High-reliability organizations invest in resilience

Dekker, Sidney; Zimmermann, Verena; Woods, David D.

DOI
10.1016/B978-0-12-420139-2.00006-X

Publication date
2022

Document Version
Final published version

Published in
Human Factors in Aviation and Aerospace, Third Edition

CHAPTER

3

High-reliability organizations invest in resilience

Sidney Dekker[a,b], *Verena Zimmermann*[c], *and David D. Woods*[d]

[a]School of Humanities, Languages and Social Science, Griffith University, Brisbane, QLD, Australia, [b]Faculty of Aerospace Engineering, Delft University of Technology, Delft, The Netherlands, [c]Department of Humanities, Social and Political Sciences, ETH Zürich, Zürich, Switzerland, [d]Integrated Systems Engineering, The Ohio State University, Columbus, OH, United States

High-reliability theory describes the extent and nature of the effort that people, at all levels in an organization, have to engage in to ensure consistently safe operations despite their inherent complexity and risks. It is founded on an empirical research base that shows how safety originates in large part in the operational and leadership activities of people. The high-reliability organizations (HROs) perspective is relevant here, since aviation has done an effective job in institutionalizing and systematizing its learning from incidents and accidents. HRO, however, tries to go further—pulling learning forward in time, studying how managerial and operational activities can encourage the exploration and exchange of safety-related information. The aim is to pick up early signs that trouble may be on the horizon and then be able to make modifications without having to wait for the more obvious signs of failure in the form of incidents or accidents. HROs ideally are able to stay curious about their own operations and keep wondering why they are successful. They stay open-minded about the sources of risk, try to remain complexly sensitized to multiple sources of safety information, keep inviting doubt and minority opinion, and stay ambivalent toward the past, so that confidence gleaned from previous results is not taken as a guarantee of future safety (Weick, 1993).

https://doi.org/10.1016/B978-0-12-420139-2.00006-X

This chapter first considers the origins and ideas of HRO, then addresses the "reliability" part of its label as applied to aviation safety, and concludes with how Resilience Engineering represents a kind of action agenda of HRO. For a more detailed description of HRO, the reader is also referred to a related article in the book "The Foundations of Safety Science" (Schochlow & Dekker, 2019) that parts of this chapter are based on. With an emerging set of techniques and models to track how organizations learn, adapt, and change without waiting for major failures, Resilience Engineering can show where overconfidence in past results may be occurring, where minority viewpoints may risk getting downplayed, and where acute performance or production demands may outweigh chronic safety concerns. These things are important for aviation because of its high safety and its complexity: Accidents do not occur much any longer and have long since ceased to be the result of single component failures. Rather, they emerge from the system's organized complexity (Amalberti, 2001). To anticipate whether aviation systems can keep coping with change and complexity takes more than tracking the system's reliability made up of individual component behaviors (Woods, 2022).

The beginnings of HRO research

The roots of HRO can be found in a group of researchers from the University of California, the so-called Berkeley group (Rochlin, LaPorte, & Roberts, 1987). Weick, a social and an organizational psychologist, also became involved in HRO research and added ideas to the approach. In 1995, Roberts started working with another group of safety scientists at the California State University in San Bernadino. This San Bernadino group in turn developed further HRO methods and applied HRO to industries where failures are business failures such as in the financial sector. Recent developments include the application of Resilience Engineering principles in HRO. As a result, HRO is not one theory described by one author. Rather, it is an evolving set of common principles, developed, adapted and renamed over time by several different authors—with slightly different methods, measures of safety or assumptions about the human factor.

La Porte, Roberts and Rochlin, as drivers of the original Berkeley group, launched their research with a project "concerned with the design and management of hazardous organizations that achieve extremely high levels of reliable and safe operations" (Roberts, 1989, p. 111). It involved the US Air Traffic Control system, electric grids and US Navy nuclear aircraft carriers (Rochlin et al., 1987). The group aimed to answer the question whether there were high risk organizations operating nearly error-free. And if so, what were the characteristics of these organizations? What did they do to prevent errors? Roberts clustered organizations on the two dimensions technological risk and reliability. Organizations were labeled

HROs when they repeatedly could have failed with catastrophic results but did not (Roberts, 1990). This was measured not in terms of absolute safety but with the help of relative measures of performance, such as collisions related to number of movements or outages related to hours of service. The organizations studied were also characterized by complexity, tight coupling, and highly interconnected technologies (LaPorte & Consolini, 1991).

When HRO researchers first set out to examine how safety is created and maintained in such complex systems, they took an approach that can be found in parts of aviation human factors today. They focused on errors and other negative indicators, such as incidents, assuming that these were the basic units that people in these organizations used to map the physical and dynamic safety properties of their production technologies, ultimately to control risk (Rochlin, 1999). The assumption turned out wrong: they were not. Operational people, those who work at the sharp end of an organization, hardly defined safety in terms of risk management or error avoidance. Ensuing empirical work by HRO, stretching across decades and a multitude of high-hazard, complex domains (aviation, nuclear power, utility grid management, navy) would paint a more complex, and in many ways a more constructive picture with safety not being the *absence* of negatives, but rather the *presence* of certain activities to manage risk. HRO began to describe how operational safety—how it is created, maintained, discussed, mythologized—should be captured as much more than the control of negatives. As Rochlin (1999, p. 1549) put it,

> the culture of safety that was observed is a dynamic, intersubjectively constructed belief in the possibility of continued operational safety, instantiated by experience with anticipation of events that could have led to serious errors, and complemented by the continuing expectation of future surprise.

The creation of safety, in other words, involves a belief about the possibility to continue operating safely. This belief is built up and shared among those who do the work every day. It is moderated or even held up in part by the constant preparation for future surprise—preparation for situations that may challenge people's current assumptions about what makes their operation risky or safe. And yes, it is also a belief punctuated by encounters with risk. But errors, or any other negatives, function at most as the narrative spice that keeps the belief flavorful and worth sharing. They turned out not to be its main substance.

Findings of the original HRO group

La Porte, Roberts and Rochlin did not focus on human performance at the level of the individual. Their interest was in human performance at a group- or organizational level (LaPorte & Consolini, 1991). And in contrast to other approaches, HRO research didn't take accidents as a dependent

variable but rather the nearly error-free operations. Researchers were not so interested in what caused accidents, but rather what contributed to high reliability and safety levels (Roberts, 1990). Their preliminary findings revealed some characteristics present in all or most of the HROs studied (La Porte, 1996; Roberts, 1989, 1990):

- *Organizational structure and rules*: The organizational structures of HROs seemed to be marked by flexibility and redundancy. This was for example shown in parallel and overlapping processes, backups, cross-checks and skills redundancy (e.g., through job rotation). Redundancy was seen as a measure to recover from errors and ensure reliability despite the shortcomings of human individuals.

  Structures were described in terms of "nested authority," which refers to the adaptive change from a hierarchical organization during routine operations to patterns deferring to experience in demanding situations. This was interpreted as one way to handle coexisting interdependence and complexity in HROs (Roberts & Gargano, 1989).

- *Operational decision-making and communication*: The flexible adaption of organizational structures also affected decision-making and communication patterns. The shift in structure went along with a shift in authority. If called for by the circumstances, decision-making migrated to the people with the information and experience at hand to the locus of the problem. This was also seen as one way to cope with the tension between quick and accurate decision making, or in other words efficiency and reliability. Several channels of communication seemed to be used constantly to inform others, to maintain the "big picture" and to keep the team integrated. Communication therefore also served as a safety function (Rochlin, 1999).
- *A "culture" of high reliability*: This culture is viewed by Rochlin (1999, p. 1549) as a "dynamic, intersubjectively constructed belief in the possibility of continued operational safety, instantiated by experience with anticipation of events that could have led to serious errors, and complemented by the continuing expectation of future surprise." People in an HRO showed high levels of personal engagement and a "sense of mission" (La Porte, 1996). Compared to other organizations the willingness to report errors and potential failures seemed to be untypically high in HROs. Reporters were not blamed but rather rewarded for their initiative and contribution.
- *The adaption to technology*: Roberts (1989, p. 121) noted a "remarkable resilience," combined with a resistance of decision-makers to adopt new technologies. In some of the empirical studies, researchers were

able to show that the resulting mix of old and new technologies ensured operations even without electricity and continuing practice of operator skills.

Based on the preliminary characteristics, Roberts (1990) described the following strategies used by this HRO to counter high interactive complexity and tight coupling (Perrow, 1984):

Complexity was countered with:

- continuous training
- job design strategies to keep function separate
- main direct information sources

Tight coupling was countered with:

- redundancy
- hierarchical differentiation
- bargaining

Complexity and tight coupling at the same time were countered with:

- redundancy
- accountability
- responsibility
- "culture" of reliability

Nevertheless, La Porte (1996) concluded that the characteristics of HROs found in the case studies were not ready for simple application to other organizations which intended to become HROs. He stated that these characteristics might be necessary but not sufficient. This again shows the more descriptive character of the initial HRO project.

Mindfulness

Working off the original empirical results from the Berkeley group, Weick and Sutcliffe introduced mindfulness into the HRO lexicon as the "capability to induce a rich awareness of discriminatory detail and a capacity for action" (Weick, Sutcliffe, & Obstfeld, 2008, p. 37). They suggested that mindfulness increases the capability of people in effective HROs to detect and handle unexpected events. This in turn leads to the reliability of the organization, hence its nearly error-free performance in a demanding social and political environment and despite the huge potential for errors in its complex processes and technologies. Weick regarded reliability as the ability to perform in the same way despite of changing conditions for operation. Thus, the operators' activity is required to vary according to the circumstances whereas the cognitive processes making sense of the activity remain stable (Weick et al., 2008). If mindfulness is a key to detect

and manage unexpected failures and therefore create reliability, how can it be encouraged or created? Weick and colleagues suggested five stable cognitive processes of mindful organizing present in effective HROs (Weick et al., 2008, p. 37).

1. *Preoccupation with failure*: This process describes that in effective HROs people have an "apparent ongoing focus on failure" (Weick et al., 2008, p. 38) instead of success. In this sense, also near failures, such as near misses, are seen as a danger to safety. Furthermore, mindfulness requires interpretative work of so-called weak signals. Weak signals were defined as signals that at the time of observation had no apparent clear or direct link to a potential danger or events that were judged to have happened in such rare, unlikely conditions that they wouldn't happen again (Vaughan, 2002). For example, after Weick et al. (2008), one lapse could already be interpreted as a weak signal for further vulnerabilities in other parts of the system. That is, in HROs people in effective HROs tend to generalize failures or weak signals respectively instead of localizing and reducing safety problems. Success is not interpreted as demonstrating competence, as this would lead to inertia, inattention and errors going undetected. Instead, people in effective HROs are even assuming that current successful operation makes future success less likely (Weick et al., 2008, p. 41).
2. *Reluctance to simplify interpretations*: Weick et al. (2008) identified simplification as a potential danger to HROs since it leads to a negligence of data, reduction of precautions and an increased likelihood of undesired surprises. Effective HROs undertake effort to limit simplifications. That is done, e.g., through processes of renewal and revision, job rotation, cross-checks or selecting new employees with nontypical prior experience.
3. *Sensitivity to operations*: This process refers to having "the integrated big picture of operations in the moment" (Weick et al., 2008, p. 43). This "big picture" is difficult to achieve and to maintain. It therefore requires active effort. Because of the individual's limited range and varying focus, the "big picture" is a shared accomplishment. Ways to achieve sensitivity to operations include collective story building, shared mental representations and knowledge of system parameters.
4. *Commitment to resilience*: Weick et al. (2008, p. 46) pointed out that "people deal with surprises, not only by anticipation that weeds them out in advance, but also by resilience that responds to them as they occur." This cognitive process is preventative in that operators in effective HROs do not wait for errors to happen, but prepare for unexpected events beforehand. Accordingly, "they pay attention both

to error-prevention and to error-containment" (Weick et al., 2008, p. 47). This is visible in the support of improvisation, the recombination of actions and the parallel belief and doubt in past experiences in effective HROs.

5. *Deference to expertise*: Another feature of effective HROs is that organizations that are normally structured in a hierarchical way are able to subordinate hierarchical rank to expertise if demanded. Depending on the situation, the organization allows an adaptive loosening of filters so that operators at the bottom of the hierarchical structure can rise to the top level of decision-making temporarily. This is also referred to as underspecification of structures (Weick et al., 2008).

Intriguingly, the label "high reliability" grew increasingly at odds with the findings this school produced. What was a research effort to examine how high-risk systems can produce high-reliability outcomes despite their inherent danger (i.e., measured in terms of reducing negatives, or failure events), transmogrified into a discovery of safety as a reflexive social construct that challenged virtually all available methodological, ontological and theoretical guidance available at the time. Safety, HRO concluded, does not exist "out there," independent from the minds or actions of the people who create it through their practice, simply to be discovered, laid bare, by those with the right measuring instrument. Knowing about safety cannot be synonymous with a tabulation of "objective" measures from real-world performance. And, indeed, the predictive value of such measures is generally quite disappointing. While ensuring consistent and reliable component performance (both human and machine) has been a hugely important contributor to the successful safety record of aviation to date, there are limits to this approach, particularly when it comes to avoiding complex system accidents that emerge from the normal functioning of already almost totally safe transportation systems (Amalberti, 2001).

Reliability and safety

Safety is of course not the same as reliability (Woods, 2022). A part can be reliable, but in and of itself it cannot be safe. It can perform its stated function to the expected level or amount, but it is context, the context of other parts, of the dynamics and the interactions and cross-adaptations between parts, that make things safe or unsafe. Reliability as an engineering property can be expressed as a component's failure rate or probabilities over a period of time. In other words, it addresses the question of whether a component lives up to its prespecified performance criteria. Organizationally, reliability is often associated with a reduction in variability, and concomitantly, with an increase in replicability: The same

process, narrowly guarded, produces the same predictable outcomes. Becoming highly reliable may be a desirable goal for unsafe or moderately safe operations (Amalberti, 2001). The guaranteed production of standard outcomes through consistent component performance is a way to reduce failure probability in those operations, and it is often expressed as a drive to eliminate errors and technical breakdowns.

In moderately safe systems, such as chemical industries, driving or chartered flights, approaches based on reliability can still generate significant safety returns (Amalberti, 2001). Regulations and safety procedures have a way of converging practice onto a common basis of proven performance. Collecting stories about negative near-miss events (errors, incidents) has the benefit in that the same encounters with risk show up in real accidents that happen to that system. There is, in other words, an overlap between the ingredients of incidents and the ingredients of accidents: recombining incident narratives has predictive (and potentially preventive) value. Finally, developing error-resistant and error-tolerant designs helps prevent errors from becoming incidents or accidents.

The monitoring of performance through operational safety audits, error counting, flight data collection, and incident tabulations has become institutionalized and in many cases required by legislation or regulation. The latest incarnation, an integrative effort to make both safety management and its inspection more streamlined with other organizational processes, is known as the Safety Management System (SMS), which is now demanded in most Western countries by regulators. Safety management systems typically encompass a process for identifying hazards to aviation safety and for evaluating and managing the associated risks, a process for ensuring that personnel are trained and competent to perform their duties and a process for the internal reporting and analyzing of hazards, incidents and accidents and for taking corrective actions to prevent their recurrence. The SMS is also about itself; about the bureaucratic accountability it both represents and spawns. Regulators typically demand that an SMS contains considerable documentation containing all safety management system processes and a process for making personnel aware of their responsibilities with respect to them. Quality assurance and safety management within the airline industry are often mentioned in the same sentence or used under one department heading. The relationship is taken as nonproblematic or even coincident. Quality assurance is seen as a fundamental activity in risk management. Good quality management will help ensure safety. This idea, together with the growing implementation of SMS, may indeed have helped aviation attain even stronger safety records than before, as SMSs help focus decision makers' attention on risk management and safety aspects of both organizational and technological change, forcing an active consideration and documentation of how that risk should be managed.

One possible downside is that pure quality assurance programs (or reliability in the original engineering sense) contain decomposition assumptions that may not really be applicable to systems that are overall as complex as aviation. For example, it suggests that each component or subsystem (layer of defense) operates reasonably independently, so that the results of a safety analysis (e.g., inspection or certification of people or components or subsystems) are not distorted when we start putting the pieces back together again. It also assumes that the principles that govern the assembly of the entire system from its constituent subsystems or components is straightforward. And that the interactions, if any, between the subsystems will be linear: not subject to unanticipated feedback loops or nonlinear interactions.

The assumptions of such a reliability (or quality assurance) approach imply that aviation must continue to strive for systems with high theoretical performance and a high safety potential. A less useful portion of this notion, of course, is the elimination of component breakdowns (e.g., human errors), but it is still a widely pursued goal, sometimes suggesting that the aviation industry today is the custodian of an already safe system that needs protection from unpredictable, erratic components that are its remaining sources of unreliability. This common sense approach, says Amalberti (2001), which indeed may have helped aviation progress to the safety levels of today, is perhaps less applicable to a system that has the levels of complexity and safety already enjoyed today. This is echoed by Vaughan (1996, p. 416):

> …we should be extremely sensitive to the limitations of known remedies. While good management and organizational design may reduce accidents in certain systems, they can never prevent them … technical system failures may be more difficult to avoid than even the most pessimistic among us would have believed. The effect of unacknowledged and invisible social forces on information, interpretation, knowledge, and—ultimately—action, are very difficult to identify and to control.

As progress on safety in aviation has become asymptotic, further optimization of this approach is not likely to generate significant safety returns. In fact, there could be indications that continued linear extensions of a traditional-componential reliability approach could paradoxically help produce a new kind of system accident at the border of almost totally safe practice (Amalberti, 2001, p. 110):

> The safety of these systems becomes asymptotic around a mythical frontier, placed somewhere around $5 \cdot 10^{7}$ risks of disastrous accident per safety unit in the system. As of today, no man-machine system has ever crossed this frontier, in fact, solutions now designed tend to have devious effects when systems border total safety.

It could be necessary to shift from a mechanistic interpretation of complex systems to a systemic one. A machine can be controlled, and it will "fail" or perform less well or run into trouble when one or more of its

components break. In contrast, a living system can be disturbed to any number of degrees. Consequently, its functioning is much less binary, and potentially much more resilient. Such resilience means that failure is not really, or cannot even really be, the result of individual or compound component breakage. Instead, it is related to the ability of the system to adapt to, and absorb variations, changes, disturbances, disruptions and surprises. If it adapts well, absorbs effectively, then even compound component breakages may not hamper chances of survival. United 232 in July 1989 is a case in point. After losing control of the aircraft's control surfaces as a result of a center engine failure that ripped fragments through all three hydraulic lines nearby, the crew figured out how to maneuver the aircraft with differential thrust on two remaining engines. They managed to put the crippled DC-10 down at Sioux City, saving 185 lives out of 293.

Simplicity and complexity

Simple things can generate very complex outcomes that could not be anticipated by just looking at the parts themselves. Small changes in the initial state of a complex system can drastically alter the final outcome. The underlying reason for this is that complex systems are dynamically stable, not statically so (like machines): instability emerges not from components, but from concurrence of functions and events in time. The essence of resilience is the ability to continue to adapt as surprises will inevitably occur to challenge the boundaries of system competencies—how a system is poised to adapt (Hollnagel, Woods, & Leveson, 2006; Woods, 2015, 2019).

Practitioners and organizations, as adaptive systems, continually assess and revise their approaches to work in an attempt to remain sensitive to the possibility of failure. Efforts to create safety, in other words, are ongoing. Not being successful is related to limits of the current model of competence, and, in a learning organization, reflects a discovery of those boundaries. Strategies that practitioners and organizations (including regulators and inspectors) maintain for coping with potential pathways to failure can either be strong or resilient (i.e., well-calibrated) or weak and mistaken (i.e., ill-calibrated). Organizations and people can also become overconfident in how well-calibrated their strategies are. High-reliability organizations remain alert for signs that circumstances exist, or are developing, in which that confidence is erroneous or misplaced (Gras, Moricot, Poirot-Delpech, & Scardigli, 1994; Rochlin, 1993). This, after all, can avoid narrow interpretations of risk and stale strategies (e.g., checking quality of components).

Resilience is the system's ability to effectively adjust to hazardous influences, rather than resist or deflect them (Hollnagel, Woods, & Leveson, 2006). The reason for this is that these influences are also ecologically adaptive and help guarantee the system's survival. Engaging crews

from different (lower-wage) countries makes it possible to keep flying even with oil prices at record highs. But effective adjustment to these potentially hazardous influences did not occur at any level in the system in this case. The systems perspective, of living organizations whose stability is dynamically emergent rather than structurally inherent, means that safety is something a system does, not something a system has—in other words, resilience is a verb that addresses the processes that create/sustain future adaptive capacitiy (Hollnagel, Woods, & Leveson, 2006; Woods, 2019). Failures represent breakdowns in adaptations directed at coping with complexity (Woods, 2003). Learning and adaptation as advocated by HRO are ongoing—without it, safety cannot be maintained in a dynamic and changing organizational setting and environment. As HRO research found, this involves multiple rationalities, reflexivity and self-consciousnesses, since the ability to identify situations that had the potential to evolve into real trouble (and separate them from the ones that did not) is in itself part of the safe operation as social construct. Differently positioned actor-groups are learning, and are learning different things at different times—never excluding their own structure or social relations from the discourse in which that learning is embedded (Rochlin, 1999).

Ensuring resilience in high-reliability organizations

HRO's recognition of the value of resilience supports the notion of safety as something that an organization does, not something that an organization has. The switch has been to focus on safety as the presence of capacities that make things go *well* rather than about the absence of negative events (e.g., incidents). The goal of HRO has expanded to describe what makes organizations reliable, robust, and resilient since these three are not the same but are all necessary for safe operations (Woods, 2022). The purpose of this chapter, and this book, is to explore how might we translate some of these research results into applicable guidance for organizations in aviation and elsewhere. How can we keep an organization's belief in its own continued safe operation curious, open-minded, complexly sensitized, inviting of doubt, and ambivalent toward the past? Resilience is one of the areas of research that offers a kind of action agenda of HRO, with some of the following items:

Not taking past success as guarantee of future safety. Does the system see continued operational success as a guarantee of future safety, as an indication that hazards are not present or that countermeasures in place suffice? In their work, HRO researchers found how safe operation in commercial aviation depends in part on frontline operators treating their operational environment not only as inherently risky, but also as actively hostile to those who misestimate that risk (Rochlin, 1993). Confidence in equipment

and training does not take away the need operators see for constant vigilance for signs that a situation is developing in which that confidence is erroneous or misplaced (Rochlin, 1999). Weick (1993) cites the example of Naskapi Indians who use caribou shoulder bones to locate game. They hold the bones over a fire until they crack and then hunt in the directions where the cracks point. This means future decisions about where to hunt are not influenced by past success, so the animal stock is not depleted and game does not get a chance to habituate to the Indians' hunting patterns. Not only are past results not taken as reason for confidence in future ones—*not* doing so actually increases future chances of success.

Distancing through differencing. In this process, organizational members look at other incidents or failures in other organizations or subunits as not relevant to them and their situation (Cook & Woods, 2006). They discard other events because they appear to be dissimilar or distant. But just because the organization or section has different technical problems, different operational settings, different managers, different histories, or can claim to already have addressed a particular safety concern revealed by the event, does not mean that they are immune to the problem. Seemingly divergent events can represent similar underlying patterns in the drift toward hazard.

Fragmented problem solving. It could be interesting to probe to what extent problem-solving activities are disjointed across organizational departments, sections or subcontractors, as discontinuities and internal handovers of tasks increase risk (Patterson, Roth, Woods, Chow, & Gomes, 2004). With information incomplete, disjointed and patchy, nobody may be able to recognize the gradual erosion of safety constraints on the design and operation of the original system (Dekker, 2011; Woods, 2006). HRO researchers have found that the importance of free-flowing information cannot be overestimated. A spontaneous and continuous exchange of information relevant to normal functioning of the system offers a background from which signs of trouble can be spotted by those with the experience to do so (Rochlin, 1999; Weick, 1993). Research done on handovers, which is one coordinative device to avert the fragmentation of problem-solving (Patterson et al., 2004) has identified some of the potential costs of failing to be told, forgetting, or misunderstanding information communicated. These costs, for the incoming crew, include:

- Having an incomplete model of the system's state;
- Being unaware of significant data or events;
- Being unprepared to deal with impacts from previous events;
- Failing to anticipate future events;
- Lacking knowledge that is necessary to perform tasks safely;
- Dropping or reworking activities that are in progress or that the team has agreed to do;
- Creating an unwarranted shift in goals, decisions, priorities, or plans.

The courage to say no. Having a person or function within the system with the authority, credibility and resources to go against common interpretations and decisions about safety and risk (Woods, 2006). A shift in organizational goal trade-offs often proceed gradually as pressure leads to a narrowing focus on some goals while obscuring the trade-off with other goals. This process usually happens when acute goals like production/efficiency take precedence over chronic goals like safety. If uncertain "warning" signs always led organizations to make sacrifices on schedule and efficiency, it would be difficult to meet competitive and stakeholder demands. By contrast, if uncertain "warning" signs are always rationalized away the organization is acting much riskier than it realizes or wishes. Sometimes people need the courage to put chronic goals ahead of acute short term goals. Thus it is necessary for organizations to support people when they have the courage to say "no" (e.g., in procedures, training, feedback on performance) as these moments serve as reminders of chronic concerns even when the organization is under acute pressures that easily can trump the warnings (see Dekker (2007), about how to create a Just Culture). Resilient systems build in this function at meaningful organizational levels, which relates to the next point.

The ability to bring in fresh perspectives. Systems that apply fresh perspectives (e.g., people from another backgrounds, diverse viewpoints) on problem-solving activities seem to be more effective: they generate more hypotheses, cover more contingencies, openly debate rationales for decision making, reveal hidden assumptions. In HRO studies of some organizations, constant rotation of personnel turned out to be valuable in part because it helped introduce fresh viewpoints in an organizationally and hierarchically legitimate fashion (Rochlin, 1999). Crucially important here is also the role of minority viewpoints, those that can be dismissed easily because they represent dissent from a smaller group. Minority viewpoints can be blocked because they deviate from the mainstream interpretation which will be able to generate many reasons the minority view misunderstands current conditions and retards the organizations formal plans. The alternative readings that minority viewpoints represent, however, can offer a fresh angle that reveals aspects of practice that were obscured from the mainstream perspective (Starbuck & Farjoun, 2005). Historically, "whistleblowers" may hail from lower ranks where the amount of knowledge about the extent of the problem is not matched by the authority or resources to do something about it or have the system change course (Vaughan, 1996). Yet in risky judgments we have to defer to those with technical expertise (and have to set up a problem-solving process that engages those practiced at recognizing anomalies in the event).

All of this can serve to *keep a discussion about risk alive* even (or especially) when everything looks safe. One way is to see whether activities associated with recalibrating models of safety and risk are going on at all. Encouraging

this behavior typically creates forums where stakeholders can discuss risks even when there is no evidence of risk present in terms of current safety statistics. As Weick (1993) illustrates, extreme confidence and extreme caution can both paralyze people and organizations because they sponsor a closed-mindedness that either shuns curiosity or deepens uncertainties. But if discussions about risk are going on even in the absence of obvious threats to safety, one could get some confidence that an organization is investing in an analysis, and possibly in a critique and subsequent update, of its models of how it creates safety.

Knowing the gap between work-as-imagined and work-as-practiced. One marker of resilience is the distance between operations as management imagines they go on and how they actually go on. A large distance indicates that organizational leadership may be miscalibrated to the challenges and risks encountered in real operations. Also, they may also miss how safety is actually created as people conduct work, construct discourse and rationality around it, and gather meaning from it.

Monitoring of safety monitoring (or meta-monitoring). In developing their safety strategies and risk countermeasures, organizations should invest in an awareness of the models of risk they believe in and apply. This is important if organizations want to avoid stale coping mechanisms, misplaced confidence in how they regulate or check safety, and if do not want to miss new possible pathways to failure. Such meta-monitoring would obviously represent an interesting new task for regulators in aviation worldwide, but it applies reflexively to themselves, too. An important aspect of engineering a resilient system is constantly testing whether ideas about risk still match with reality; whether the model of operations (and what makes them safe or unsafe) is still up to date—at every level in the operational, managerial and regulatory hierarchy.

High resilience organizations

Over the past 2 decades, high-reliability research and resilience engineering research showed how organizations can manage acute pressures of performance and production in a constantly dynamic balance with chronic concern for safety. Safety is not something that these organizations have, it is something that organizations do. Practitioners and organizations, as adaptive systems, continually assess and revise their work so as to remain sensitive to the possibility of failure. Efforts to create safety are ongoing, but not always successfully so. An organization usually is unable to change its model of itself unless and until overwhelming evidence accumulates that demands revising the model. This is a guarantee that the organization will tend to learn late, that is, revise its model of risk

only after serious events occur. The crux is to notice the information that changes past models of risk and calls into question the effectiveness of previous risk reduction actions, without having to wait for complete clear cut evidence. If revision only occurs when evidence is overwhelming, there is a grave risk of an organization acting too risky and finding out only from near misses, serious incidents, or even actual harm. The practice of revising assessments of risk needs to be continuous.

High-reliability organization research is continually evolving as its language for accommodating and communicating the results evolves all the time. It is already obvious, though, that traditional engineering notions of reliability (that safety can be maintained by keeping system component performance inside acceptable and prespecified bandwidths) have very little to do with what makes organizations highly reliable, or, rather, resilient (Woods, 2022). As progress on safety in aviation has become asymptotic, further optimization of this reliability approach is not likely to generate significant safety returns. Failure in aviation today is not really, or not in any interesting or predictively powerful way, the result of individual or compound component breakage. Instead, it is related to the ability of the industry to effectively adapt to, and absorb variations, changes, disturbances, disruptions, and surprises.

One contributor to the rise of Resilience Engineering was findings from the HRO work described in Sutcliffe & Vogus (2003) and Weick et al. (2008). Resilience Engineering is concerned with assessing organizational risk, that is, the risk that holes in organizational decision making will produce unrecognized drift toward failure boundaries where brittle collapse occurs (Woods, 2015, 2019). While assessing technical hazards is one kind of input into Resilience Engineering, the goal is to monitor organizational pressures produce conflicts that drive local adaptations. For example, Resilience Engineering would monitor for evidence that organizational pressures undermine effective cross checks are being degraded when risky decisions are made, or that economic pressures limit the investment organizations make in practicing how to handle surprising anomalies following design changes (as occurred in the Boeing 737 Max accidents).

Other dimensions of organizational risk include the commitment of the management to balance the acute pressures of production with the chronic pressures of protection (Woods, 2006). Their willingness to invest in safety and to allocate resources to safety improvement in a timely, proactive manner, despite pressures on production and efficiency, are key factors in ensuring a resilient organization. The degree to which the reporting of safety concerns and problems is truly open and encouraged provides another significant source of resilience within the organization. Assessing the organization's response to incidents indicates if there is a

learning culture or a culture of denial. Other dimensions of organizations which could be monitored include:

- *Preparedness/Anticipation*—is the organization proactive in picking up on evidence of developing problems versus only reacting after problems become significant?
- *Opacity/Observability*—does the organization monitor safety boundaries and recognize how close it is to "the edge" in terms of degraded defenses and barriers? To what extent is information about safety concerns widely distributed throughout the organization at all levels versus closely held by a few individuals?
- *Flexibility/Stiffness*—how does the organization adapt to change, disruptions, and opportunities?
- Successful, highly reliable aviation organizations in the future will have become skilled at the three basics of Resilience Engineering:
  1) detecting signs of increasing organizational risk, especially when production pressures are intense or increasing;
  2) having the resources and authority to make extra investments in safety at precisely the times when it appears least affordable;
  3) having a means to recognize when and where to make targeted investments to control rising signs of organizational risk and rebalance the safety and production trade-off.

These mechanisms can help produce an organization that creates foresight about changing risks before failures and harm occur—continuing the original impetus and spirit of high-reliability theory once more.

References

Amalberti, R. (2001). The paradoxes of almost totally safe transportation systems. *Safety Science*, *37*(2–3), 109–126.

Cook, R. I., & Woods, D. D. (2006). Distancing through differencing: An obstacle to organizational learning following accidents. In *Resilience Engineering* (pp. 329–338). CRC Press.

Dekker, S. (2007). *Just culture: balancing safety and accountability*. Ashgate Publishing, Ltd.

Dekker, S. W. A. (2011). *Drift into failure: From hunting broken components to understanding complex systems*. Farnham, UK: Ashgate Publishing Co.

Gras, A., Moricot, C., Poirot-Delpech, S., & Scardigli, V. (1994). *Face à l'automate: Le pilote, le contrôleur et l'ingénieur*. Éditions de la Sorbonne.

Hollnagel, E., Woods, D. D., & Leveson, N. (Eds.). (2006). *Resilience engineering: Concepts and precepts* Ashgate Publishing, Ltd.

La Porte, T. R. (1996). High reliability organizations: Unlikely, demanding and at risk. *Journal of Contingencies & Crisis Management*, *4*(2), 60–71.

LaPorte, T. R., & Consolini, P. M. (1991). Working in practice but not in theory: Theoretical challenges of "high-reliability organizations". *Journal of Public Administration Research and Theory*, *1*(1), 19–48.

Patterson, E. S., Roth, E. M., Woods, D. D., Chow, R., & Gomes, J. O. (2004). Handoff strategies in settings with high consequences for failure: Lessons for health care operations. *International Journal for Quality in Health Care*, *16*(2), 125–132.

Perrow, C. (1984). *Normal accidents: Living with high-risk technologies*. New York: Basic Books.
Roberts, K. H. (1989). New challenges in organizational research: High reliability organizations. *Organization & Environment*, *3*(2), 111–125.
Roberts, K. H. (1990). Some characteristics of one type of high reliability organization. *Organization Science*, *1*(2), 160–176.
Roberts, K. H., & Gargano, G. (1989). Managing a high reliability organization: A case for interdependence. In *Managing complexity in high technology industries: Systems and people*. New York: Oxford University Press.
Rochlin, G. I. (1993). Essential friction: Error-control in organizational behavior. In *The necessity of friction* (pp. 196–232). Physica-Verlag HD.
Rochlin, G. I. (1999). Safe operation as a social construct. *Ergonomics*, *42*(11), 1549–1560.
Rochlin, G. I., LaPorte, T. R., & Roberts, K. H. (1987). The self-designing high reliability organization: Aircraft carrier flight operations at sea. *Naval War College Review*, *40*, 76–90.
Schochlow, V., & Dekker, S. W. A. (2019). The 1980s and onward: Normal accidents and high reliability organizations. In *Foundations of safety science* (pp. 267–304). Routledge.
Starbuck, W. H., & Farjoun, M. (2005). *Organization at the limit: Lessons from the Columbia disaster*. Malden, MA: Blackwell Pub.
Sutcliffe, K. M., & Vogus, T. J. (2003). Organizing for resilience. In K. S. Cameron, J. E. Dutton, & R. E. Quinn (Eds.), *Positive organizational scholarship: Foundations of a new discipline* (p. 94). San Francisco, CA: Berret-Koehler Publishers, Inc.
Vaughan, D. (1996). *The challenger launch decision: Risky technology, culture, and deviance at NASA*. Chicago: University of Chicago Press.
Vaughan, D. (2002). Signals and interpretive work: The role of culture in a theory of practical action. In *Culture in mind: Toward a sociology of culture and cognition* (pp. 28–54). Routledge.
Weick, K. E. (1993). The collapse of sensemaking in organizations: The Mann gulch disaster. *Administrative Science Quarterly*, *38*(4), 628–652.
Weick, K. E., Sutcliffe, K. M., & Obstfeld, D. (2008). Organizing for high reliability: Processes of collective mindfulness. *Crisis Management*, *3*, 81–123.
Woods, D. D. (2003). Discovering how distributed cognitive systems work. *Handbook of Cognitive Task Design*, 37–53.
Woods, D. D. (2006). Essential characteristics of resilience. In E. Hollnagel, D. D. Woods, & N. G. Leveson (Eds.), *Resilience engineering: Concepts and precepts* (pp. 21–34). Aldershot: Ashgate Publishing Co.
Woods, D. D. (2015). Four concepts of resilience and the implications for resilience engineering. *Reliability Engineering and Systems Safety*, *141*, 5–9. https://doi.org/10.1016/j.ress.2015.03.018.
Woods, D. D. (2019). Essentials of resilience, revisited. In M. Ruth, & S. G. Reisemann (Eds.), *Handbook on resilience of socio-technical systems* (pp. 52–65). Edward Elgar Publishing.
Woods, D. D. (2022). *Why do reliable systems fail? Adobe Research Summit, May*. https://www.youtube.com/watch?v=fbwDnpuys7w&t=16s. (Accessed 6 September 2022).