

Chapter 12

Stress-Strain Plots as a Basis for Assessing System Resilience

David D. Woods & John Wreathall

Introduction

Ways to characterize and measure an organization's resilience can be based on an analogy from the world of materials engineering: that of the relationship between stress – the varying loads placed on a mechanical structure, and the resulting strain – how the structure stretches in response. This aspect of material science is related to the basic finding in Cognitive Systems Engineering (CSE) that demand factors are critical (Woods, 1988; Rasmussen et al., 1993). Thus, to characterize a cognitive system of people and machines one should examine how that joint system responds to different demands on work, in other words, plot how a system stretches in response to changes in demands. For example, Woods & Patterson (2000; cf. Woods & Hollnagel, 2006, chapter 9) used this idea to propose that one should evaluate and predict system performance by relating how demands increase and cascade to the ability of the joint system to bring more knowledge and expertise to bear. In effect, they suggested a joint system is characterized by a mapping between how demands increase relative to how the system stretches to accommodate the increasing demands (tempo, cascade of effects, and the potential for bottlenecks).

We noticed a similarity between this aspect of CSE and discussions about how organizations adapt and about how to engineer resilience into organizations (Weick et al., 1999; Sutcliffe and Vogus, 2003; Hollnagel et al., 2006; Hollnagel and Rigaud, 2006). Descriptions that contrasted cases of resilience and brittleness often use similar language to express how an organization stretches as demands increase (Cook and Nemeth, 2006). The early explorations of resilience seemed to be based on an old method for assessing the adaptive capacity of any

system – observing how it responds to disrupting events. Since adaptive capacity is finite, patterns in the response to disrupting events provide information about limits and how the system behaves when events push it near or over those boundaries (Rasmussen et al., 1994; Cook and Rasmussen, 2005).

These similarities led us to explore the potential of the analogy between an organization as adaptive system and stress-strain plots in material science. We develop the stress-strain state space analogy of organizational resilience where the parameters and regions within the state space characterize the properties of the organization as an adaptive system. The stress-strain state space analogy for organizations includes different regimes – a uniform region where the organization stretches smoothly and uniformly in response to an increase in demands; and an extra region (x-region) where sources of resilience are drawn on to compensate for non-uniform stretching (risks of gaps in the work) in response to increases in demands. Different parameters in the state space capture different aspects of brittleness and resilience. For example, resilient organizations can make smooth transitions between regions and sub-regions in the state space, whereas those that are less resilient experience increasingly disruptive shifts, or fail to adapt at all and operations suddenly collapse (this pattern of brittleness occurred recently in the multi-day disruption of airline operations of jetBlue Airways in February 2007 resulting from a minor winter storm.)

Organizations can be well-calibrated or mis-calibrated as to how they actually operate as an adaptive system. Poorly calibrated organizations tend to perceive their system is placed well within the uniform region when the organization is actually operating in the x-region given the demands that actually occur in operations. Studies of high and low reliability organizations have documented the problems created when organizations are poorly calibrated with respect to their operating point within the state space (Weick et al., 1999).

The stress-strain state space provides an analogy of the adaptive capacity of organizations and proposes a set of parameters which characterize how organizations adapt as demands change. Estimating these parameters, including calibration, provides a means to measure

different aspects of brittleness and resilience for specific organizations that carry out risky activities.

Stress-strain States Space and Adaptive Capacity

Regions in Stress-Strain State Space

Our starting point is the need to assess how an organization responds to an increase in load or demands. Following the conventions of stress-strain plots in material sciences, the y-axis is the stress axis. We will here label the y-axis as the demand axis (D) and the basic unit of analysis is how the organization responds to an increase in D relative to a base level of D (Figure 1). The x-axis captures how the material stretches when placed under a given load or a change in load. In the extension to organizations, the x-axis captures how the organization stretches to handle an increase in demands (S relative to some base).

In materials there are two different regions of behavior: the elastic region where the material stretches uniformly under increasing load and a plastic region where the material begins to stretch non-uniformly until the distortions and gap accumulate and a fracture or failure point is reached. In the elastic or uniform region the response to increasing demands is proportional; in the plastic region the material cannot stretch completely to meet the demand (Figure 1).

In the first region – which we will term the *uniform* response region – the organization has developed plans, procedures, training, personnel and related operational resources that can stretch uniformly as demand varies in this region. This is the *on-plan* performance area or what Woods (2006) referred to as the *competence envelope*.

There are specific parameters that capture the capacity of a system to handle demands in uniform region (the slope and its length of a linear or proportional relationship) before the transition to the second region of behavior. These parameters of the uniform region are implicit in the *yield height* and represent the first-order adaptive capacity of the organization.

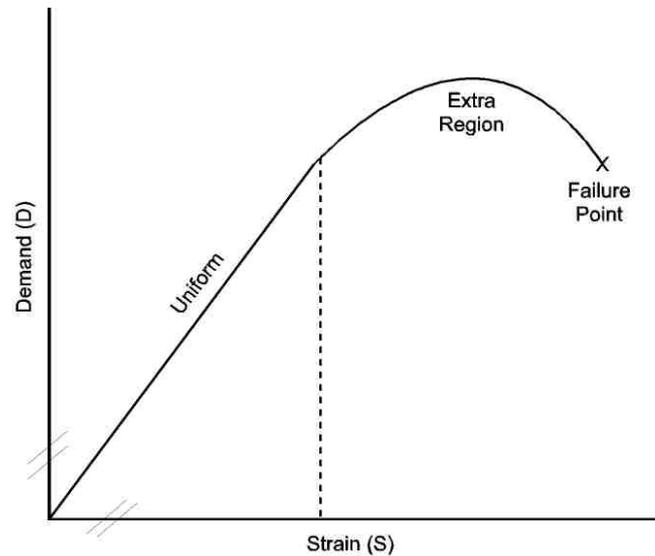


Figure 1: Basic stress-strain or demand-stretch state-space.

In the second region non-uniform stretching begins; in other words, ‘gaps’ begin to appear in the ability to maintain safe and effective production (as defined within the competence envelope) as the change in demands exceeds the ability of the organization to adapt within the competence envelope. At this point, the demands exceed the limit of the first order adaptations built into the plan-ful operation of the system in question. To avoid an accumulation of gaps that would lead to a system failure, active steps are needed to compensate for the gaps or to extend the ability of the system to stretch in response to increasing demands. These local adaptations are provided by people and groups as they actively adjust strategies and recruit resources so that the system can continue to stretch. We term this the ‘extra’ region (or more compactly, the x-region) as compensation requires extra work, extra resources, and new (extra) strategies. These local adaptations draw on sources of resilience to provide the extra adaptiveness the system requires to function under increasing demands without gaps accumulating to the failure point. This process continues to cope with increasing demands until either the second-order sources of

adaptiveness are exhausted and the system reaches a failure point, or until the system re-organizes and then functions in a new mode. In the latter case, the system re-structures into a new form with a new slope and length of uniform or on-plan performance in its new mode which is captured by a shift to a third region in Figure 2.

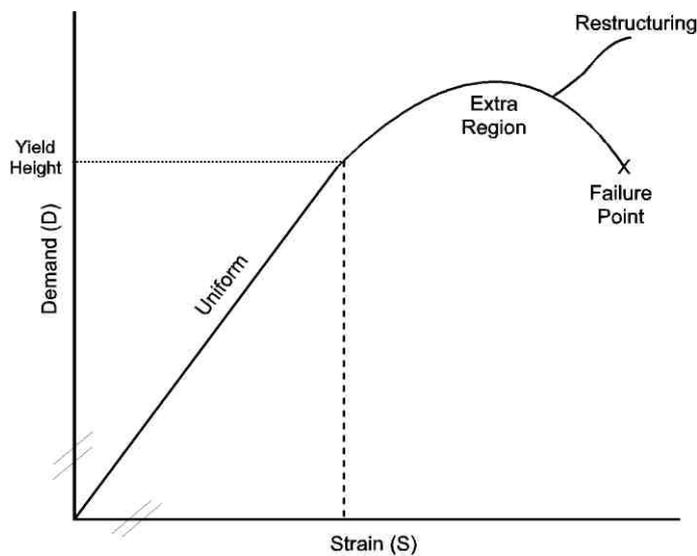


Figure 2: Expanded stress-strain state-space.

In the former case, exhausting sources of resilience in the x-region leads to the decompensation pattern of adaptive system breakdown as described in Woods and Cook (2006). The decompensation pattern is most easily seen in highly automated systems such as aircraft and cardiovascular physiology and the pattern consists of a two phase signature. In the first phase the automation works to compensate for the growing disturbance or disruption (e.g., a slow engine failure or increasing wing icing; Woods and Sarter, 2000). For the general case, this phase represents extra adaptive capacity being added to allow the system to continue to stretch as demands increase (the upswing in the basic x-region curve). The second phase occurs because the automation has limits on its capacity to adapt. When the automation's capacity is

exhausted, control of the system or parameter in question collapses without additional interventions. For the general case, the exhaustion of capacity to adapt as demands grow is represented by the movement to a failure point. This second phase is represented by the slope and distance to the failure point (the downswing portion of the x-region curve). Rapid collapse is one kind of brittleness; more resilient systems can anticipate the eventual decline or recognize that capacity is becoming exhausted and recruit additional resources and methods for adaptation or switch to a re-structured mode of operations (Figures 2 and 3). Gracefully degrading systems can defer movement toward a failure point by continuing to act to add extra adaptive capacity.

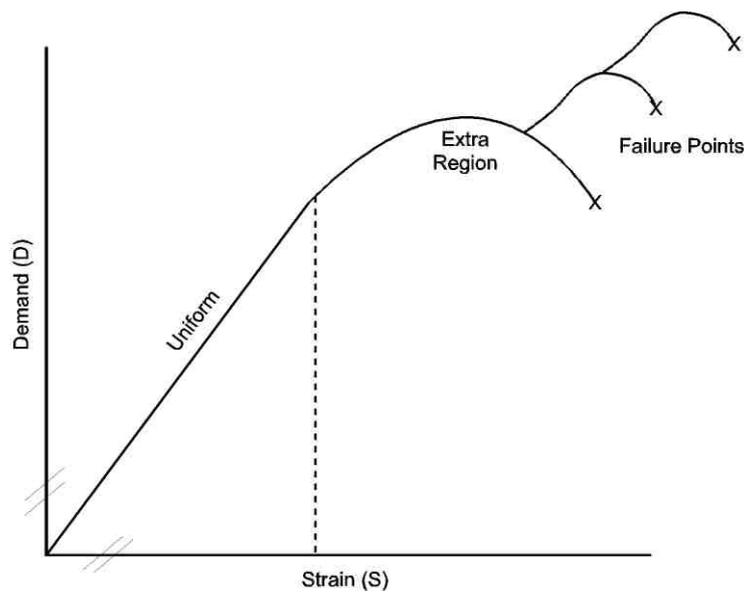


Figure 3: Sub-regions within the extra-region of the stress-strain state-space.

Characteristics of Organizational Stretching

For organizations the process of stretching in response to increasing demands is an active process. In the x-region, people in various roles in the organization can pick up on signs that it has moved into the x-

region. Recognizing that this has occurred (or is about to occur) leads people in these various roles to actively adapt to make up for the non-uniform stretching (or to signal the need for active adaptation to others). They inject new resources, tactics, and strategies to stretch adaptive capacity beyond the base built into on-plan behavior. People are the usual source of these gap-filling adaptations and these people are often inventive in finding ways to adapt when they have experience with particular gaps (Cook et al., 2000). Experienced people generally anticipate the need for these gap-filling adaptations to forestall or to be prepared for upcoming events (Klein et al., 2005; Woods and Hollnagel, 2006), though they may have to adapt reactively on some occasions after the consequences of gaps have begun to appear. (The critical role of anticipation was missed in some early work that noticed the importance of resilient performance, e.g., Wildavsky, 1988.)

The gap-filling adaptations in the x-region have costs associated with them in the form of:

- additional workload with the potential for workload and attentional bottlenecks,
- knowledge costs associated with developing and practicing shifts to new forms of work,
- expertise and the costs of preparing to have access to additional expertise when it is needed,
- opportunity costs when a resource is occupied and no longer available to meet future (or alternate concurrent) demands
- the financial costs associated with using resources when x-region adaptations are needed,
- the efficiency loss of having extra resources available in some form or through some mechanism to support x-region adaptation.

In contrast, behavior in the uniform region represents the adaptive capacity built into on-plan behavior. All systems have some capacity to adapt to changing demands built into the plans, procedures, and roles designed into the system. The costs of stretching to meet these forms or levels of demands are built into how the organization operates. People are provided means to learn to behave to roles in the on-plan system on the assumption that, if people will fulfill these roles, and not

go further, the system will work as designed and perform to expectations.

Leaders, groups, and organizations can monitor and assess their adaptive capacity and, as a result, self-consciously modify this capacity either in terms of expanding the range of demands on-plan performance can handle (moving the yield height) or in terms of how well the system is prepared to adapt. New methods can be learned and adopted throughout the organization, new contingencies can be planned for and tested in case they are needed, and the need to re-organize and function in new ways can be anticipated for extreme conditions (e.g., disaster plans build in new organizational structures and new authority for groups to make decisions autonomously).

There are feedback connections between demands placed on organizations and the success of organizations in expanding first order adaptive capacity. As an organization is successful, more may be demanded of it ('faster, better, cheaper' pressures) pushing the organization to handle demands that will exceed its uniform range. In part this relationship is captured in the Law of Stretched Systems (Woods and Hollnagel, 2006) – with new capabilities, effective leaders will adapt to exploit the new margins by demanding higher tempos, greater efficiency, new levels of performance, and more complex ways of work. This connection between past success and future demands also occur in modeling of complex adaptive systems (Carlson and Doyle, 2000; 2002; Csete and Doyle, 2002). These analyzes have demonstrated a trade-off where effort to increase optimal performance (essentially improving the parameters of the uniform region) paradoxically increases fragility or brittleness when unanticipated disturbances arise. Most models have assumed that efforts to improve system performance do not affect resilience/brittleness and that some efforts to make systems perform more optimally will also produce benefits on system resilience. Doyle and his colleagues' results reveal that these assumptions are wrong; instead, efforts to make systems perform more optimally on some dimensions and demands will increase the systems brittleness when it encounters situations or demands that fall outside that design envelope. In the stress-strain plot, improving the performance in the uniform region (increasing yield

height) squeezes out extra resources some of which may be much more than inefficiencies, but rather, function critically as buffers or sources of resilience that support adaptation in the x-region.

Another very important characteristic of the stress-strain state space representation of an organization's adaptive capacity is *calibration or mis-calibration*. Organizations can misunderstand where they are operating in the stress-strain space. If an organization thinks it is operating in the uniform region, but where the actual events occurring are forcing the work system into the x-region, the organization's model of how it works is mis-calibrated.

Calibration

Calibration refers to how accurate is one's model of one's own performance or capability. A basic observation about resilience / brittleness is that organizations and distant observers are mis-calibrated; that is they overestimate their capability of on-plan behavior to handle the situations and disruptions that can arise. Calibration can be captured very economically in the stress-strain space: the distant observer or manager implicitly overestimates the slope and the length parameters of the uniform region (they believe yield height is much higher than it actually is; Figure 2), or they have no accurate sense of the adaptations already made to accommodate demand factors. In other words, poorly calibrated organizations misrepresent capabilities of on-plan behavior to meet demands and they misrepresent the demands that will occur and how they would challenge on-plan behavior (Dekker, 2006).

Effective organizations are constantly looking for signs that specify how the organization actually operates and to use this information to be better calibrated (e.g., Weick et al., 1999; Sutcliffe and Vogus, 2001). One use of the stress-strain state space approach is to guide how organizations search for information and provide a means of integrating the results into an overall picture of changes in adaptive capacity.

Performance indicators are a specific form of measurement for organizations designed to provide continuing information about where the system is in relation to the boundaries that define its performance

envelope. Performance indicators have become a popular tool in performance-based management systems. They have been adopted in many diverse industries from healthcare to aviation and nuclear power since the 1980's (see Poister, 1983 for general background; Connelly et al., 1993 for process safety management; Majeed & Voss, 1995 and Blatt et al., 2006 for healthcare, for example).

Notwithstanding their popularity, there is a paucity of guidance on how to select specific indicators and how to combine them to form an overall picture. In the past indicators may be chosen based on considerations of tractability of data gathering, based on a consensus process involving various stakeholders, or based on organizational survey instruments developed to assess organizational culture such as safety culture. Examples of specific indicators include time since the last accident (so-called "event clocks") or the number of lost-time injuries in a period. These can be used for bench-marking between various sites in an organization or between competing companies, but these data tend to focus on performance in the past. Typically, performance indicators provide little information for project management who must make resource allocation decisions, and little information about real time changes that can effect safety. Whether based on actuarial or survey data, these indicators provide little anticipatory information to support proactive safety management (they only provide a sense of where you have been, like driving using only the rear-view mirror for directional guidance).

Many groups are trying to overcome these limitations with better survey instruments that attempt to capture critical characteristics of proactive organizations on safety (e.g., Blatt et al., 2006). Wreathall (2001; 2006) and Wreathall and Merritt (2003) has tried to select sets of indicators that map onto aspects of resilience: measures that indicate the onset of gaps in normal work practices as pressures grow and that reveal where workers develop gap-filling adaptations to compensate. In particular, these indicators are selected to reveal the potential for management to be unaware of the challenges, either in terms of changing demands or in terms of the need for workplace adaptations, and to be over-confident that current plans cover the changing demand profile.

The stress-strain state space analogy provides an integrated view of the set of parameters to be estimated to measure the adaptive capacity of an organization and the level of mis-calibration present in the organization. Gap-filling adaptations and incidents are indicators that the demands are exceeding the adaptive range of the uniform region or that the adaptive range of the uniform region is much lower than the organization believes (as captured in parameters related to yield height). An excellent example of gap-filling adaptations is a hospital emergency department (ED) when it is being “slammed” with a high number of difficult patients (Wears and Perry, 2006; Chapter X). As the number of patients needing care rose, the normal protocol for increasing numbers of patients (first using ED beds, then chairs, then using gurneys in the hallway) was insufficient to cope. As a result other aspects of the environment had to be utilized to accommodate the physical demands of handling extra patients (using hospital transports and seats in the physicians’ office as patient “beds”). Additional gap-filling adaptations occurred in how the staff track patients status, tests, and progress and how staff share that information as they coordinate care during surges. These situations are occurring so often in emergency departments that the Institute for Medicine in the US have recognized the ED as a critical brittle point in the national health care system (IOM, 2006).

The onset of the transition to the x-region can be detected by monitoring for the onset of gaps in work and by uncovering gap-filling adaptations (Cook et al., 2000). If an organization assumes or believes it is operating in the uniform region, such incidents and gap-filling adaptations look like very puzzling departures from standard methods, procedures and roles. This assumption leads the organization to badly misread the feedback. In the stress-strain analogy, incidents and other signs of adaptive behavior in the x-region are critical indicators that help estimate the parameters of the state space. Signs of gap-filling adaptations indicate what classes of disruptions or demands challenge the uniform region and indicate what sources of resilience are present to help accommodate these demands (Cook and Nemeth, 2006; Woods and Cook, 2006). Overall, the organization or other observers of the organization can be well-calibrated or mis-calibrated with respect to any

or all of the parameters associated with stress-strain state-space for that organization.

Adapting in the X-Region

Moving from uniform response requires some actors in the system to recognize the emerging shortfall as demands increase or change from some base. These people in various roles initiate adaptations as they draw on sources of resilience. For example, an individual might adapt their strategies and utilize other resources to make up for an adaptive shortfall. Or a key person might recognize the shortfall and redirect a team of people to handle the evolving situation. With still higher demands the group as a whole might recognize the shortfall and adapt as a unit to stretch with demand. In this way, the x-region is made up of a series of adaptive stretches: any one adaptive shift can become exhausted with the danger of decompensating toward a failure point; agents can recognize the shortfall and initiate another adaptive shift to forestall moving to close to failure points, and Chapter X illustrates two such shifts. Wears et al., (2006; this volume) describe this kind of sequence within hospital emergency departments as demands increase. Note that each shift within the x-region also relates to who controls access to the resources that need to be recruited: within an individual's control; within a team's control; within an organizational unit. Going outside one's range of control may necessitate restructuring as some parts of the organization may have to permit or authorize changes in how resources are controlled and released to various actors/roles (as occurs in hurricane emergency response).

If stresses continue to build, the system may shift to a new form through restructuring the work practices and organization. A critical factor for resilient organizations may be how much the organization practices recognizing shortfalls and making transitions. Studies of mission control as a successful adaptive organization strongly suggest this conclusion (Patterson et al., 1999). Note that the transitions within the x-region and from x-region to a re-structured form have to be planned for and practiced in general, as the specific situations that might arise are too varied and potentially unique to train for one by one.

One can think about the relationships over the different regions in the state space in terms of the costs associated with stretching in response to demand changes. In the uniform region the marginal cost of stretching ($\Delta c/\Delta d$) is very low or near zero (the costs – c – are built into the on-plan structure). In the x-region, the marginal cost of stretching is real and, in many cases, increases as demands increase (as cognitive or physical resources have to be recruited and deployed to compensate for the limited ability to stretch and therefore no longer are available for other tasks or contingencies). The cost of stretching in the x-region may go up as a series of steps that correspond to the shifts in range of adaptive response (Figure 3). For example, in the hospital emergency department there is a cost associated with an individual in a role adapting (and therefore using various kinds of cognitive and physical resources), a different but higher cost for a group or team adapting and a further different/higher cost associated with an entire unit adapting as demands increase. Shifts in the cost associated with stretching in the x-region may be the key marker for the series of adaptive responses that can make up the x-region.

There are costs associated with recruiting resources, and other parts of the organization may constrain or even resist releasing those resources (i.e., political costs). Taking resources for adaptive responses consumes them with respect to other activities. Note that costs can apply to more than to the owner/operators of the system. For example, one adaptation to an increased demand on a hospital emergency department is to close the department to new patients in order not to compromise the provision of care to the existing patients in the waiting room. The ‘cost’ here includes the shift in risk to those who are denied entry and must travel further and wait longer for care. The stress-strain state space suggests some ways to characterize the costs associated with bringing resources to bear.

A resource drawn on for x-region adaptive responses (a source of resilience) may be seen as an organizational inefficiency when viewed from the perspective of typical behavior in the uniform region. There are constraints between extra resources when the system is operating in the uniform range and when the system is operating in the x-region. Developments in complexity theory show that improving/expanding

the uniform region inevitably makes the system vulnerable to sudden large failures when situations arise that fall outside the competence envelope (Carlson and Doyle, 2002; Zhou et al., 2005). In the stress-strain analogy this means that as organizations improve they tend to erode or remove sources of resilience mistaking them for simple inefficiencies. Adapting in the x-region requires resources which easily can be eroded when organizations are under intense or continuing pressure to be “faster, better, cheaper” (Woods, 2006).

Resilience as Skill at Recognizing and Managing Transitions

The stress-strain space suggests a variety of parameters, regions and exhibits a number of characteristics (parameters that specify adaptive capacity in the uniform region such as yield height; parameters that specify adaptive capacity in the x-region such as how sharply decompensation occurs; and the state space approach can be extended to propose other possible parameters, for example, for graceful degradation). Which of these or other possibilities should be labeled “resilience”? The stress-strain space makes one thing very clear: calling the performance in the uniform region ‘resilience’ is wrong (Woods, 2006). The uniform region captures how all organizations develop planful behavior through a variety of means that exhibits some adaptive capacity. This is a first order adaptive capacity to absorb or respond to some disturbances and maintain system integrity (e.g., Hollings 1973). Effective management can expand the competence envelope improving the organization’s ability to perform in the face of designed for uncertainties and to expand the types or range of uncertainties that fall within planful responses – first order adaptive capacity. This might lead some to think that x-region adaptations are a kind of residual capacity to handle rare special situations that are not yet incorporated into the competence envelope (but will be as processes of continuous improvement proceed). But Csete and Doyle (2002) has shown that only focusing energy on improving first order adaptive capacity has a built-in limit. The inevitable commercial pressures that lead to increasing optimality of response to anticipated demands turn out to increase the brittleness of the organization to unplanned-for demands

(Carlson and Doyle, 2000). This optimality-fragility trade-off means that effective organizations in volatile or changing environments need to be able to invest in the potential for effective x-region or second-order adaptive behavior.

This leads to the possibility that some parameters for the x-region constitute resilience/brittleness (e.g., how rapidly the system moves toward a failure point as the extra adaptive resources become exhausted). When demands exceed the range of the uniform region, how well does the system bring new resources, tactics, and strategies to bear to cope? Brittle systems would quickly run out of second order adaptive capacity and move toward the failure point, while more resilient systems are able to recognize the need to shift and have practiced mechanisms available to provide the extra stretching needed for that situation. This is one plausible use of the label ‘resilience.’

Another plausible approach would be to reserve the label ‘resilience’ as a reference to the stress-strain state space in general for an organization. The stress-strain state space and its associated parameters capture a fairly full picture of the adaptive capacity of the organization in question. Rather than argue about which parameter or set of parameters should be labeled resilience seems somewhat beside the point. What is valuable is developing indicators that allow one to specify the space and how it is changing for a particular organization. In addition, resilient organizations work to be well-calibrated despite change. This means such organizations exert effort at overcoming the uncertainties and difficulties in estimating the parameters captured in the stress-strain state space. Using ‘resilience’ to refer to entire stress-strain state space would also emphasize the importance of being well-calibrated – knowing where you are, knowing how you are changing, and knowing your limits. If organizations are well-calibrated they will learn quickly and accurately where they have adaptive shortfalls and demonstrate the ability to make targeted effective interventions to enhance the particular aspect of adaptive capacity that is weak.

While both of these have merits, at this stage we would emphasize a third referent for the label “resilience.” It may be best to use the label “resilience” to refer to how well an organization can make transitions between regions or between sub-regions. Resilient organizations

practice transitions in the face of increases in demands: (a) from uniform to x-region; (b) from individual-based to team-based to unit-based in the x-region; and (c) from x-region to a restructured form involving other parts and levels of the organization. Evidence from successful organizations tends to indicate that they invest energy and resources to make these transitions work well (Patterson et al., 1999; Wears et al., this volume).

Limits and Extensions to the Stress-Strain Analogy

The stress-strain state space and the analogy to material science is a fruitful directions to explore in order to characterize an organization's adaptive capacity. The analogy makes plain there are many facets to a system's adaptive capacity and provides a framework to begin to organize them and to debate their relationships. The analogy also makes clear the distinction between first and second order adaptive capacity which often is blurred (Carlson and Doyle, 2002; Woods, 2006). This approach begins the process of parameterizing an organization's adaptive capacity so that discussions of resilience can move beyond vague allusions to how well a system bounces back from a stressor. The analogy also provides a framework that includes previous attempts to parameterize resilience. The parameters proposed by Holling (1973) and the parameters of "ball and cup" engineering models of resilience can be shown to be identical to the slope and length of the uniform region in the state space.

We also would like to note several limits to the stress-strain analogy. First, it maps all demands onto a single dimension when there are clearly different kinds of demands that disrupt on going plans in different ways and that challenge adaptive capacity in different ways. Future work is needed to break the overall demand dimension into different general sub-classes (taxonomies of classes of complicating factors provide a first step). Second, the analogy does not directly address the design problem. How does one set up or modify an organization to be more resilient across the parameters captured in stress-strain relationships? The analogy currently does not specify how to size or deploy reserves. Third, the analogy does not take advantage of the advances in modeling complex adaptive systems (Zhou et al.,

2005; Page, 2007). New adaptive systems modelling concepts and methods will need to be used because adaptive capacity and resilience concepts are non-intuitive and require us to step outside of the tendency to fall back on linear causal thinking.

Nevertheless, we have found the stress-strain plot to be very useful and promising way to look at the adaptive capacity of an organization. The stress-strain state space incorporates most of the different referents for the label resilience. It provides a set of formal parameters to be estimated by gathering data on actual operations in real organizations and provides a guide for why some indicators of resilience have proven valuable if indirect. But the ultimate test for this (or any) proposal to model/measure organizational resilience is whether it can show management how resources that otherwise look like an inefficiency to be squeezed out are actually part of what makes the system resilient as changing demands require new forms of adaptation.

4 REFERENCES

- Blatt, R., Christianson, M. K., Sutcliffe, K. M. and Rosenthal, M. M. (2006). A sensemaking lens on reliability. *Journal of Organizational Behaviour*, 27, 897–917.
- Carlson, J. M. and Doyle, J. C. (2000). Highly Optimized Tolerance: Robustness and Design in Complex Systems. *Physical Review Letters*, 84(11), 2529 – 2532.
- Carlson, J. M. and Doyle, J. C. (2002). Complexity and Robustness. *Proceedings of the National Academy of Sciences*, 99, 2538 – 2545.
- Carter, N., Klein, R., & Day, P. (1992). *How organizations measure success: the use of performance indicators in government*. London: Routledge.
- Connelly, E., et al. (1993). Method for Building Performance Measures for Process Safety Management. Paper presented at the International Process Safety Management Conference and Workshop, San Francisco, CA.
- Cook, R. I., Render M. L. and Woods, D. D. (2000). Gaps in the continuity of care and progress on patient safety. *British Medical Journal*, 320, 791 – 794, March 18, 2000.
- Cook, R. I. and Nemeth, C. (2006). Taking Things in Stride: Cognitive Features of Two Resilient Performances. In E. Hollnagel, D. D. Woods and N. Leveson, eds., *Resilience Engineering: Concepts and Precepts*. Ashgate, Aldershot, UK.
- Cook, R. and Rasmussen, J. (2005). “Going Solid”: A Model of System Dynamics and Consequences for Patient Safety. *Quality and Safety in Health Care*. 14, 130-134.
- Csete, M. E. & Doyle, J. C. (2002). Reverse engineering of biological complexity. *Science*, 295, 1664-1669.
- Dekker, S. W. A. (2006). In E. Hollnagel, D.D. Woods & N. Leveson (Eds.), *Resilience Engineering: Concepts and Precepts* (pp. 69-76). Aldershot, UK: Ashgate.
- Holling, C.S. (1973). Resilience and Stability of Ecological Systems, *Annual Review of Ecological Systems*. 4:1-23.
- Hollnagel, E., Woods, D. D., & Leveson, N. (Eds.) (2006). *Resilience Engineering: Concepts and Precepts*. Aldershot UK: Ashgate.
- IOM Committee on the Future of Emergency Care in the US. (2006).

Hospital-Based Emergency Care: At the Breaking Point. Washington, DC: National Academies Press.

Klein, G., Pliske, R., Crandall, B. and Woods, D. D. (2005). Problem Detection. *Cognition, Technology, and Work*. 7(1), 14-28.

Majeed, F. A., & Voss, S. (1995). Performance indicators for general practice (Editorial). *British Medical Journal*, 311, 209-210.

Page, S. E. (2007). *The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies*. Princeton NJ: Princeton University Press.

Patterson, E. S., Watts-Perotti, J. & Woods, D. D. (1999). Voice loops as coordination aids in space shuttle mission control. *Computer Supported Cooperative Work: The Journal of Collaborative Computing*. 8(4), 353-71.

Perry, S. J., Wears, R. L. & Anderson, B. (2006). Extemporaneous adaptation to evolving complexity: A case study of resilience in healthcare. In E. Hollnagel and E. Rigaud (Eds.) *Proceedings of the Second International Symposium on Resilience Engineering*. Juan-les-Pins, France, Nov. 8-10, 2006.

Poister, T. H. (1983). *Performance Monitoring*. Lexington, MA: Lexington Books.

Rasmussen, J., Pejtersen, A. M. and Goodstein, L. (1994). *Cognitive Systems Engineering*. New York: Wiley.

Schenk, J., Allen, T. and Woods, D. D. (this volume).

Sutcliffe, K. and Vogus, T. (2003). Organizing for resilience. In K. S. Cameron, I. E. Dutton, & R. E. Quinn (Eds.), *Positive Organizational Scholarship*. San Francisco: Berrett-Koehler, p. 94-110.

Wears, R. L. and Perry, S. J. (2006). "Free fall" – a case study of resilience, its degradation, and recovery in an emergency department. In E. Rigaud & E. Hollnagel (Eds.) *Second Symposium on Resilience Engineering*. Juan-les-Pins, France, Nov. 8-10, 2006.

Wears, R. L., Perry, S. J. and Anders, S. (this volume).

Weick, K. E., Sutcliffe, K. M., & Obstfeld, D. (1999). Organizing for high reliability: Processes of collective mindfulness. *Research in Organizational Behavior*, 21, 13-81.

Wildavsky, A. (1988). *Searching for safety*. New Brunswick: Transaction Books.

- Woods, D. D. (1988). Coping with complexity: The psychology of human behavior in complex systems. In L.P. Goodstein, H. B. Andersen, and S. E. Olsen (Eds.), *Mental Models, Tasks and Errors* (pp. 128-148). London: Taylor & Francis.
- Woods, D. D. (2006). Essential Characteristics of Resilience for Organizations. In E. Hollnagel, D. D. Woods & N. Leveson (Eds.), *Resilience Engineering: Concepts and Precepts* (pp. 69-76). Aldershot, UK: Ashgate.
- Woods, D. D. and Cook, R. I. (2006). Incidents: Are they markers of resilience or brittleness? In E. Hollnagel, D. D. Woods & N. Leveson (Eds.), *Resilience Engineering: Concepts and Precepts* (pp. 69-76). Aldershot, UK: Ashgate.
- Woods, D. D. and Hollnagel, E. (2006). *Joint Cognitive Systems: Patterns in Cognitive Systems Engineering*. Boca Raton FL: Taylor & Francis.
- Woods, D. D. and Patterson, E. S. (2000). How Unexpected Events Produce an Escalation of Cognitive and Coordinative Demands. In P. A. Hancock and P. Desmond (Eds.), *Stress Workload and Fatigue*. Mahwah, NJ: Lawrence Erlbaum Associates.
- Woods, D. D. and Sarter, N. (2000). Learning from Automation Surprises and Going Sour Accidents. In N. Sarter and R. Amalberti (Eds.), *Cognitive Engineering in the Aviation Domain*. Mahwah, NJ: Lawrence Erlbaum Associates.
- Wreathall, J. (2001). Systemic Safety Assessment of Production Installations. In *World Congress: Safety of Modern Technical Systems*, Saarbrücken, Germany, TUV-Verlag GmbH, Cologne, Germany.
- Wreathall, J. (2006). Properties of Resilient Organizations: An Initial View. In E. Hollnagel, D. D. Woods and N. Leveson (Eds.), *Resilience Engineering: Concepts and Precepts* (pp. 275-285). Aldershot, UK: Ashgate.
- Wreathall, J. and Merritt, A. C. (2003). Managing Human Performance in the Modern World: Developments in the US Nuclear Industry. In G. Edkins and P. Pfister (Eds.) *Innovation and Consolidation in Aviation*. Aldershot (UK): Ashgate.
- Zhou, T., Carlson, J. M., and Doyle, J. (2005). Evolutionary Dynamics and Highly Optimized Tolerance, *Journal of Theoretical Biology*. 236, 438-447.