

---

## 4 Essentials of resilience, revisited

*David D. Woods*

---

### INTRODUCTION

The idea that systems have a property called ‘resilience’ has emerged and grown extremely popular in the last two decades. The idea arose from multiple sources and has been examined from multiple disciplinary perspectives including: systems safety (Hollnagel et al. 2006), complexity (Doyle and Csete 20011), human organizations (Sutcliffe and Vogus 2003), ecology (Walker and Salt 2006), and others. However, with popularity has come noise and confusion as the label continues to be used in multiple, diverse and, sometimes, incompatible ways.

I advocated for studying resilience beginning in 2000 and then 2003 in the aftermath of accidents, notably the National Aeronautics and Space Administration (NASA) space exploration mishaps in 1999 and the space shuttle *Columbia* accident in 2003 (Woods 2003, 2005a). The idea was that (1) systems become increasingly brittle under ‘faster, better, cheaper’ pressure; (2) signs of increasing brittleness are discounted under ‘faster, better, cheaper’ pressure; (3) investments in building and sustaining sources of resilient performance need to be valued and managed in order to offset the brittleness of systems and to manage increasing complexity (Woods 2006).

In the first *Resilience Engineering* book (Hollnagel et al. 2006) I opened the volume with a chapter called ‘Essentials’, which introduced a few initial fundamentals. The intent was to begin and stimulate discovery of the fundamentals that lead to resilient performance in complex systems. Based on results from interdisciplinary inquiry since 2006, this chapter expands on the ‘essentials’ because there has been progress, despite the noise generated by hyper-popularity. The intent again is to stimulate further inquiry and discovery of fundamentals. This goal remains crucial given the confusion that has resulted from overuse of resilience as a label. Understanding what generates resilient performance in an increasingly interdependent and multi-scaled world continues to be important for every industry sector and for communities at all levels.

This chapter covers some additional ‘essentials’ to demonstrate that we are able to see through the diversity of systems and events to understand fundamentals that govern adaptive systems across multiple scales. Studies of resilience in action have revealed a rich set of patterns about how some systems overcome brittleness. Concepts have emerged on what makes the difference between resilient performance and brittle collapse of performance as surprises occur and disturbances cascade. The chapter revisits and updates several of the essentials introduced in the 2006 chapter and then covers additional ones – initiative, managing the expression of initiative, and reciprocity. These are covered because they raise scientific challenges, at the same time as they provide practical guidance (or could do so) to working organizations plagued by complexity, surprise and brittleness (which is virtually all of them).

## ESSENTIALS 1

### Poised to Adapt

One of the first things we noted about resilience is that it refers to the potential for future adaptive action when conditions change – resilience concerns the capabilities a system needs to respond to inevitable surprises. Responding to surprise requires preparatory investments that provide the potential for adaptive action in the future.

*Definition* Adaptive capacity is the potential for adjusting patterns of activities to handle future changes in the kinds of events, opportunities and disruptions experienced, therefore, adaptive capacities exist before changes and disruptions call upon those capacities.

Systems possess varieties of adaptive capacity, and resilience engineering seeks to understand how are these built, sustained, degraded and lost. Adaptive capacity means a system is poised to adapt, it has some readiness or potential to change how it currently works – its models, plans, processes and behaviors (Woods 2018). Adaptation is not about always changing the plan, model or previous approaches, but about the potential to modify plans to continue to fit changing situations. Space mission control is the positive case study for this capability, especially how space shuttle mission control developed its skill at handling anomalies, even as they expected that the next anomaly to be handled would not match any of those they had planned and practiced for (Watts-Perotti and Woods 2009). Studies of how successful military organizations adapted to handle surprises provide another rich set of contrasting cases, as in Finkel (2011).

Adaptive capacity does not mean a system is constantly changing what it has planned or does so all the time, but that the system has some ability to recognize when it is adequate to continue the plan, to continue to work in the usual way and when it is not adequate to continue on, given the demands, changes and context ongoing or upcoming. Adaptation can mean continuing to work to plan, but, and this is a very important but, with the continuing ability to reassess whether the plan fits the situation confronted – even as evidence about the nature of the situation changes and evidence from the effects of interventions changes. The ability to recognize and to stretch, extend or change what you are doing or what you have planned, has to be there in advance of adapting, even when there are no adjustments to behavior visible to outside observers. There are great difficulties to be overcome in studying how a system is poised to adapt, in assessing how much capability is available or will be needed and in uncovering what factors will contribute to resilient performance when future challenges arise. This capability can be extended or constricted as challenges arise, expanded or degraded over cycles of change and re-directed or become stuck as conditions evolve into new configurations.

### Brittleness

All systems have an envelope of performance, or a range of adaptive behavior, owing to finite resources and the inherent variability of its environment in a continuously changing world. A bounded competence envelope raises the question of how systems perform when events push the system near the edge of its envelope. Descriptively, brittleness is

how rapidly a system's performance declines when it nears and reaches its boundary. Brittle systems experience rapid performance collapses, or failures, when events challenge boundaries. One difficulty is that the location of the boundary is normally uncertain and moves as capabilities and conditions change.

In 2006, I was using resilience to mean the opposite of brittleness, or, how to extend adaptive capacity in the face of surprise (Woods 2005a, 2009). I was asking how systems stretch to handle surprises, since systems with finite resources in changing environments are always experiencing and stretching to accommodate events that challenge boundaries. Without some capability to continue to stretch in the face of events that challenge boundaries, systems are more brittle than stakeholders realize (Woods and Branlat 2011). All systems, however successful, have boundaries and experience events that fall outside these boundaries – surprises, given that no system escapes the constraints of finite resources and changing conditions for long.

However, I found I could no longer use resilience as the opposite of brittleness; resilience-as-label pointed to too many different issues and concepts. As a result, I coined the term graceful extensibility to refer to the opposite of brittleness (Woods 2015). Graceful extensibility is the ability of a system to extend its capacity and to adapt when surprise events challenge its boundaries. It can be thought of as a blend of two traditional terms – graceful degradation and software extensibility. Software engineering emphasizes the need to design in advance properties that support the ability to extend capabilities later, without requiring a major revisions to the basic architecture, as conditions, contexts, uses, risks, goals, and relationships change. Ironically, the best examples of graceful extensibility come from biology (for example, Fairhall et al. 2001; Csete and Doyle 2002; Meyers and Bull 2002; Wark et al. 2007). To enhance safety, we desire a system exhibit graceful degradation where the system's performance declines slowly as disturbances grow and cascade, rather than experience a brittle fracture. The combination of software extensibility and graceful degradation highlights how adaptation at the boundaries is active and critical for how a system grows and adjusts in a changing world, not simply a softer degradation curve when events challenge base competencies.

With low graceful extensibility, systems exhaust their ability to respond as challenges grow and cascade. As the ability to continue to respond declines in the face of growing demands, systems with low graceful extensibility risk a sudden collapse in performance. With high graceful extensibility, systems have capabilities to anticipate bottlenecks ahead, to learn about the changing shape of disturbances or challenges prior to acute events, and possess the readiness-to-respond to meet new challenges (Woods and Branlat, 2011). As a result, systems with high graceful extensibility are able to continue to meet critical goals and even recognize and seize new opportunities to meet pressing goals.

Boundary refers to the transition zone where systems shift regimes of performance. This boundary area can be more crisp or blurred, more stable or dynamic, well-modeled or misunderstood. Brittleness and graceful extensibility refer to the behavior of the system as it transitions across this boundary area. The latter refers to a system's ability to adapt how it works to extend performance past the boundary area into a new regime of performance invoking new resources, responses, relationships and priorities. This process has been studied in medicine; for example, Wears et al. (2008) describe how medical emergency rooms adapt to changing and high patient loads, and Chuang et al. (2018) describes how emergency departments adapted during a mass casualty event.

## Trade-Offs

The third essential I noted in 2006 is that complex adaptive systems face fundamental trade-offs, and how systems manage these trade-offs determines the ability to demonstrate resilient performance in complex environments. I specifically examined one candidate – the acute–chronic trade-off – that was particularly relevant to NASA's accidents in 1999 and 2003. The failure cases highlighted the sacrifice judgment that arises in managing goal conflicts and how production pressure leads to discounting of evidence that a system is encountering events at or beyond its competence envelope (Woods 2009).

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 taking too much risk. However, if uncertain warning signs always lead to sacrifices on acute goals, the organization may be unable to meet the pressures for greater efficiency and throughput. It is easy for organizations that are working hard to advance the acute goal set to see responding to such signs as risks to efficiency goals or as low-probability events that do not need a response given current pressures. Also, they usually can point to a record of apparent success and improvement. This occurred notably in the run-up to the *Columbia* accident as the risk of insulating foam (debris) strikes during launch were discounted by managers under schedule pressure. Afterwards, it was clear the foam debris strikes were events well outside the boundaries established by safety and risk analyses – energy in this source of debris strikes was two orders of magnitude greater than the assumed maximum for debris strikes, and this was a different source of debris in a different phase of flight than had been assumed in risk analyses.

Other work in the 2000s also identified fundamental trade-offs (Hoffman and Woods 2011) – optimality–brittleness, termed robust yet fragile (RyF) in Doyle's work (Doyle et al. 2005; Alderson and Doyle 2010) and efficiency–thoroughness (Hollnagel 2009). The question was, what types of governance or architectures would balance these fundamental trade-offs over multiple cycles of change? Doyle asked this question in terms of what are the properties of architectures that balance basic trade-offs to sustain adaptive capacities for layered networks in biological and physical systems using complexity science and control engineering techniques (for example, Csete and Doyle 2002; Doyle and Csete 2011; Chen et al. 2015). Ostrom asked this question in terms of how human social systems develop and sustain polycentric governance mechanisms over long time scales despite pressures for groups to act selfishly (Ostrom 1998, 1999, 2012).

The ability to continue to adapt over multiple cycles of change is a type of adaptive capacity – sustained adaptability, which refers to the ability to continue to adapt to changing environments, stakeholders, demands, contexts and constraints, that is, the ability to adapt how the system in question adapts (Woods 2015). Some layered networks, biological systems, complex adaptive systems and human systems demonstrate sustained adaptability. Well modeled examples in biology can be found in Li et al. (2014) on how the cardiovascular system adapts to handle widely varying loads, and Wark et al. (2007) on how sensory systems adapt to widely varying contexts. However, most layered networks can fail to sustain adaptability when confronting new periods of change, that is, they get stuck in adaptive shortfalls, unravel and collapse, regardless of their past record of successes (for example, failures in matching markets in Roth 2008).

Ostrom's and Doyle's work, from very different starting points, leads resilience engineering to some basic scientific and practical challenges:

1. What governance or architectural characteristics explain the difference between networks that produce sustained adaptability and those that fail to sustain adaptability?
2. What design principles and techniques would allow someone to design or manage a layered network that can produce sustained adaptability?
3. How would someone assess or know if they succeeded in their design or management efforts to endow a system with the ability to sustain adaptability over time (similar to evolvability from a biological perspective and like a new kind of stability from a control engineering perspective)?

In socio-technical systems, sustained adaptability addresses a system's dynamics over life cycles or multiple cycles of change. The architecture of the system needs to be equipped at earlier stages with the wherewithal to adapt or be adaptable when it faces predictable changes and challenges across its life cycle. Central to resilience is identifying what basic architectural principles provide the needed flexibility to continue to adapt over long scales. Over life cycles, and the system in question will have to adapt to seize opportunities and respond to challenges by readjusting itself and its relationships in the layered network.

## ESSENTIALS 2

Progress has been made on the governance policies that operate across layered networks in biological systems, human systems and technological systems – what governance policies sustain the ability of the network to continue to adapt and avoid falling into traps in the trade spaces as conditions change over long timescales. This progress indicates new research challenges. In the shorter term, the results also point to pragmatic guidance for those who manage socio-technical systems. The chapter considers two key essential characteristics of layered networks that demonstrate sustained adaptability – initiative and reciprocity.

### **Initiative**

In work on human adaptive and complex systems, it turns out that:

1. initiative is critical to adaptive capacity as a unit has to possess some degree and form of initiative to contribute to graceful extensibility; and
2. interactions across roles and units in the network affect the expression of initiative, and those effects depend on the potential for surprise in that setting.

For example, the interactions across roles tend to support initiative when the potential for surprise is high. This occurs in military operations in adversarial situations, as in Shattuck and Woods (2000) and Finkel (2011), and in emergency departments' ability to respond to surges in patient load, as in Wears et al. (2008) and Chuang et al. (2018).

Interestingly, initiative arises in studies of human systems, but it is an explicit factor much less often in engineered systems. However, even in human systems it is difficult to identify what capabilities initiative refers to (for example, Shattuck and Woods 2000). What is initiative in the context of human adaptive systems? Note: in the following I use ‘plan’ as a generic label to refer to any plan and any form of embodiment of a plan in procedures, machines, computations, or automation. Also I use the label ‘unit’ to refer generically to any role or group at any layer.

Initiative is

1. the ability of a unit to adapt when the plan no longer fits the situation, as seen from that unit’s perspective ;
2. the willingness (even the audacity) to adapt planned activities to work around impedes or to seize opportunities in order to better meet the goals/intent behind the plan; and
3. when taking the initiative, the unit begins to adapts on its own, using information and knowledge available at that point, without asking for and then waiting for explicit authorization or tasking from other units.

First, initiative requires the ability to see or recognize that the plan is not making progress or that the plan does not fit the situation. Has the situation changed relative to the assumptions behind the plan? Is an impasse or gap looming ahead or becoming clear? Is there an opportunity to better achieve the intent of goals behind the plan? The unit in the situation has to be able to see that the plan is no longer working to achieve goals – and has to do this prior to a complete failure of the plan – that is, there is some anticipatory or looking ahead component to initiative.

Second, initiative means that adaptations to the plan occur in order to better meet goals and intent when difficulties accumulate or when new opportunities arise. The latter is particularly important – recognizing and seizing opportunities is a critical part of what it means to show initiative.

Interestingly, people who have experienced the necessity for initiative in order to be sufficiently adaptive given real pressures and the possibility of tangible failures, often want to use the word ‘audacity’ to capture the meaning of initiative. Audacity in some ways refers to a characteristic of the people adapting – a readiness to go beyond the plan or be poised to adapt. The use of audacity also refers to the risks that we take on when we adaptes plans to better fit situations.

Third, once a unit has recognized the need to adapt, it has to be willing to begin that process on its own, based on its view of the situation and its perception of the goals, intent, priorities and trade-offs. The unit that is adapting may seek additional guidance and expand coordination over more roles, but it cannot wait for new tasking, additional partners or authorization before initiating actions.

### **Why Does Initiative Matter in Human Adaptive Systems?**

Initiative is almost synonymous with an ability to adapt, especially the role of anticipation. Simply working to plan or carrying out standard activities for a role is not sufficient to handle exceptions, anomalies and surprises, regardless of the contingencies built into the

standard practice. All plans have a bounded competence envelope, given finite resources in a world of continuing change (Woods 2006).

However, even more important, initiative is necessary in order for adaptive systems to keep pace with events and the potential for difficulties to cascade. When a unit confronts situations that challenge plans, delays are inevitable and significant if the unit must first inform others and then wait for new instructions before initiating a response. If this is the method for revising plans in progress, performance is guaranteed to be slow and stale, with limited ability to keep pace with change. The risk is that adapting in time cannot match the tempo of events, and breakdowns occur as ‘challenges grow and cascade faster than responses can be decided on and deployed to effect’ (Woods and Branlat 2011, p. 130). This risk is called decompensation (when the system exhausts its capacity to adapt as disturbances or challenges cascade), and it is the primary of three basic failure modes for adaptive systems (Woods and Branlat, 2011).

The 2010 Knight Capital runaway automation collapse in financial trading is an example of this breakdown in a modern highly automated layered organization.<sup>1</sup> In this case, one part of the organization deployed new software in order to keep up with and take advantage of changes in the industry – all changes regardless of type become changes in software for computerized financial trading. The rollout did not go as expected and produced anomalous behavior. The team then tried to roll back to a previous software configuration as is standard practice for reliability, but the rollback produced more anomalous behavior. The team did not understand why this was occurring and was not sure what kind of interventions would be needed to resolve the situation. The team did not take the initiative and felt it did not have the authority to stop trading. By the time the team was able to decide to go to upper management and tell upper management that there was a problem, that they did not understand the problem and that the only action available was to stop trading, tens of minutes had gone by. When upper management approved and trading was stopped, it was already too late – the company had lost almost a half a billion dollars and was effectively bankrupt. The interaction and communication vertically across layers was slow and stale, unable to keep pace with the cascade of effects. This case illustrates a combination of the two basic failure modes for adaptive systems: inability to keep pace with events and working at cross-purposes, in this case across the vertical layers, combined to produce a slow and stale response that resulted in the collapse of the organization.

In contrast, studies of how emergency departments handle large surges in patient load, as occurs in mass casualty events, reveal the critical role of initiative in keeping pace or staying ahead of a cascade (for example, Cook and Nemeth 2006; Wears et al. 2008; Chuang et al. 2018). Initiative in these situations proves critical to anticipate and prepare for bottlenecks ahead, to avoid or reduce the risk of poor performance during a workload crunch, and to synchronize coordination across roles and units.

Deary examined how a large transportation firm had learned to reconfigure relationships across roles and layers to keep pace with unpredictable demands, and how the organization used these techniques during hurricane Sandy in autumn 2012 (Deary et al.

---

<sup>1</sup> See <https://michaelhamilton.quora.com/How-a-software-bug-made-Knight-Capital-lose-500M-in-a-day-al-most-go-bankrupt> and <https://www.kitchensoap.com/2013/10/29/counterfactuals-knight-capital/> (both accessed 24 April 2017).

2013). To adapt effectively, the organization had to re-prioritize over multiple conflicting goals, sacrifice cost control processes in the face of safety risks, value timely responsive decisions and actions, coordinate horizontally across functions to reduce the risk of missing critical information or side effects when re-planning under time pressure, control the cost of coordination to avoid overloading already busy people and communication channels, and push initiative and authority down to the lowest unit of action in the situation to increase the readiness to respond when new challenges arose. New temporary teams were created quickly to provide critical information updates (weather impact analysis teams). They created temporary local command centers where key personnel from different functions worked together to keep track of the evolving situation and to re-plan. The horizontal coordination in these centers worked to balance the efficiency–thoroughness trade-off (Hollnagel 2009) in a new way for a situation that presented surprising challenges and demanded high responsiveness.

In the case of disruptions, this highly adaptive firm was able to synchronize different groups at different echelons even with time pressure, surprises, goals conflicts and trade-offs. The firm had developed mechanisms to keep pace with cascades and expand or speed coordination across roles, though these sacrificed economics and standard processes, because this firm's business model, environment, clientele and external events regularly required adaptation in smaller or larger ways.

### **What Governs the Expression of Initiative?**

Initiative is fundamental to adaptive capacity. Units, regardless of the scale at which they operate, have to possess some degree and form of initiative to contribute to graceful extensibility. However, initiative can run too wide when undirected, leading to fragmentation, working at cross-purposes and mis-synchronization across roles. Also, initiative can be reduced or eliminated by pressure to work-to-rule or work-to-plan, especially by threats of sanctions should adaptations prove ineffective or erroneous. Emphasis on work-to-rule or work-to-plan – compliance cultures – limits adaptive capacity when events occur that do not meet assumptions in the plan, when it passes block progress or when unforeseen opportunities arise (Shattuck and Woods 2000; Woods and Shattuck 2000; Perry and Wears 2012).

The hinge between these is the potential for surprise (Woods 2009; Woods and Branlat 2011). In worlds where the experience and risk of surprise is tangible – for example, military operations (Finkel 2011), space mission operations (Watts-Perotti and Woods 2009) and emergency medicine (Wears et al. 2008) – initiative is pushed down to the unit of action in that world to build adaptive capacity that goes beyond what was accounted for in the plans themselves. Worlds that appear stable, that trust in their models of what produces successes and that invest in building comprehensive plans tend to explain poor performance as the result of failures to comply with rules or plans. In the former, revision and adaptation is valued in the recognition that plans are limited and surprise is inevitable; in the latter, compliance is valued in the belief that the world is well understood and changes will be announced clearly and well in advance, with time to prepare. The rise of resilience as a critical societal need recognizes that the latter state is at best temporary, if not an illusion. Over cycles of change, (1) the competence boundaries of plans and automata will be challenged by surprises, (2) compliance-based systems will be brittle and

exhibit the characteristics of sudden performance collapses as a result, and (3) graceful extensibility is a required fundamental capacity for adaptive systems at all scales.

Resilience engineering is then left with the task of specifying how a governance system, or layered network architecture, balances the expression of initiative as the potential for surprise waxes and wanes. The pressures generated by other interdependent units either energizes or reduces initiative and, therefore, the capacity to adapt. These pressures also change how initiative is synchronized across adapting units given their perspective, role and relationships. The pressures constrain and direct how the expression of initiative prioritizes some goals and sacrifices others goals when conflicts in the trade spaces intensify, as in the stories of NASA's failures that precipitated one branch of inquiry into resilient systems. Importantly, changing pressure on a role influences what goals are sacrificed when demands and risks grow. Understanding these processes is a scientific challenge and a practical necessity.

### **Reciprocity**

Social science, network science and computational social science have all identified that reciprocity across roles, units and layers is essential in human adaptive systems (for example, Ostrom and Walker 2003). The characterization of reciprocity varies: sometimes it is used in general to make the point that cooperation is built on trust and altruism; at other times, particularly for computational simulations of network behavior, reciprocity is defined merely as mutual linkage or influence between two nodes. Ostrom (2003) argues that polycentric governance systems require mechanisms and norms that support reciprocity.

How can or should resilience engineering make use of reciprocity? On the one hand, social norms and trust are part of building collaboration across units in a network; on the other, how do we operationalize this to design systems that overcome brittleness and build in the capacity to adapt over multiple cycles of change? Reducing reciprocity to mutual linkage in a directed graph is computable, but seems to throw away all of the insight associated with the concept. Or do solutions already exist in the large body of work in experimental microeconomics (for example, Roth 2008)?

Ironically, inspired by the work of an eminent social scientist (Ostrom) and by an eminent control engineering theorist (Doyle), and starting with the idea of resilience as graceful extensibility (which derived from my previous work on how human and machine cognitive systems cope with complexity), an integration with actionable potential emerged (Woods 2018).

A classic sense of reciprocity in collaborative work is commitment to mutual assistance (for example, Klein et al. 2005). Reciprocity can be explained in terms of the interaction between two roles. Units 1 and 2 demonstrate reciprocity when unit 1 takes an action to help unit 2 that gives up some amount of immediate benefit relative to its scope of responsibility. The sacrifice by unit 1 relative to a narrow view of its role allows for a larger, longer-run benefit for both units 1 and 2 relative to the broader goals of the network in which these two units exist. However, in helping another unit cope with its difficulties, unit 1 is relying on unit 2 to reciprocate in the future – when unit 1 needs help, unit 2 will be responsive and willing to take actions that will give up some benefit to that role in the short run in order to make both roles better off relative to common goals and constraints.

When reciprocity is seen as commitment to mutual assistance, one unit is donating from their limited resources now to help another in their role, in order for both to achieve

benefits for overarching goals. Each unit operates under limited resources in terms of energy, workload, time and attention for carrying out each role. Diverting some of these resources creates opportunity costs and workload management costs for the donating unit. Alternatively, a unit can ignore other interdependent roles and focus their resources on meeting just the performance standards set for their role alone, even though this can be short-sighted and parochial. In the latter case, 'role retreat' undermines coordinated activity and increases brittleness. Negative accountability systems tend to produce role retreat (Woods 2005b).

A potential instability arises because there is a lag between donating limited resources and when that investment will pay off for the donating unit. It could pay off in better performance on larger system goals – an investment toward common pool goals – but that effect might be difficult to see and improve matters for the donating unit. It could also pay off in the future when other units make donations of their limited resources to help the unit that is donating now when it experiences challenges. The investment is now, definite and specific; the benefit is uncertain and in the future. Also, the receiving unit could act selfishly, exploiting the donating unit, by not being willing to reciprocate in the future. Aligning the multiple goals will always require relaxing (a sacrifice) some local short-term (acute) goals in order to permit more global and longer-term (chronic) goals to be addressed. Interdependent units in a network should show a willingness to invest energy to accommodate other units specifically when the other units' performance is at risk (again, look at how emergency departments handle surges in patients). Alternatively, pressure to meet or improve on role-specific performance criteria can lead a unit to just performing alone, walled off inside its narrow scope and sub-goals (or to 'defect' from collaborations). Pressures for compliance and productivity (leaner systems struggling to be faster, better and cheaper) undermine a willingness to reach across roles and coordinate when anomalies and surprises occur. The latter tendency increases brittleness when that system operates as part of interdependent units in a larger complex network of influences.

All I have provided above is another description of what is already understood in the social sciences. This description now links reciprocity to brittleness and to graceful extensibility, which will prove of some importance. We can go further; add in a basic idea from control engineering – risk of saturation – and throw in some linkages to biological systems (Woods 2018).

Control systems fail when their capacity to respond saturates, that is, they run out of capacity to continue to respond as disturbances continue to grow. This is always very bad if you want to control a physical process or vehicle. However, all control systems, both physical and biological, have limits. Brittleness occurs if events push these systems near their limits unless another mechanism can be called upon to stretch performance. These other mechanisms constitute the capacity for graceful extensibility. Thus, all control systems need to be able to recognize when risk of saturation is growing and then deploy and mobilize new response capabilities. These processes of anticipation and building up a readiness to respond are what we observe in studying settings such as emergency medicine, military operations and space mission operations.

In these settings we see some human actors and teams are able to anticipate bottlenecks ahead and adapt activities to generate the capacity to handle that bottleneck, or challenge, should it arise. To fail to anticipate and prepare for the bottleneck ends up putting the team in a situation where they have to generate the means to respond in the middle of the

challenge event – greatly increasing the risk of failure to keep up with the pace and tempo of events. This is the decompensation failure mode for adaptive systems, the risk of being slow and stale to respond to the pace of events, and I have already covered how initiative is essential for systems that anticipate rather than decompensate.

Thus, resilience is not about how well a system adapts to meets its targets and constraints, and not about how well such a system regulates its processes. Instead, however an adaptive system regulates processes and however well it meets targets and constraints, (1) its capacity to do these things is bounded and (2) the environment will present events that fall outside its bounds. The system in question has to have some ability to continue to function when this happens; if not, the system is too brittle and vulnerable to sudden collapse, so that long-term viability declines. The risk of saturation is existential. The ability of a system to continue to respond near or beyond the boundaries is its capacity for graceful extensibility. Viability, over time, requires extensibility and this constraint exists regardless of pressures in increased optimality on some specific criteria. Viability of an entity in the long run requires the ability to stretch – to gracefully move its capabilities as change, challenge and surprise reoccurs.

How does this connect to reciprocity? To see the connection, we look at the hospital emergency department, or ER, as a well-studied exemplar of graceful extensibility (Miller and Xiao 2007; Wears et al. 2008; Perry and Wears 2012; Stephens et al. 2015; Patterson and Wears 2015). All systems have some built-in capacity matched to handle regularities and variations and variation of variations. This defines their base adaptive capacity or competence envelope – what they can handle without risk of saturation. Some systems are designed to be able to handle a range of changing demands. For example, ERs are able to adjust their resources to handle a range of patient problems and varying numbers of patients across every shift as well as over longer rhythms of human activity. Each ER still has finite ability to handle surges in patient load. Situations occur that challenge the boundaries even of systems, such as the ER, that are equipped to handle changing loads. Then the issue is what happens when situations challenge boundaries – when risk of saturation is high? Emergency departments regularly have experience with situations that challenge their ability to respond, and they recognize increasing risk of saturation and adapt in a variety of ways. However, failures can and do occur where triage goes astray, patients are mis-prioritized or under-monitored, and patient condition deteriorates faster than ER staff can recognize or respond.

The ER as an adaptive unit demonstrates graceful extensibility, but studies also illustrate that there are limits to how much graceful extensibility an ER can exhibit when challenged. As an ER approaches saturation, extensibility requires help from other actors in other parts of the health-care system, otherwise the risk of under-monitoring and under-treating critically ill patients grows too large. When the ER risks saturation, continued graceful extensibility requires other neighboring units in the health-care network to recognize this risk and adjust their behavior to help. Chuang et al. (2018) studied adaptations to handle a large mass casualty event, and the results highlighted the critical role of reciprocity and initiative in resilient performance. Stephens et al. (2015) studied a problem in ERs that arises when reciprocity across parts of the hospital fails – working at cross-purposes.

Any unit can be in either position, depending on events and context – the unit at risk of saturation in need of assistance from neighbors, or the neighbor with the potential to assist another at risk of saturation. In networks with high graceful extensibility, adaptive

units demonstrate reciprocity – each unit can anticipate assistance from neighbors should the risk of saturation loom larger, even though that assistance requires adaptations by the assisting units – adaptations that increase the costs and risks experienced by those assisting units. This form of reciprocity is needed to ensure synchronization across units in times of stress. Coordination across units depends on how nearby adaptive units respond when another unit experiences increasing risk of saturating its ability to continue to respond to meet demands. Will the neighboring units adapt in ways that extend the response capability of the unit at risk, or will the neighboring units behave in ways that further constrict the response capability of the unit at risk?

### Complementarity across Engineering, Biological and Social Science Perspectives

Understanding the factors that govern the expression of initiative and reciprocity is a line of work for building resilience that derives from the social sciences. Engineering approaches need to incorporate these social science results in their work. Currently, there is no explicit means to incorporate either initiative or reciprocity in layered architectures for physical systems, for example, Doyle and Csete (2011). Also, when engineers such as Doyle look at social science results they find some after-the-fact explanatory power, but little traction on how to design architectures beyond general soft guidance (Ostrom 2012).

However, biological systems have to address both initiative and reciprocity, though using other terms (for example, Meyers and Bull, 2002). As some commentators on evolution have put it, ‘Life did not take over the globe by combat, but by *networking*’ (Margulis and Sagan 1986, p.15, emphasis added). In the evolution of complexity in biological systems, models have to address processes of goal alignment and sacrifice in a network as conflicts arise or intensify (Michod 2003; ‘conflict mediation was a crucial aspect of the emergence of complexity’ Blackstone 2016, p.3). Ironically, Doyle and colleagues have produced elegant mathematical accounts of systems of biological complexity (Chandra et al. 2011; Li et al. 2014).

Reciprocity points to a space where the gulf across engineering and social science perspectives looms large. I have found myself often moderating a mixed group of engineers and social scientists trying to integrate findings and concepts in order to characterize what makes systems resilient. The discussions inevitably come to an impasse where social scientists protest the engineers are not taking into account critical findings, and where the engineers counter that those findings are too vague and not actionable to design or modify real systems. In one setting, I started to talk about fundamental technical results in human and cognitive systems when I realized the engineers in the room were all staring at me confused. I asked why the puzzled looks, and they responded, ‘technical means engineering’. At which point I stared back equally puzzled, ‘What do you mean? Of course, social and cognitive sciences have technical results, empirical findings and lawful generalizations that can, should and have guided design’.

The synthesis of the social science findings about reciprocity and the control engineering constraints on the risk of saturation, explained above, provides an example of a way forward. In my own work on fundamentals that produce resilient performance across settings, technology and scales, all of the insights come from combining engineering, biological, cognitive and social science perspectives (Woods and Branlat 2011; Woods 2015a). The synthesis respects basic technical results in social science, while creating actionable

directions to design new architectures or polycentric governance mechanisms (for example, Farjadian et al. 2018). These types of syntheses expand the set of tractable but real settings that can be addressed (Seager et al. 2017). Syntheses such as the examples above provide new ways to model, measure and design adaptive networks that are beginning to show both scientific and practical results. The centrality of initiative and reciprocity in resilient performance of complex human systems reveals basic mechanisms that determine whether systems, of all types and at all scales, steer toward resilience or toward brittleness.

## REFERENCES

- Alderson, D.L. and J.C. Doyle (2010), 'Contrasting views of complexity and their implications for network-centric infrastructures', *IEEE SMC – Part A*, **40** (4), 839–52.
- Blackstone, N.W. (2016), 'An evolutionary framework for understanding the origin of eukaryotes', *Biology*, **5** (2), 18, doi:10.3390/biology5020018.
- Chandra, F., G. Buzi and J.C. Doyle (2011), 'Glycolytic oscillations and limits on robust efficiency', *Science*, **333** (6039), 187–92.
- Chen, Y.-Z., Z.-G. Huang, H.-F. Zhang, D. Eisenberg, T.P. Seager and Y.-C. Lai (2015), 'Extreme events in multi-layer, interdependent complex networks and control', *Scientific Reports*, **5**, art. 17277, doi:10.1038/srep17277.
- Chuang, S., D.D. Woods, D. Ting, R.I. Cook and J.-C. Hsu (2018), 'Emergency department response to mass burn casualties of the Formosa Color Dust Explosion in a public hospital', submitted to *Academic Emergency Medicine*.
- Cook R.I. and C. Nemeth (2006), 'Taking things in one's stride: cognitive features of two resilient performances', in E. Hollnagel, D.D. Woods and N. Leveson (eds), *Resilience Engineering: Concepts and Precepts*, Aldershot: Ashgate, pp.205–20.
- Csete, M.E. and J.C. Doyle (2002), 'Reverse engineering of biological complexity', *Science*, **295** (5560), 1664–9.
- Deary, D.S., K.E. Walker and D.D. Woods (2013), 'Resilience in the face of a superstorm: a transportation firm confronts hurricane Sandy', *Proceedings of the Human Factors and Ergonomics Society 57th Annual Meeting*, Santa Monica, CA: Human Factors and Ergonomics Society, pp.329–33, doi.org/10.1177/1541931213571072.
- Doyle, J.C. and M.E. Csete (2011), 'Architecture, constraints, and behavior', *Proceedings of the National Academy of Science USA*, **108** (suppl. 3), S15624–30.
- Doyle, J.C., D.L. Alderson, L. Li, S. Low, M. Roughan, S. Shalunov, et al. (2005), 'The "robust yet fragile" nature of the Internet', *Proceedings of the National Academy of Sciences USA*, **102** (41), 14497–502.
- Fairhall, A.L., G.D. Lewen, W. Bialek, R.R. de Ruyter Van Steveninck (2001), 'Efficiency and ambiguity in an adaptive neural code', *Nature*, **412** (6849), 787–92.
- Farjadian, A.B., A.M. Annaswamy and D.D. Woods (2018), 'A shared pilot-autopilot control architecture for resilient flight', *IEEE Transactions on Control Systems Technology*, submitted.
- Finkel, M. (2011), *On Flexibility: Recovery from Technological and Doctrinal Surprise on the Battlefield*, Palo Alto, CA: Stanford Security Studies.
- Hoffman, R.R. and D.D. Woods (2011), 'Beyond Simon's slice: five fundamental trade-offs that bound the performance of human macrocognitive work systems', *IEEE Intelligent Systems*, **26** (6), 67–71.
- Hollnagel, E. (2009), *The ETTO Principle: Efficiency-Thoroughness Trade-Off: Why Things that Go Right Sometimes Go Wrong*, Farnham: Ashgate.
- Hollnagel, E., D.D. Woods and N. Leveson (eds) (2006), *Resilience Engineering: Concepts and Precepts*, Aldershot: Ashgate.
- Klein, G., P. Feltovich, J.M. Bradshaw and D.D. Woods (2005), 'Common ground and coordination in joint activity', in W. Rouse and K. Boff (ed.), *Organizational Simulation*, Hoboken, NJ: Wiley, pp.139–78.
- Li, N., J. Cruz, S.C. Chenghao, S. Somayeh, B. Recht, D. Stone, et al. (2014), 'Robust efficiency and actuator saturation explain healthy heart rate control and variability', *Proceedings of the National Academy of Science*, **111** (33), E3476–85, accessed 23 July 2018 at <http://www.pnas.org/content/111/33/E3476>.
- Margulis, L. and D. Sagan (1986), *Microcosmos*, New York: Summit.
- Meyers, L.A. and J.J. Bull (2002), 'Fighting change with change: adaptive variation in an uncertain world', *Trends in Ecology & Evolution*, **17** (12), 551–7.
- Michod, R.E. (2003), 'Cooperation and conflict mediation in the evolution of multicellularity', in P. Hammerstein (ed.), *Genetic and Cultural Evolution of Cooperation*, Cambridge, MA: MIT Press.
- Miller, A. and Y. Xiao (2007), 'Multi-level strategies to achieve resilience for an organization operating at capacity: a case study at a trauma centre', *Cognition Technology and Work*, **9** (2), 51–66.

- Ostrom, E. (1998), 'Scales, polycentricity, and incentives: designing complexity to govern complexity', in L.D. Guruswamy and J. McNeely (eds), *Protection of Global Biodiversity: Converging Strategies*, Durham, NC: Duke University Press, pp. 149–67.
- Ostrom, E. (1999), 'Coping with tragedies of the commons', *Annual Reviews in Political Science*, **2** (1), 493–535.
- Ostrom, E. (2003), 'Toward a behavioral theory linking trust, reciprocity, and reputation', in E. Ostrom and J. Walker (eds), *Trust and Reciprocity: Interdisciplinary Lessons from Experimental Research*, New York: Russell Sage Foundation.
- Ostrom, E. (2012), 'Polycentric systems: multilevel governance involving a diversity of organizations', in E. Brousseau, T. Dedeurwaerdere, P.-A. Jouvét and M. Willinger (eds), *Global Environmental Commons: Analytical and Political Challenges in Building Governance Mechanisms*, Oxford: Oxford University Press, pp. 105–25.
- Ostrom, E. and J. Walker (eds) (2003), *Trust and Reciprocity: Interdisciplinary Lessons from Experimental Research*, New York: Russell Sage Foundation.
- Patterson, M.D. and R.L. Wears (2015), 'Resilience and precarious success', *Reliability Engineering and System Safety*, **141** (September), 45–53.
- Perry, S. and R. Wears (2012), 'Underground adaptations: cases from health care', *Cognition Technology and Work*, **14** (3), 253–60, doi.org/10.1007/s10111-011-0207-2.
- Roth, A.E. (2008), 'What have we learned from market design?', *Economic Journal*, **118** (527), 285–310.
- Seager, T.P., S.S. Clark, D.A. Eisenberg, J.E. Thomas, M.M. Hinrichs, R. Kofron, et al. (2017), 'Redesigning resilient infrastructure research', in I. Linkov and J. Palma Oliveira (eds), *Resilience and Risk: Methods and Application in Environment, Cyber and Social Domains*, Dordrecht: Springer, pp. 81–119.
- Shattuck, L.G. and D.D. Woods (2000), 'Communication of Intent in Military Command and Control Systems', in C. McCann and R. Pigeau (eds), *The Human in Command: Exploring the Modern Military Experience*, New York: Kluwer Academic and Plenum, pp. 279–92.
- Stephens, R.J., D.D. Woods and E.S. Patterson (2015), 'Patient boarding in the emergency department as a symptom of complexity-induced risks', in R.L. Wears, E. Hollnagel and J. Braithwaite (eds), *Resilience in Everyday Clinical Work*, Farnham: Ashgate, pp. 129–44.
- Sutcliffe, K.M. and T.J. Vogus (2003), 'Organizing for resilience', in K.S. Cameron, I.E. Dutton and R.E. Quinn (eds), *Positive Organizational Scholarship*, San Francisco, CA: Berrett-Koehler, pp. 94–110.
- Walker, B.H. and D. Salt (2006), *Resilience Thinking: Sustaining Ecosystems and People in a Changing World*, Washington, DC: Island Press.
- Wark, B., B.N. Lundstrom and A. Fairhall (2007), 'Sensory adaptation', *Current Opinion in Neurobiology*, **17** (4), 423–9.
- Watts-Perotti, J. and D.D. Woods (2009), 'Cooperative advocacy: a strategy for integrating diverse perspectives in anomaly response', *Computer Supported Cooperative Work: The Journal of Collaborative Computing*, **18** (2), 175–98.
- Wears, R.L., S.J. Perry, S. Anders and D.D. Woods (2008), 'Resilience in the emergency department', in E. Hollnagel, C. Nemeth and S.W.A. Dekker (eds), *Resilience Engineering Perspectives 1: Remaining Sensitive to the Possibility of Failure*, Aldershot: Ashgate, pp. 193–209.
- Woods, D.D. (2003), 'Creating foresight: how resilience engineering can transform NASA's approach to risky decision making. Testimony on the future of NASA to Senate Committee on Commerce, Science and Transportation, John McCain, Chair, Washington D.C., October 29, 2003', accessed 24 April 2017 at <http://cseel.org.ohio-state.edu/blog/dave-before-congress/>
- Woods, D.D. (2005a), 'Creating foresight: lessons for resilience from Columbia', in W.H. Starbuck and M. Farjoun (eds), *Organization at the Limit: NASA and the Columbia Disaster*, Malden, MA: Blackwell, pp. 289–308.
- Woods, D.D. (2005b), 'Conflicts between learning and accountability in patient safety', *DePaul Law Review*, **54** (2), 485–502.
- Woods, D.D. (2006), 'Essential characteristics of resilience for organizations', in E. Hollnagel, D.D. Woods and N. Leveson (eds), *Resilience Engineering: Concepts and Precepts*, Aldershot: Ashgate, pp. 21–34.
- Woods, D.D. (2009), 'Escaping failures of foresight', *Safety Science*, **47** (4), 498–501.
- Woods, D.D. (2015), 'Four concepts of resilience and the implications for resilience engineering', *Reliability Engineering and System Safety*, **141** (April), 5–9, doi:10.1016/j.res.2015.03.018.
- Woods, D.D. (2018), 'The theory of graceful extensibility', *Environment Systems and Decisions*, **38**, doi:10.1007/s10669-018-9708-3.
- Woods, D.D. and M. Branlat (2011), 'How adaptive systems fail', in E. Hollnagel, J. Paries, D.D. Woods and J. Wrathall (eds), *Resilience Engineering in Practice*, Aldershot: Ashgate, pp. 127–43.
- Woods, D.D. and L.G. Shattuck (2000), 'Distant supervision – local action given the potential for surprise', *Cognition, Technology and Work*, **2** (4), 242–5.