

Chapter 2

Essential Characteristics of Resilience

David D. Woods

Avoiding the Error of the Third Kind

When one uses the label ‘resilience,’ the first reaction is to think of resilience as if it were adaptability, i.e., as the ability to absorb or adapt to disturbance, disruption and change. But all systems adapt (though sometimes these processes can be quite slow and difficult to discern) so resilience cannot simply be the adaptive capacity of a system. I want to reserve resilience to refer to the broader capability – how well can a system handle disruptions and variations that fall outside of the base mechanisms/model for being adaptive as defined in that system.

This depends on a distinction between understanding how a system is competent at designed-for-uncertainties, which defines a ‘textbook’ performance envelope and how a system recognizes when situations challenge or fall outside that envelope – unanticipated variability or perturbations (see parallel analyses in Woods et al., 1990 and Carlson & Doyle, 2000; Csete & Doyle, 2002). Most discussions of definitions of ‘robustness’ in adaptive systems debate whether resilience refers to first or second order adaptability (Jen, 2003). In the end, the debates tend to settle on emphasizing the system’s ability to handle events that fall outside its design envelope and debate what is a design envelope, what events challenge or fall outside that envelope, and how does a system see what it has failed to build into its design (e.g., see url: <http://discuss.santafe.edu/robustness/>)

The area of textbook competence is in effect a model of variability/uncertainty and a model of how the strategies/plans

/countermeasures in play handle these, mostly successfully. Unanticipated perturbations arise (a) because the model implicit and explicit in the competence envelope is incomplete, limited or wrong and (b) because the environment changes so that new demands, pressures, and vulnerabilities arise that undermine the effectiveness of the competence measures in play.

Resilience then concerns the ability to recognize and adapt to handle unanticipated perturbations that call into question the model of competence, and demand a shift of processes, strategies and coordination. When evidence of holes in the organization's model builds up, the risk is what Ian Mitroff called many years ago, the error of the third kind, or solving the wrong problem (Mitroff, 1974). This is a kind of under-adaptation failure where people persist in applying textbook plans and activities in the face of evidence of changing circumstances that demand a qualitative shift in assessment, priorities, or response strategy.

This means resilience is concerned with monitoring the boundary conditions of the current model for competence (how strategies are matched to demands) and adjusting or expanding that model to better accommodate changing demands. The focus is on assessing the organization's adaptive capacity relative to challenges to that capacity – what sustains or erodes the organization's adaptive capacities? Is it degrading or lower than the changing demands of its environment? What dynamics challenge or go beyond the boundaries of the competence envelope? Is the organization as well adapted as it thinks it is? Note that boundaries are properties of the model that defines the textbook competence envelope relative to the uncertainties and perturbations it is designed for (Rasmussen, 1990a). Hence, resilience engineering devotes effort to make observable the organization's model of how it creates safety, in order to see when the model is in need of revision.

To do this, Resilience Engineering must monitor organizational decision-making to assess the risk that the organization is operating nearer to safety boundaries than it realizes (Woods, 2005a). Monitoring resilience should lead to interventions to manage and adjust the adaptive capacity as the system faces new forms of variation and challenges.

Monitoring and managing resilience, or its absence, brittleness, is concerned with understanding how the system adapts and to what kinds of disturbances in the environment, including properties such as:

- buffering capacity: the size or kinds of disruptions the system can absorb or adapt to without a fundamental breakdown in performance or in the system's structure;
- flexibility versus stiffness: the system's ability to restructure itself in response to external changes or pressures;
- margin: how closely or how precarious the system is currently operating relative to one or another kind of performance boundary;
- tolerance: how a system behaves near a boundary – whether the system gracefully degrades as stress/pressure increase or collapses quickly when pressure exceeds adaptive capacity.

In addition, cross-scale interactions are critical, as the resilience of a system defined at one scale depends on influences from scales above and below:

- Downward, resilience is affected by how organizational context creates or facilitates resolution of pressures/goal conflicts/dilemmas, for example, mismanaging goal conflicts or poor automation design can create authority-responsibility double binds for operational personnel (Woods et al., 1994; Woods, 2005b).
- Upward, resilience is affected by how adaptations by local actors in the form of workarounds or innovative tactics reverberate and influence more strategic goals and interactions (e.g., workload bottlenecks at the operational scale can lead to practitioner workarounds that make management's attempts to command compliance with broad standards unworkable; Cook et al., 2000).

As illustrated in the cases of resilience or brittleness described or referred to in this book, all systems have some degree of resilience and sources for resilience. Even cases with negative outcomes, when seen as breakdowns in adaptation, reveal the complicating dynamics that stress the textbook envelope and the often hidden sources of resilience used to cope with these complexities.

Accidents have been noted by many analysts as ‘fundamentally surprising’ events because they call into question the organization’s model of the risks they face and the effectiveness of the countermeasure deployed (Lanir, 1986; Woods et al., 1994, chapter 5; Rochlin, 1999; Woods, 2005b). In other words, the organization is unable to recognize or interpret evidence of new vulnerabilities or ineffective countermeasures until a visible accident occurs. At this stage the organization can engage in fundamental learning but this window of opportunity comes at a high price and is fragile given the consequences of the harm and losses. The shift demanded following an accident is a reframing process. In reframing one notices initial signs that call into question ongoing models, plans and routines, and begins processes of inquiry to test if revision is warranted (Klein et al., 2005). Resilience Engineering aims to provide support for the cognitive processes of reframing an organization’s model of how safety is created before accidents occur by developing measures and indicators of contributors to resilience such as the properties of buffers, flexibility, precariousness, and tolerance and patterns of interactions across scales such as responsibility-authority double binds.

Monitoring resilience is monitoring for the changing boundary conditions of the textbook competence envelope – how a system is competent at handling designed-for-uncertainties – to recognize forms of unanticipated perturbations – dynamics that challenge or go beyond the envelope. This is a kind of broadening check that identifies when the organization needs to learn and change. Resilience engineering needs to identify the classes of dynamics that undermine resilience and result in organizations that act riskier than they realize. This chapter focuses on dynamics related to safety-production goal conflicts.

Coping with Pressure to be Faster, Better, Cheaper

Consider recent NASA experience, in particular, the consequences of NASA’s adoption of a policy called ‘faster, better, cheaper’ (FBC). Several years later a series of mishaps in space science missions rocked the organization and called into question that policy. In a remarkable ‘organizational accident’ report, an independent team investigated the organizational factors that spawned the set of mishaps (Spear, 2000).

The investigation realized that FBC was not a policy choice, but the acknowledgement that the organization was under fundamental

pressure from stakeholders. The report and the follow-up, but short-lived, 'Design for Safety' program noted that NASA had to cope with a changing environment with increasing performance demands combined with reduced resources: drive down the cost of launches, meet shorter, more aggressive mission schedules, do work in a new organizational structure that required people to shift roles and coordinate with new partners, eroding levels of personnel experience and skills. Plus, all of these changes were occurring against a backdrop of heightened public and congressional interest that threatened the viability of the space program. The MCO investigation board concluded: NASA, which had a history of 'successfully carrying out some of the most challenging and complex engineering tasks ever faced by this nation,' was being asked to 'sustain this level of success while continually cutting costs, personnel and development time ... these demands have stressed the system to the limit' due to 'insufficient time to reflect on unintended consequences of day-to-day decisions, insufficient time and workforce available to provide the levels of checks and balances normally found, breakdowns in inter-group communications, too much emphasis on cost and schedule reduction.' The MCO Board diagnosed the mishaps as indicators of an increasingly brittle system as production pressure eroded sources of resilience and led to decisions that were riskier than anyone wanted or realized. Given this diagnosis, the Board went on to re-conceptualize the issue as how to provide tools for proactively monitoring and managing project risk throughout a project life-cycle and how to use these tools to balance safety with the pressure to be faster, better, cheaper.

The experience of NASA under FBC is an example of the law of stretched systems: every system is stretched to operate at its capacity; as soon as there is some improvement, for example in the form of new technology, it will be exploited to achieve a new intensity and tempo of activity (Woods, 2003). Under pressure from performance and efficiency demands (FBC pressure), advances are consumed to ask operational personnel 'to do more, do it faster or do it in more complex ways', as the Mars Climate Orbiter Mishap Investigation Board report determined. With or without cheerleading from prestigious groups, pressures to be 'faster, better, cheaper' increase. Furthermore, pressures to be 'faster, better, cheaper' introduce changes, some of which are new capabilities (the term does include 'better'), and these changes modify the vulnerabilities or paths toward failure. How conflicts and trade-offs

like these are recognized and handled in the context of vectors of change is an important aspect of managing resilience.

Balancing Acute and Chronic Goals

Problems in the US healthcare delivery system provide another informative case where faster, better, cheaper pressures conflict with safety and other chronic goals. The Institute of Medicine in a calculated strategy to guide national improvements in health care delivery conducted a series of assessments. One of these, *Crossing the Quality Chasm: A New Health System for the 21st Century* (IOM, 2001), stated six goals needed to be achieved simultaneously: the national health care system should be – Safe, Effective, Patient-centered, Timely, Efficient, Equitable.¹ Each goal is worthy and generates thunderous agreement. The next step seems quite direct and obvious – how to identify and implement quick steps to advance each goal (the classic search for so-called ‘low hanging fruit’). But as in the NASA case, this set of goals is not a new policy direction but rather an acknowledgement of demanding pressures already operating on health care practitioners and organizations. Even more difficult, the six goals represent a set of interacting and often conflicting pressures so that in adapting to reach

¹ The IOM states the quality goals as –
‘Health Care Should Be:

- Safe – avoiding injuries to patients from the care that is intended to help them.
- Effective – providing services based on scientific knowledge to all who could benefit and refraining from providing services to those not likely to benefit (avoiding underuse and overuse, respectively).
- Patient-centered – providing care that is respectful of and responsive to individual patient preferences, needs, and values and ensuring that patient values guide all clinical decisions.
- Timely – reducing waits and sometimes harmful delays for both those who receive and those who give care.
- Efficient – avoiding waste, including waste of equipment, supplies, ideas, and energy.
- Equitable – providing care that does not vary in quality because of personal characteristics such as gender, ethnicity, geographic location, and socioeconomic status.’

for one of these goals it is very easy to undermine or squeeze others. To improve on all simultaneously is quite tricky.

As I have worked on safety in health care, I hear many highly placed voices for change express a basic belief that these six goals can be synergistic. Their agenda is to energize a search for and adoption of specific mechanisms that simultaneously advance multiple goals within the six and that do not conflict with others – ‘silver bullets’. For example, much of the patient safety discussion in US health care continues to be a search for specific mechanisms that appear to simultaneously save money and reduce injuries as a result of care. Similarly, NASA senior leaders thought that including ‘better’ along with faster and cheaper meant that techniques were available to achieve progress on being faster, better, and cheaper together (for almost comic rationalizations of ‘faster, better, cheaper’ following the series of Mars science mission mishaps and an attempt to protect the reputation of the NASA administrator at the time, see Spear, 2000). The IOM and NASA senior management believed that quality improvements began with the search for these ‘silver bullet’ mechanisms (sometimes called ‘best practices’ in health care). Once such practices are identified, the question becomes how to get practitioners and organizations to adopt these practices. Other fields can help provide the means to develop and document new best practices by describing successes from other industries (health care frequently uses aviation and space efforts to justify similar programs in health care organizations). The IOM in particular has had a public strategy to generate this set of silver bullet practices and accompanying justifications (like creating a quality catalog) and then pressure health care delivery decision makers to adopt them all in the firm belief that, as a result, all six goals will be advanced simultaneously and all stakeholders and participants will benefit (one example is computerized physician order entry).

However, the findings of the Columbia accident investigation board (CAIB) report should reveal to all that the silver bullet strategy is a mirage. The heart of the matter is not silver bullets that eliminate conflicts across goals, but developing new mechanisms that balance the inherent tensions and trade-offs across these goals (Woods et al., 1994). The general trade-off occurs between the family of acute goals – timely, efficient, effective (or after NASA’s policy, the Faster, Better, Cheaper or FBC goals) and the family of chronic goals, for the health care case consisting of safety, patient-centeredness, and equitable access.

The tension between acute production goals and chronic safety risks is seen dramatically in the Columbia accident which the investigation board found was the result of pressure on acute goals eroding attention, energy and investments on chronic goals related to controlling safety risks (Gehman, 2003). Hollnagel (2004, p. 160) compactly captured the tension between the two sets of goals with the comment that:

If anything is unreasonable, it is the requirement to be both efficient and thorough at the same time – or rather to be thorough when with hindsight it was wrong to be efficient.

The FBC goal set is acute in the sense that they happen in the short term and can be assessed through pointed data collection that aggregates element counts (shorter hospital stays, delay times). Note that ‘better’ is in this set, though better in this family means increasing capabilities in a focused or narrow way, e.g., cardiac patients are treated more consistently with a standard protocol. The development of new therapies and diagnostic capabilities belongs in the acute sense of ‘better.’

Safety, access, patient-centeredness are chronic goals in the sense that they are system properties that emerge from the interaction of elements in the system and play out over longer time frames. For example, safety is an emergent system property, arising in the interactions across components, subsystems, software, organizations, and human behavior.

By focusing on the tensions across the two sets, we can better see the current situation in health care. It seems to be lurching from crisis to crisis as efforts to improve or respond in one area are accompanied by new tensions at the intersections of other goals (or the tensions are there all along and the visible crisis point shifts as stakeholders and the press shift their attention to different manifestations of the underlying conflicts). The tensions and trade-offs are seen when improvements or investments in one area contribute to greater squeezes in another area. The conflicts are stirred by the changing background of capabilities and economic pressure. The shifting points of crisis can be seen first in 1995-6 as dramatic well publicized deaths due to care helped create the patient safety crisis (ultimately documented in Kohn et al., 1999). The patient safety movement was energized by patients feeling vulnerable as

health care changed to meet cost control pressures. Today attention has shifted to an access crisis as malpractice rates and prescription drug costs undermine patients' access to physicians in high risk specialties and challenge seniors' ability to balance medication costs with limited personal budgets.

Dynamic Balancing Acts

If the tension view is correct, then progress revolves around how to dynamically balance the potential trade-offs so that all six goals can advance (as opposed to the current situation where improvements or investments in one area create greater squeezes in another area). It is important to remember that trade-offs are defined by two parameters, one that captures discrimination power or how well one can make the underlying judgement, and a second that defines where to place a criterion for making a decision or taking an action along the trade-off curve, criterion placement or movement. The parameters of a trade-off cannot be estimated by a single case, but require integration over behavior in sets of cases and over time.

One aspect of the difficulty of goal conflicts is that the default or typical ways to advance the acute goals often make it harder to achieve chronic goals simultaneously. For example, increasing therapeutic capabilities can easily appear as new silos of care that do not redress and can even exacerbate fragmentation of care (undermining the patient-centeredness goal). To advance all of the goals, ironically, the chronic set of goals of patient centered, safety and access must be put first, with secondary concern for efficient and timely methods. To do otherwise will fall prey to the natural tendency to value the more immediate and direct consequences (which, by the way, are easier to measure) of the acute set over the chronic and produce an unintentional sacrifice on the chronic set. Effective balance seems to arise when organizations shift from seeing safety as one of a set of goals to be measured (is it going up or down?) to considering safety as a basic value. The point is that for chronic goals to be given enough weight in the interaction with acute goals, the chronic needs to be approached much more like establishing a core cultural value.

For example, valuing the chronic set in health care puts patient centeredness first with its fellow travelers safety and access. The central

issue under patient centeredness is emergent continuity of care, as the patient makes different encounters with the health care system and as disease processes develop over time. The opposite of continuity is fragmentation. Many of the tensions across goals exacerbate fragmentation, e.g., ironically, new capabilities on specific aspects of health care can lead to more specialization and more silos of care. Placing priority on continuity of care vs. fragmentation focuses attention (a) on health care issues related to chronic diseases which require continuity and which are inherently difficult in a fragmented system of care and (b) on cognitive system issues which address coordination over time, over practitioners, over organizations, and over specialized knowledge sources. Consider the different ways new technology can have an effect on patient care. Depending on how computer systems are built and adapted over time, more computerization can lead to less contact with patients and more contact with the image of the patient in the database. This is a likely outcome when FBC pressure leads acute goals to dominate chronic ones (the benefits of the advance in information technology will tend to be consumed to meet pressures for productivity or efficiency). When a chronic goal such as continuity of care, functions as the leading value, the emphasis shifts to finding uses of computer capabilities that increase attention and tailoring of general standards to a specific patient over time (increasing the effective continuity) and only then developing these capabilities to meet cost considerations.

The tension diagnosis is part of the more general diagnosis that past success has led to increasingly complex systems with new forms of problems and failure risks. The basic issue for organizational design is how large-scale systems can cope with complexity, especially the pace of change and coupling across parts that accompany the methods that advance the acute goals. To miss the complexity diagnosis will make otherwise well-intentioned efforts fail as each attempt to advance goals simultaneously through silver bullets will rebound as new crises where goal trade-offs create new dissatisfactions and tensions.

Sacrifice Judgements

To illustrate a safety culture, leaders tell stories about an individual making tough decisions when goals conflict. The stories always have the same basic form even though the details may come from a personal

experience or from re-telling of a story gathered from another domain with a high reputation for safety (e.g., health care leaders often use aerospace stories):

Someone noticed there might be a problem developing, but the evidence is subtle or ambiguous. This person has the courage to speak up and stop the production process underway. After the aircraft gets back on the ground or after the system is dismantled or after the hint is chased down with additional data, then all discover the courageous voice was correct. There was a problem that would otherwise have been missed and to have continued would have resulted in failure, losses, and injuries. The story closes with an image of accolades for the courageous voice.

When the speaker finishes the story, the audience sighs with appreciation – that was an admirable voice and it illustrates how a great organization encourages people to speak up about potential safety problems. You can almost see people in the audience thinking, ‘I wish my organization had a culture that helped people act this way.’

But this common story line has the wrong ending. It is a quite different ending that provides the true test for a high resilience organization.

When they go look, after the landing or after dismantling the device or after the extra tests were run, everything turns out to be OK. The evidence of a problem isn’t there or may be ambiguous; production apparently did not need to be stopped. Now, how does the organization’s management react? How do the courageous voice’s peers react?

For there to be high resilience, the organization has to recognize the voice as courageous and valuable even though the result was apparently an unnecessary sacrifice on production and efficiency goals. Otherwise, people balancing multiple goals will tend to act riskier than we want them to, or riskier than they themselves really want to.

These contrasting story lines illustrate the difficulties of balancing acute goals with chronic ones. Given a backdrop of schedule pressure, how should an organization react to potential ‘warning’ signs and seek to handle the issues the signs point to? If organizations never sacrifice production pressure to follow up warning signs, they are acting much too risky. On the other hand, if uncertain ‘warning’ signs always lead to sacrifices on acute goals, can the organization operate within reasonable parameters or stakeholder demands? It is easy for organizations that are working hard to advance the acute goal set to see such warning signs as risking inefficiencies or as low probability of concern as they point to a record of apparent success and improvement. Ironically, these same signs after-the-fact of an accident appear to all as clear cut undeniable warning signs of imminent dangers.

To proactively manage risk prior to outcome requires ways to know when to relax the pressure on throughput and efficiency goals, i.e., making a sacrifice judgement. Resilience engineering needs to provide organizations with help on how to decide when to relax production pressure to reduce risk (Woods, 2000). I refer to these trade-off decisions as sacrifice judgements because acute production or efficiency related goals are temporarily sacrificed, or the pressure to achieve these goals is relaxed, in order to reduce the risks of approaching too near safety boundaries. Sacrifice judgements occur in many settings: when to convert from laparoscopic surgery to an open procedure (Dominguez et al., 2004 and the discussion in Cook et al., 1998), when to break off an approach to an airport during weather that increases the risks of wind shear, and when to have a local slowdown in production operations to avoid risks as complications build up.

New research is needed to understand this judgement process in individuals and in organizations. Previous research on such decisions (e.g., production/safety trade-off decisions in laparoscopic surgery) indicates that the decision to value production over safety is implicit and unrecognized. The result is that individuals and organizations act much riskier than they would ever desire. A sacrifice judgement is especially difficult because the hindsight view will indicate that the sacrifice or relaxation may have been unnecessary since ‘nothing happened.’ This means that it is important to assess how peers and superiors react to such decisions.

The goal is to develop explicit guidance on how to help people make the relaxation/sacrifice judgement under uncertainty, to maintain

a desired level of risk acceptance/risk averseness, and to recognize changing levels of risk acceptance/risk averseness. For example, what indicators reveal a safety/production trade-off sliding out of balance as pressure rises to achieve acute production and efficiency goals? Ironically, it is these very times of higher organizational tempo and focus on acute goals that require extra investments in sources of resilience to keep production/safety trade-offs in balance – valuing thoroughness despite the potential for sacrifices on efficiency required to meet stakeholder demands.

Note how the recommendation to aid sacrifice judgements is a specialization of general methods for aiding any system confronting a trade-off: (a) improve the discrimination power of the system confronting the trade-off, and (b) help the system dynamically match its placement of a decision criterion with the assessment of changing risk and uncertainty.

Resilience Engineering should provide the means for dynamically adjusting the balance across the sets of acute and chronic goals. The dilemma of production pressure/safety trade-offs is that we need to pay the most attention to, and devote scarce resources to, potential future safety risks when they are least affordable due to increasing pressures to produce or economize. As a result, organizations unknowingly act riskier than they would normally accept. The first step is tools to monitor the boundary between competence at designed-for-uncertainties and unanticipated perturbations that challenge or fall outside that envelope. Recognizing signs of unanticipated perturbations consuming or stretching the sources of resilience in the system can lead actions to re-charge a system's resilience. How can we increase, maintain, or re-establish resilience when buffers are being depleted, margins are precarious, processes become stiff, and squeezes become tighter?

Acknowledgements

This work was supported in part by grant NNA04CK45A from NASA Ames Research Center to develop resilience engineering concepts for managing organizational risk. The ideas presented benefited from discussions in the NASA's Design for Safety workshop and Workshop

on organizational risk. Discussions with John Wreathall helped develop the model of trade-offs across acute and chronic goals.