poor or brittle, new concepts were identified to support and broaden the application of successful strategies.

Corollary IIC. Be Practice-Centred – Avoid the Psychologist's Fallacy

Understand the nature of practice from the practitioner's point of view.

It is easy to commit what William James called over one hundred years ago the Psychologist's Fallacy (1890). Updated to today, this fallacy occurs when well-intentioned observers think that their distant view of the workplace captures the actual experience of those who perform technical work in context. Distant views can miss important aspects of the actual work situation and thus can miss critical factors that determine human performance in that field of practice. To avoid the danger of this fallacy, cognitive anthropologists use research techniques based on an 'emic' or practice-centred perspective (Hutchins, 1995). Researchers on human problem solving and decision making refer to the same concept with labels such as process tracing and naturalistic decision making (Klein et al 1993).

It is important to distinguish clearly that doing technical work expertly is not the same thing as expert understanding of the basis for technical work. This means that practitioners' descriptions of how they accomplish their work are often biased and cannot be taken at face value. For example, there can be a significant gap between people's descriptions (or self-analysis) of how they do something and observations of what they actually do.

Since technical work in context is grounded in the details of the domain itself, it is also insufficient to be expert in human performance in general. Understanding technical work in context requires (1) in-depth appreciation of the pressures and dilemmas practitioners face and the resources and adaptations practitioners bring to bear to

accomplish their goals, and also (2) the ability to step back and reflect on the deep structure of factors that influence human performance in that setting. Individual observers rarely possess all of the relevant skills, so that progress on understanding technical work in context and the sources of safety inevitably requires interdisciplinary cooperation.

In the final analysis, successful practice-centred inquiry requires a marriage between the following three factors:

- the view of practitioners in context;
- technical knowledge in that area of practice; and
- knowledge of general results/concepts about the various aspects of human performance that play out in that setting.

Interdisciplinary collaborations have played a central role as health care has begun to make progress on iatrogenic risks and patient safety recently (e.g., Hende 1999).

This leads us to note a third maxim:

III. The Interdisciplinary Synthesis Maxim

Progress on safety depends on facilitating interdisciplinary investigations.

4. SEARCH FOR SYSTEMIC VULNERABILITIES

Through practice-centred observation and studies of technical work in context, safety is not found in a single person, device or department of an organisation. Instead, safety is created and sometimes broken in systems, not individuals (Cook et al 2000). The issue is finding systemic vulnerabilities, not flawed individuals.

IV. The Systems Maxim

Safety is an emergent property of systems and not of their components.

Examining technical work in context with safety as our purpose, one will notice many hazards, complexities, gaps, trade-offs, dilemmas and points where failure is possible. One will also begin to see how practice has evolved to cope with these kinds of complexities. After elucidating complexities and coping strategies, one can examine how these adaptations are limited, brittle and vulnerable to breakdown under differing circumstances. Discovering these vulnerabilities and making them visible to the organisation is crucial if we are to anticipate future failures and institute change to head them off.

A repeated finding from research on complex systems is that practitioners and organisations have opportunities to recognise and react to threats to safety. Precursor events may serve as unrecognised 'dress rehearsals' for future accidents. The accident itself often evolves through time so that practitioners can intervene to prevent negative

outcomes or to reduce their consequences. Doing this depends on being able to recognise accidents-in-the-making. However, it is difficult to act on information about systemic vulnerabilities as potential interventions often require sacrificing some goals under certain circumstances (e.g., productivity) and therefore generate conflicts within the organisation.

Detection and recovery from incipient failures is a crucial part of achieving safety at all levels of an organisation – a corollary to the Systems Maxim. Successful individuals, groups and organisations, from a safety point of view, learn about complexities and the limits of current adaptations and then have mechanisms to act on what is learned, despite the implications for other goals (Rochlin 1999; Weick and Roberts 1993).

Corollary IVA. Detection and Recovery Are Critical to Success

Understand how the system of interest supports (or fails to support) detection and recovery from incipient failures.

In addition, this process of feedback, learning and adaptation should go on continuously across all levels of an organisation. With change, some vulnerabilities decay while new paths to failure emerge. To track the shifting pattern requires getting information about the effects of change on sharp end practice and about new kinds of incidents that begin to emerge. If the information is rich enough and fresh enough, it is possible to forecast future forms of failure, to share schemes to secure success in the face of changing vulnerabilities. Producing and widely sharing this sort of information may be one of the hallmarks of a culture of safety (Weick et al. 1999).

However, establishing a flow of information about systemic vulnerabilities is quite difficult because it is frightening to consider how all of us, as part of the system of interest, can fail. Repeatedly, research notes that blame and punishment will drive this critical information underground. Without a safety culture, systemic vulnerabilities become visible only after catastrophic accidents. In the aftermath of accidents, learning also is limited because the consequences provoke first stories, simplistic attributions and shortsighted fixes.

Understanding the 'systems' part of safety involves understanding how the system itself learns about safety and responds to threats and opportunities. In organisational safety cultures, this activity is prominent, sustained and highly valued (Cook 1999). The learning processes must be tuned to the future to recognise and compensate for negative side effects of change and to monitor the changing landscape of potential paths to failure. Thus, the Systems Maxim leads to the corollary to examine how the organisation at different levels of analysis supports or fails to support the process of feedback, learning and adaptation.

Corollary IVB. Learning how to Learn

Safe organisations deliberately search for and learn about systemic vulnerabilities.

The future culture all aspire to is one where stakeholders can learn together about systemic vulnerabilities and work together to address those vulnerabilities, *before celebrated failures occur* (Woods, 2000).

5. STUDY HOW PRACTICE CREATES SAFETY

Typically, reactions to failure assume the system is 'safe' (or has been made safe) inherently and that overt failures are only the mark of an unreliable component. But what is irreducible is uncertainty about the future, change and finite resources. As a result, all systems confront inherent hazards, trade-offs and are vulnerable to failure. Second stories reveal how practice is organised to allow practitioners to create success in the face of threats. Individuals, teams and organisations are aware of hazards and adapt their practices and tools to guard against or defuse these threats to safety. It is these efforts that 'make safety'. This view of the human role in safety has been a part of complex systems research since its origins (see Rasmussen et al 1994, ch. 6). The Technical Work in Context maxim tell us to study how practice copes with hazards and resolves trade-offs, for the most part succeeding yet in some situations failing.

However, the adaptations of individuals, teams and organisations can be limited or stale so that feedback about how well adaptations are working or about how the environment is changing is critical. Examining the weaknesses and strengths, costs and benefits of these adaptations points to the areas ripe for improvement. As a result, progress depends on studying *how practice creates safety in the face of challenges – expertise in context* (Feltovich et al 1997; Klein, 1998).

6. SEARCH FOR UNDERLYING PATTERNS

In the discussions of some particular episode or 'hot button' issue it is easy for commentators to examine only surface characteristics of the area in question. Progress has come from going beyond the surface descriptions (the phenotypes of failures) to discover underlying patterns of systemic factors (genotypical patterns; see Hollnagel 1993; 1998).

V. The Genotypes Maxim

Progress on safety comes from going beyond the surface descriptions (the phenotypes of failures) to discover underlying patterns of systemic factors (genotypical patterns).

Genotypes are concepts and models about how people, teams and organisations coordinate information and

activities to handle evolving situations and cope with the complexities of that work domain. These underlying patterns are not simply about knowledge of one area in a particular field of practice. Rather, they apply, test and extend knowledge about how people contribute to safety and failure and how complex systems fail by addressing the factors at work in this particular setting. As a result, when we examine technical work, search for underlying patterns by contrasting sets of cases.

7. EXAMINE HOW ECONOMIC, ORGANISATIONAL AND TECHNOLOGICAL CHANGE WILL PRODUCE NEW VULNERABILITIES AND PATHS TO FAILURE

As capabilities, tools, organisations and economic pressures change, vulnerabilities to failure change as well.

VI. Safety is a Dynamic Process Maxim

The state of safety in any system always is dynamic.

Systems exist in a changing world. The environment, organisation, economics, capabilities, technology, management and regulatory context all change over time. This backdrop of continuous systemic change ensures that hazards and how they are managed are constantly changing. Plus, the basic pattern in complex systems is a drift toward failure as planned defences erode in the face of production pressures and change. As a result, when we examine technical work in context, we need to understand how economic, organisational and technological change can create new vulnerabilities in spite of or in addition to providing new benefits.

Research reveals that organisations that manage potentially hazardous technical operations remarkably successfully create safety by anticipating and planning for unexpected events and future surprises. These organisations did not take past success as a reason for confidence. Instead they continued to invest in anticipating the changing potential for failure because of the deeply held understanding that their knowledge base was fragile in the face of the hazards inherent in their work and the changes omnipresent in their environment (Rochlin 1999).

Research results have pointed to several corollaries to the Dynamic Process Maxim.

Corollary VIA. Law of Stretched Systems

Under resource pressure, the benefits of change are taken in increased productivity, pushing the system back to the edge of the performance envelope.

Change occurs to improve systems. However, because the system is under resource and performance pressures from

stakeholders, we tend to take the benefits of change in the form of increased productivity and efficiency and not in the form of a more resilient, robust and therefore safer system (Rasmussen 1986). Researchers in the field speak of this observation as follows: systems under pressure move back to the 'edge of the performance envelope' or the Law of Stretched Systems (Woods 2002):

... we are talking about a law of systems development, which is every system operates, always at its capacity. As soon as there is some improvement, some new technology, we stretch it ... (Hirschhorn 1997)

Change under resource and performance pressures tends to increase coupling, that is, the interconnections between parts and activities, in order to achieve greater efficiency and productivity. However, research has found that increasing coupling also increases operational complexity and increases the difficulty of the problems practitioners can face. Jens Rasmussen (1986) and Charles Perrow (1984) provided some of the first accounts of the role of coupling and complexity in modern system failures.

Corollary VIB. Increasing Coupling Increases Complexity

Increased coupling creates new cognitive and collaborative demands and new forms of failure.

Increasing the coupling between parts in a process changes how problems manifest, creating or increasing complexities such as more effects at a distance, more and faster cascades of effects, tighter goal conflicts, more latent factors. As a result, increased coupling between parts creates new cognitive and collaborative demands which contribute to new forms of failure (Woods 1988; Woods and Patterson 2000).

Because all organisations are resource limited to one degree or another, we are often concerned with how to prioritise issues related to safety. The Dynamics Process Maxim suggests that we should consider focusing our resources on anticipating how economic, organisational and technological change could create new vulnerabilities and paths to failure. Armed with this knowledge we can address or eliminate these new vulnerabilities at a time when intervention is less difficult and less expensive (because the system is already in the process of change). In addition, these points of change are at the same time opportunities to learn how the system actually functions.

VII. The Window of Opportunity Maxim

Use periods of change as windows of opportunity to anticipate and treat new systemic vulnerabilities.

8. USE NEW TECHNOLOGY TO SUPPORT AND ENHANCE HUMAN EXPERTISE

The notion that it is easy to get 'substantial gains' through computerisation is common in many fields. The implication is that computerisation *by itself* reduces human error and system breakdown. Any difficulties that are raised about the computerisation process become mere details to be worked out later.

VIII. Joint Systems Maxim

But this idea, which Woods stated a long time ago as 'a little more technology will be enough', has not turned out to be the case in practice (for an overview see Woods et al 1994, ch. 5 or Woods and Tinapple 1999). Those pesky details turn out to be critical in whether the computerisation creates new forms of failure. New technology can help and can hurt, often at the same time depending on how the technology is used to support technical work in context.

Basically, it is the underlying complexity of operations that contributes to the human performance problems. Improper computerisation can simply exacerbate or create new forms of complexity to plague operations. The situation is complicated by the fact the new technology often has benefits at the same time that it creates new vulnerabilities.

People and computers are not separate and independent, but are interwoven into a distributed system that performs cognitive work in context.

The key to skilful as opposed to clumsy use of technological possibilities lies in understanding the factors that lead to expert performance and the factors that challenge expert performance. The irony is that once we understand the factors that contribute to expertise and to breakdown, we then will understand how to use the powers of the computer to enhance expertise. This is illustrated in uncelebrated second stories in research on human performance in medicine, explored in Cook et al (1998). On the one hand, new technology creates new dilemmas and demands new judgments, but, on the other hand, once the basis for human expertise and the threats to that expertise had been studied, technology was an important means to the end of enhanced system performance.

We can achieve substantial gains by understanding the factors that lead to expert performance and the factors that challenge expert performance. This provides the basis to change the system, for example, through new computer support systems and other ways to enhance expertise in practice.

As a result, when we examine technical work, *understand the sources of and challenges to expertise in context*. This is

crucial to guide the skilful, as opposed to clumsy use of technological possibilities.

Corollary VIII.A. There is no Neutral in Design

In design, we either support or hobble people's natural ability to express forms of expertise (Woods 2002).

9. TAME COMPLEXITY THROUGH NEW FORMS OF FEEDBACK

The theme that leaps out from past results is that failure represents *breakdowns in adaptations* directed at coping with complexity. Success relates to organisations, groups and individuals who are skilful at recognising the need to adapt in a changing, variable world and in developing ways to adapt plans to meet these changing conditions despite the risk of negative side effects.

Recovery before negative consequences occur, adapting plans to handle variations and surprise, and recognising side effects of change are all critical to high resilience in human and organisational performance. Yet, all of these processes depend fundamentally on the ability to see the emerging effects of decisions, actions, policies – feedback, especially feedback about the future. In general, increasing complexity can be balanced with improved feedback. Improving feedback is a critical investment area for improving human performance and guarding against paths toward failure. The constructive response to issues on safety is to study where and how to invest in better feedback.

This is a complicated subject since better feedback is

- integrated to capture relationships and patterns, not simply a large set of available data elements;
- event based to capture change and sequence, not simply the current values on each data channel;
- future oriented to help people assess what could happen next, not simply what has happened;
- context sensitive and tuned to the interests and expectations of the monitor.

Feedback at all levels of the organisation is critical because the basic pattern in complex systems is a drift toward failure as planned defences erode in the face of production pressures and change. The feedback is needed to support adaptation and learning processes. Ironically, feedback must be tuned to the future to detect the emergence of the drift toward failure pattern, to explore and compensate for negative side effects of change, and to monitor the changing landscape of potential paths to failure. To achieve this organisations need to develop and support mechanisms that *create foresight* about the changing shape of risks, before anyone is injured.

References

- Barley S, Orr J (eds) (1997). *Between craft and science: technical work in US settings*. IRL Press, Ithaca, NY.
- Cook RI (1999). Two years Before the Mast: Learning How to Learn about Patient Safety. In W. Hendee, (ed.), *Enhancing Patient Safety and Reducing Errors in Health Care*. National Patient Safety Foundation, Chicago IL.
- Cook RI, Woods DD, Miller C. A tale of two stories: contrasting views on patient safety. National Patient Safety Foundation, Chicago IL, April 1998 (available at www.npsf.org/exec/report.html).
- Cook RI, Render M, Woods DD (2000). Gaps: learning how practitioners create safety. *British Medical Journal* 320:791–794.
- Feltovich P, Ford K, Hoffman R (eds) (1997). *Expertise in context*. MIT Press, Cambridge MA.
- Hendee W (ed) (1999). *Enhancing patient safety and reducing errors in health care*. National Patient Safety Foundation, Chicago, IL.
- Hirschhorn L (1997). Quoted in Cook RI, Woods DD and Miller C (1998). *A Tale of Two Stories: Contrasting Views on Patient Safety*. National Patient Safety Foundation, Chicago IL, April 1998.
- Hollnagel E (1993). *Human reliability analysis: context and control*. Academic Press, London.
- Hollnagel E (1998). *Cognitive reliability method and error analysis method*. Elsevier, New York.
- Hutchins E (1995). *Cognition in the wild*. MIT Press, Cambridge, MA.
- James W (1890). *Principles of psychology*. H. Holt & Co. NY.
- Klein G (1998). *Sources of power: how people make decisions*. MIT Press, Cambridge, MA.
- Klein GA, Orasanu J, Calderwood R (eds) (1993). *Decision making in action: models and methods*. Ablex, Norwood, NJ.
- Perrow C (1984). *Normal accidents*. Basic Books, NY.
- Rasmussen J (1985). Trends in human reliability analysis. *Ergonomics* 28(8):1185–1196.
- Rasmussen J (1986). *Information processing and human-machine interaction: an approach to cognitive engineering*. North-Holland, New York.
- Rasmussen J, Pejtersen AM, Goodstein LP. At the periphery of effective coupling: human error. In *Cognitive systems engineering*. Wiley, New York, pp 135–159.
- Rochlin GI (1999). Safe operation as a social construct. *Ergonomics* 42(11):1549–1560.
- Schon DA (1995). Causality and causal inference in the study of organizations. In Goodman RF, Fisher WR (eds). *Rethinking knowledge: reflections across the disciplines*. State University of New York Press, Albany, pp 000–000.
- Weick KE, Roberts KH (1993). Collective mind and organizational reliability: the case of flight operations on an aircraft carrier deck. *Administration Science Quarterly* 38:357–381.
- Weick KE, Sutcliffe KM, Obstfeld D (1999). Organizing for high reliability: processes of collective mindfulness. *Research in Organizational Behavior* 21:81–123.
- Woods DD (2000). Behind human error: human factors research to improve patient safety. In National summit on medical errors and patient safety research, Quality Interagency Coordination Task Force and Agency for Healthcare Research and Quality, 11 September 2000.
- Woods DD (2002). Steering the Reverberations of Technology Change on Fields of Practice: Laws that Govern Cognitive Work. In Proceedings of the 24th Annual Meeting of the Cognitive Science Society, August 2002. [Plenary Address].
- Woods DD, Johannesen L, Cook RI, Sarter N (1994). Behind human error: cognitive systems, computers and hindsight. Crew Systems Ergonomic Information and Analysis Center, WPAFB, Dayton OH, 1994 (at <http://iac.dtic.mil/hsiac/product/BEHIND.htm>)
- Woods DD (1988). Coping with complexity: the psychology of human behavior in complex systems. In Goodstein LP, Andersen HB, Olsen SE

- (eds). Mental models, tasks and errors. Taylor & Francis, London, pp 128–148.
- Woods DD, Patterson ES (2002). How unexpected events produce an escalation of cognitive and coordinative demands. In Hancock PA, Desmond P (eds). Stress workload and fatigue. Erlbaum L, Hillsdale, NJ (in press).
- Woods DD, Tinapple D (1999). W^3 : watching human factors watch people at work. Presidential address, 43rd annual meeting of the Human Factors

and Ergonomics Society, 28 September 1999 (multimedia production at <http://csel.eng.ohio-state.edu/hf99/>).

Correspondence and offprint requests to: D. D. Woods, Cognitive Systems Engineering Laboratory, Department of Industrial and Systems Engineering, Ohio State University, 1971 Neil Avenue, Columbus, OH 43210, USA. Email: woods.2@osu.edu