

Prologue: Resilience Engineering Concepts

David D. Woods
Erik Hollnagel

Hindsight and Safety

Efforts to improve the safety of systems have often – some might say always – been dominated by hindsight. This is so both in research and in practice, perhaps more surprising in the former than in the latter. The practical concern for safety is usually driven by events that have happened, either in one's own company or in the industry as such. There is a natural motivation to prevent such events from happening again, in concrete cases because they may incur severe losses – of equipment and/or of life – and in general cases because they may lead to new demands for safety from regulatory bodies, such as national and international administrations and agencies. New demands are invariably seen as translating into increased costs for companies and are for that reason alone undesirable. (This is, however, not an inevitable consequence, especially if the company takes a longer time perspective. Indeed, for some businesses it makes sense to invest proactively in safety, although cases of that are uncommon. The reason for this is that sacrificing decisions usually are considered over a short time horizon, in terms of months rather than years or in terms of years rather than decades.)

In the case of research, i.e., activities that take place at academic institutions rather than in industries and are driven by intellectual rather than economic motives, the effects of hindsight ought to be less marked. Research should by its very nature be looking to problems that go beyond the immediate practical needs, and hence address issues that are of a more principal nature. Yet even research – or perhaps one

should say: researchers – are prone to the effects of hindsight, as pointed out by Fischhoff (1975). It is practically a characteristic of human nature – and an inescapable one at that – to try to make sense of what has happened, to try to make the perceived world comprehensible. We are consequently constrained to look at the future in the light of the past. In this way our experience or understanding of what has happened inevitably colours our anticipation and preparation for what could go wrong and thereby holds back the requisite imagination that is so essential for safety (Adamski & Westrum, 2003). Approaches to safety and risk prediction furthermore develop in an incremental manner, i.e., the tried and trusted approaches are only changed when they fail and then usually by adding one more factor or element to account for the unexplained variability. Examples are easy to find such as ‘human error’, ‘organisational failures’, ‘safety culture’, ‘complacency’, etc. The general principle seems to be that we add or change just enough to be able to explain that which defies the established framework of explanations. In contrast, resilience engineering tries to take a major step forward, not by adding one more concept to the existing vocabulary, but by proposing a completely new vocabulary, and therefore also a completely new way of thinking about safety. With the risk of appearing overly pretentious, it may be compared to a paradigm shift in the Kuhnian sense (Kuhn, 1970).

When research escapes from hindsight and from trying merely to explain what *has* happened, studies reveal the sources of resilience that usually allow people to produce success when failure threatens. Methods to understand the basis for technical work shows how workers are struggling to anticipate paths that may lead to failure, actively creating and sustaining failure-sensitive strategies, and working to maintain margins in the face of pressures to do more and to do it faster (Woods & Cook, 2002). In other words, doing things safely always has been and always will be part of operational practices – on the individual as well as the organisational level. It is, indeed, almost a biological law that organisms or systems (including organisations) that spend all efforts at the task at hand and thereby neglect to look out for the unexpected, run a high risk of being obliterated, of meeting a speedy and unpleasant demise. (To realise that, you only need to look at how wild birds strike a balance between head-down and head-up time when eating.) People in their different roles within an organisation are aware of potential paths to failure and therefore develop failure-

sensitive strategies to forestall these possibilities. Failing to do that brings them into a reactive mode, a condition of constant fire-fighting. But fires, whether real or metaphorical, can only be effectively quelled if the fire-fighters are proactive and able to make the necessary sacrifices (McLennan et al., 2005).

Against this background, failures occur when multiple contributors – each necessary but only jointly sufficient – combine. Work processes or people do not choose failure, but the likelihood of failures grow when production pressures do not allow sufficient time – and effort – to develop and maintain the precautions that normally keep failure at bay. Prime among these precautions is to check all necessary conditions and to take nothing important for granted. Being thorough as well as efficient is the hallmark of success. Being efficient without being thorough may gradually or abruptly create conditions where even small variations can have serious consequences. Being thorough without being efficient rarely lasts long, as organisations are pressured to meet new demands on resources. 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 gaps, hazards, trade-offs, and multiple goals (Cook et al., 2000).

The thesis that leaps out from these results is that failure, as individual failure or performance failure on the system level, represents the temporary inability to cope effectively with complexity. Success belongs to organisations, groups and individuals who are resilient in the sense that they recognise, adapt to and absorb variations, changes, disturbances, disruptions, and surprises – especially disruptions that fall outside of the set of disturbances the system is designed to handle (Rasmussen, 1990; Rochlin, 1999; Weick et al., 1999; Sutcliffe & Vogus, 2003).

From Reactive to Proactive Safety

This book marks the maturation of a new approach to safety management. In a world of finite resources, of irreducible uncertainty, and of multiple conflicting goals, safety is created through proactive resilient processes rather than through reactive barriers and defences. The chapters in this book explore different facets of resilience as the

ability of systems to anticipate and adapt to the potential for surprise and failure.

Until recently, the dominant safety paradigm was based on searching for ways in which limited or erratic human performance could degrade an otherwise well designed and 'safe system'. Techniques from many areas such as reliability engineering and management theory were used to develop 'demonstrably safe' systems. The assumption seemed to be that safety, once established, could be maintained by requiring that human performance stayed within prescribed boundaries or norms. Since 'safe' systems needed to include mechanisms that guarded against people as unreliable components, understanding how human performance could stray outside these boundaries became important.

According to this paradigm, 'error' was something that could be categorised and counted. This led to numerous proposals for taxonomies, estimation procedures, and ways to provide the much needed data for error tabulation and extrapolation. Studies of human limits became important to guide the creation of remedial or prosthetic systems that would make up for the deficiencies of people. Since humans, as unreliable and limited system components, were assumed to degrade what would otherwise be flawless system performance, this paradigm often prescribed automation as a means to safeguard the system from the people in it. In other words, in the 'error counting' paradigm, work on safety comprised protecting the system from unreliable, erratic, and limited human components (or, more clearly, protecting the people at the blunt end – in their roles as managers, regulators and consumers of systems – from unreliable 'other' people at the sharp end – who operate and maintain those systems).

When researchers in the early 1980s began to re-examine human error and collect data on how complex systems had failed, it soon became apparent that people actually provided a positive contribution to safety through their ability to adapt to changes, gaps in system design, and unplanned for situations. Hollnagel (1983), for instance, argued for the need of a theory of action, including an account of performance variability, rather than a theory of 'error', while Rasmussen (1983) noted that 'the operator's role is to make up for holes in designers' work.' Many studies of how complex systems succeeded and sometimes failed found that the formal descriptions of work embodied in policies, regulations, procedures, and automation were incomplete as

models of expertise and success. Analyses of the gap between formal work prescriptions and actual work practices revealed how people in their various roles throughout systems always struggled to anticipate paths toward failure, to create and sustain failure-sensitive strategies, and to maintain margins in the face of pressures to increase efficiency (e.g., Cook et al, 2000). Overall, analysis of such ‘second stories’ taught us that failures represented breakdowns in adaptations directed at coping with complexity while success was usually obtained as people learned and adapted to create safety in a world fraught with hazards, trade-offs, and multiple goals (Rasmussen, 1997). In summary, these studies revealed:

- How workers and organisations continually revise their approach to work in an effort to remain sensitive to the possibility for failure.
- How distant observers of work, and the workers themselves, are only partially aware of the current potential for failure.
- How ‘improvements’ and changes create new paths to failure and new demands on workers, despite or because of new capabilities.
- How the strategies for coping with these potential paths can be either strong and resilient or weak and mistaken.
- How missing the side effects of change is the most common form of failure for organisations and individuals.
- How a culture of safety depends on remaining dynamically engaged in new assessments and avoiding stale, narrow, or static representations of the changing paths (revising or reframing the understanding of paths toward failure over time).
- How overconfident people can be that they have already anticipated the types and mechanisms of failure, and that the strategies they have devised are effective and will remain so.
- How continual effort after success in a world of changing pressures and hazards is fundamental to creating safety.

In the final analysis, safety is not a commodity that can be tabulated. It is rather a chronic value ‘under our feet’ that infuses all aspects of practice. Safety is, in the words of Karl Weick, a dynamic non-event. Progress on safety therefore ultimately depends on providing workers and managers with information about changing vulnerabilities and the ability to develop new means for meeting these.

Resilience

Resilience engineering is a paradigm for safety management that focuses on how to help people cope with complexity under pressure to achieve success. It strongly contrasts with what is typical today – a paradigm of tabulating error as if it were a thing, followed by interventions to reduce this count. A resilient organisation treats safety as a core value, not a commodity that can be counted. Indeed, safety shows itself only by the events that do not happen! Rather than view past success as a reason to ramp down investments, such organisations continue to invest in anticipating the changing potential for failure because they appreciate that their knowledge of the gaps is imperfect and that their environment constantly changes. One measure of resilience is therefore the ability to create foresight – to anticipate the changing shape of risk, before failure and harm occurs (Woods, 2005a).

The initial steps in developing a practice of Resilience Engineering have focused on methods and tools:

- to analyse, measure and monitor the resilience of organisations in their operating environment.
- to improve an organisation's resilience vis-à-vis the environment.
- to model and predict the short- and long-term effects of change and line management decisions on resilience and therefore on risk.

This book charts the efforts being made by researchers, practitioners and safety managers to enhance resilience by looking for ways to understand the changing vulnerabilities and pathways to failure. These efforts begin with studies of how people cope with complexity – usually successfully. Analyses of successes, incidents, and breakdowns reveal the normal sources of resilience that allow systems to produce success when failure threatens. These events and other measures indicate the level and kinds of brittleness/resilience the system in question exhibits. Such indicators will allow organisations to develop the mechanisms to create foresight, to recognise, anticipate, and defend against paths to failure that arise as organisations and technology change.