Analyzing System Resilience to Adversary Kinetic and Cyber Actions

Clayton A. Smith

ABSTRACT

Quantifying the resilience of an engineered system is a risk proposition. Definitions of resilience encompass notions of loss of capability and the time required to restore it. The loss and restoration of capability are responses to off-nominal conditions (be they random events or actions initiated by an adversary), and these conditions are probabilistic in nature. The Johns Hopkins University Applied Physics Laboratory (APL) developed a framework that quantifies system resilience as a risk profile by identifying the set of threats and analyzing the potential system responses. Although the framework was applied to space architectures, it is agnostic to the specific domain and therefore can be implemented in other areas where resilience is needed at the enterprise level. The analysis of space architecture resilience led to understanding threats (kinetic, RF, cyber, or random failures) to the systems. Mitigations to these threats were prioritized using quantitative resilience metrics that evaluate the potential performance impacts and time to recover and overcome these impacts.

INTRODUCTION

> As we invest in next generation space capabilities and fill gaps in current capabilities, we will include resilience as a key criterion in evaluating alternative architectures.
>
> —*National Security Space Strategy* (2011)

For much of the Cold War, U.S. space systems focused primarily on supporting the strategic missions, including missile warning, technical intelligence, and nuclear command and control, which enabled the strategic détente between the United States and the Soviet Union. Since the end of the Cold War, the space domain has become increasingly crowded and contested. More than 60 nations now own or operate satellites, and virtually all nations depend on space-based capabilities for civilian applications such as weather forecasting and navigation. An implicit assumption in the space domain was that strategic deterrence would prevent space systems from being attacked during conventional conflicts.

The 1991 Gulf War marked a substantial shift in the way the U.S. military uses space systems. This conflict demonstrated the value of fusing space-based capabilities, such as precision navigation and timing and satellite communications, with conventional weapon systems. Potential adversaries are not as reliant on space-based capabilities and do not have symmetric vulnerabilities, making traditional deterrence in space a difficult proposition. Moreover, the U.S. military's critical dependence on space-based capabilities for global power projection means that counter-space capabilities may figure prominently in an adversary's anti-access/area denial operations.

The community's response has been to attempt to address the resilience of the entire system. Indeed, as stated in the 2011 *National Security Space Strategy* "Our military . . . must be prepared to 'fight through' a degraded environment and defeat attacks targeted at our space systems and supporting infrastructure." General C. Robert Kehler, as commander of U.S. Strategic Command, more specifically stated, "Beyond awareness in space we need robust, resilient architectures—both space-based constellations and terrestrial assets—to ensure today's essential space-based services are available to accomplish the mission."[1] General William L. Shelton, Commander of Air Force Space Command from 2011 to 2014, said "resilience in the face of . . . growing space threats is an imperative. If space assets come under attack, either as a precursor to conflict or as an integral part of terrestrial hostilities, our architectures must be resilient enough to assure mission accomplishment."[2] Understanding and improving the resilience of our nation's space assets is a military priority.

Space has changed from a relatively safe operational domain to one that is increasingly congested with space debris, increasingly contested by a growing range of foreign counter-space capabilities, and increasingly competitive as more entities operate in it. U.S. national strategy relies on maintaining and enhancing our space-derived advantages while confronting the challenges of an evolving space environment. Improving resilience is one of the ways to address this challenge. Indeed, the 2012 Space Policy Directive (DoDD 3100.10) states that "reliability, protection, and resilience of required space capabilities, including information systems and networks and other infrastructure required to support sustained operations, will be considered in all architecture planning and evaluation."[3]

Currently, there is no agreed-upon quantitative method for measuring resilience. As such, most discussions revolve around methods to improve resilience in the abstract through disaggregation, hosted payloads, hardening, or avoidance. In reality, a combination of policy; architecture; orbit; and tactics, techniques, and procedures may yield the best results within the constraints of declining budgets.[4] A quantified metric allows for the trade space exploration encompassing all these attributes.

The Johns Hopkins University Applied Physics Laboratory (APL) developed a framework that quantifies resilience by establishing essential characteristics and identifying a constructive model that can incorporate existing analysis and simulation efforts being used within the space community. The framework described herein is an extension of the work originally started as a study to quantify project risks as a multi-attribute problem[5] and later expanded to quantify resilience using potential threats instead of risk for the Space Security and Defense Program (SSDP). As a result, the methodology was created for understanding and quantifying resilience in a context appropriate to SSDP. It measures resilience of on-orbit assets, up-links/down-links/cross-links, and ground infrastructure against a series of known and emerging threats. Under a contract with the Office of the Under Secretary of Defense for Acquisition, Technology, and Logistics for the Research, Development Agency Task Force, the framework was later expanded to include cyberattacks. Although the framework was applied to space architectures, it is agnostic to the specific domain and therefore can be implemented in other areas where resilience is needed at the enterprise level.

RESILIENCE

The basis of the methodology uses the key ideas of resilience as described in the Office of the Secretary of Defense fact sheet entitled *Resilience of Space Capabilities*[6]:

> The purpose of resilience is to assure performance of military and related intelligence functions at a level necessary to execute assigned missions within an acceptable tolerance for risk. This functional mission assurance must account for the full range of anticipated scenarios, conditions, and threats that drive our planning.

The space community has coalesced around four fundamental objectives of resilient systems[6]:

1. **Avoidance:** Reduce the likelihood and consequence of adverse conditions or hostile actions.

2. **Robustness:** Resist capability degradation when faced by adverse conditions or hostile action.

3. **Reconstitution:** Replenish lost or diminished capabilities to an acceptable level, for a particular purpose, subsequent to an adverse condition or hostile action.

4. **Recovery:** Reestablish full operational capability subsequent to reconstitution.

Resilience at its core is a risk proposition. Resilience is the ability of a system to continue providing required capabilities in the face of system failures, environmental challenges, or adversary action. Risk is the measure of future uncertainties in achieving performance goals and objectives within defined cost, schedule, and performance constraints.[7] As such, risk is the measure for resilience once the threats and mission objective are defined.

When assessing risk, four questions are asked: (*i*) What can happen? (*ii*) How likely is it to happen? (*iii*) If it does happen, what are the consequences? (*iv*) How confident are we in the results? The answers define a set of possible outcomes, probabilities (with uncertainty), and consequences (measures of damage) for each scenario.[8] Probability plays a key role in the quantification of resilience. It is a measure encompassing the concepts of both relative frequency and degree of belief. The latter is the degree of support a body of evidence gives to a hypothesis about the occurrence of an event. Resilience needs to be assessed in conjunction with architecture trade studies to reduce the probability of potentially mission-ending consequences.

When decisions are complex, information is uncertain, impacts are ambiguous, and consequences have multiple attributes, a formal structured systematic assessment is useful for leadership. Certain characteristics of state-of-the-art risk assessments are now standard,[9–17] but they are not reflected in current frameworks for resilience. The APL-developed framework is built on the following characteristics essential to evaluating system resilience:

- **Scenario based:** Assessments that are scenario based contain model logic tracing events that perturb the nominal functioning of a system and examine how the system responds. Multiple potential outcomes are addressed. These scenarios often include potential mitigation alternatives that enable evaluation of mitigation effectiveness. This characteristic allows analysts to "think" through situations before they occur so that they can prepare for contingencies, and it provides a platform to quantify consequences and likelihoods.

- **Integrated:** Analyses address a wide scope of concerns. Integrated frameworks address the entire system/architecture/mission, combined effects of multiple threats, dependencies among and within threat scenarios, and interfaces with other existing models. A systematic integrated framework enables trade-off studies and sensitivity analyses across multiple scenarios and constituent system elements.

- **Quantitative:** Model-based assessments compute measures representing risk and resilience. Emphasis is on expressing the likelihood of observable performance measures or events. Probabilities are assigned based on systematic proven processes for data analysis and probability calculus.

- **Probabilistic:** Any model involves assumptions, simplifications, and data variability. It is essential that decision-makers receive information about all three in the form of uncertainties. Both aleatory (randomness due to inherent variability in the system) and epistemic (imprecision due to lack of knowledge and information on the system of models) uncertainties are included as integral parts of the assessments.

- **Multi-attribute:** Decision-makers struggle with multiple measures of effectiveness that must be unified. They rarely make decisions on single values alone, but rather on an integrated view of the world. Assessments account for decision-makers' risk aversion levels and preferences for those attributes.

- **Actionable:** The framework's scope and level of detail must be coincident with decision-makers' ability to take action. This focuses on actions that can be implemented in the design, operations, or acquisition cycles.

- **Probative:** The ability to rank order important drivers and perform sensitivity analysis enables insights about current situations and future alternatives. Importance measures provide a quantitative view of model elements.

In addition, this resilience framework needs to incorporate the many dimensions of resilience when assessing various architectures:

- Anticipated ability of adversary
- Functional capability goals necessary to support the mission
- Probability that these goals may not be met at a given level of adversity
- Severity of the functional shortfall to the mission
- Period of time for which the shortfall can be tolerated by the mission
- Uncertainty about various parameters in the model

The temporal component is particularly important, since time primarily quantifies the reconstitution component of resilience.

Pedagogical literature shows a number of different attempts to develop and quantify metrics for resilience.[18–21] The cybersecurity community is actively using this concept and language as it implements resilience into networks and other interdependent systems and systems of systems.[22] A body of literature discusses the benefits of these techniques in many industries.

RESILIENCE QUANTIFICATION FRAMEWORK

Measuring the efficacy of a system's resilience can be achieved, for example, through the unique functionality of that system and its responses (outputs) to specific inputs. Given that such inputs are probabilistic, so are the outputs, meaning that the system's resilience—because it is measured in terms of responses to the inputs—can be measured (quantified) in probabilistic terms and for specific inputs. We can thus adduce the following premises for scoring resilience: (*i*) The probabilistic nature of threats and thus their associated outputs necessitates a holistic, multidimensional probabilistic scoring system of resilience. Furthermore, the myriad plausible threat sce-

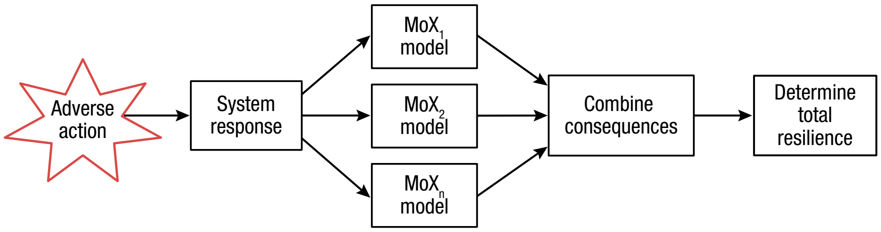

Figure 1. Adverse action propagated through system models. System responses to adverse actions are represented by multiple MoXs to determine resilience.

narios, each with associated magnitudes and durations, necessarily limit any such scoring system of resilience to specific classes of input threats. (*ii*) Resilience, as a vector of the states of all physical and natural systems, is time variant; given the inherent characteristics of such systems, their resilience will deteriorate over time. Thus, even an input-limited scope of any scoring will be further constrained by the inherent time variant resilience of the system.

To properly understand and improve the resilience of a system, it must be measured, which means a metric must be developed. Resilience is tied to the various changes to measures of interest (MoXs, which is shorthand for measures of operation, measures of performance, or measures of effectiveness dependent on the scope and level of detail for the analysis) and the time it takes to bring them back to required levels. From a modeling perspective (see Fig. 1), an adverse action elicits a response from the system. The model then collects and combines the responses to develop an impact on the system as a whole in order to determine the overall resilience. For simplicity in describing our basic concept, the discussion is initially limited to a single MoX.

Adverse Actions

A list of threats to the system is developed using intelligence sources and systematic methods from reliability engineering to identify the vulnerabilities and how they can be exploited. For each adverse action, be it a kinetic attack or a cyberattack, a naturally occurring change in the environment, or a random failure, the system will respond in a variety of ways.

System Response

To model the system response, scenarios are developed. Scenarios are stories about postulated series of events constructed to focus attention on causal processes and decision points. Scenarios are coherent descriptions of alternative images of the future, created from mental maps reflecting various perspectives on past, present, and future developments. They are used within resilience assessments to broaden views and raise questions about conventional success-oriented thinking.

Modeling scenarios begins with a description of the success sequence. This represents a sequence of events executed by the system. At each point in this sequence, we can ask what can go wrong. The answer to that question is termed an adverse action. Given its occurrence, and, depending on what happens next, a set of paths emerges and terminates at an end state. Many engineering systems include safety or backup systems meant to be activated in response to the various events. If backups work as intended, consequences are typically insignificant. However, if the event occurs and corresponding backup systems fail, there could be serious consequences. Probabilities are assigned to every event in the scenarios, allowing for a rollup of probability to be computed at the end states.

Scenarios are useful tools in articulating key considerations, assumptions, and constraints. They provide a platform to blend qualitative and quantitative knowledge of systems and their interactions. Analysts still need to be cautious to avoid common traps such as narrowly examining a situation, applying assumptions inconsistently, or not fully documenting assumptions, thereby reducing transparency or overly constraining the problem space. One thing to note about scenarios in this context is that they are meant to describe a class of situations that can occur. They are not meant to explicitly describe every possible permutation of events, an infinite set of permutations.

MoX Models

At the end state of each scenario, the MoX versus time is modeled. Consider a system operating in a nominal state (M_0) as shown in Fig. 2. Here, the performance (measured by some MoX) is at the required value. Suppose

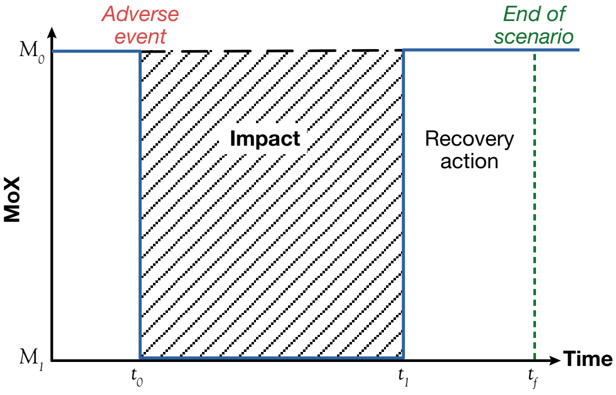

Figure 2. Basic resilience schematic of a scenario. Resilience quantifies to change in system impact over time.

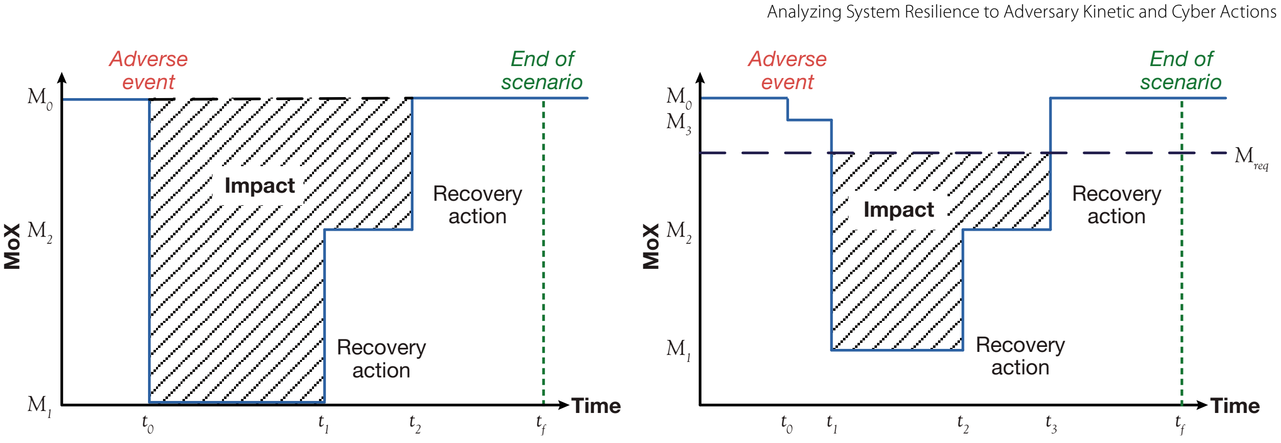

Figure 3. More complex view of resilience. Change in impact varies depending of the event and recovery scenario.

that at time t_0 an adverse event brings down the system to state (M_1), providing no capability, and then at some later time, t_1, the system is restored to full capability. The total impact to the system is characterized by the shaded area, Impact = Δ MoX × Δ time.

The construct agrees in principle with those discussed in Refs. 20 and 23–27. Resilience is now defined as the inverse of the impact × the probability of encountering the state:

$$\text{Consequence} = \frac{1}{\text{Impact}} = \frac{\%\text{Capability}}{\text{Time to recover}}.$$

This expression correctly relates capability to resilience (the higher the capability retained, the more resilient the system). Likewise, resilience is higher when the recovery time is shorter. As this concept expands, the probabilistic nature is reflected in the uncertainties related to the probability of being in the state, the nondeterministic degradation in capability due to the attack, and the variability of the recovery time. The make-up of impact can also look different for various scenarios. For example, Fig. 3 shows a two-step recovery process on the left. The situation on the right is more complicated. First, the system's full capability is greater than the required level, changing the impact computation with respect to M_{req}. Next, the depth of capability lost is not total, indicating some level of degraded operation (whether that degraded capability has value is a question later addressed with utility functions). Also, this schematic shows an initial reduction in capability that is still more than required and therefore does not have an impact.

Resilience metrics defined this way require knowledge of several aspects of the scenario and system/architecture/mission. First, an understanding of how the system functions in nominal operations is needed so that a baseline can be determined. Once this understanding is established, the system responses to various adverse conditions can be understood as deviations from nominal operations as measured by the many MoXs. Because of the potentially large number of MoXs, a small subset, preferably one or two, needs to be identified as a surrogate for the total set. The system responses will vary based on the nature of the attack, and therefore an understanding of the threat space is also needed.

The framework relies on assigning probabilities to the events in the scenario. Evidence on which to base these probability assignments will come from one or more of the following:

- Observations
- Intelligence assessments
- Experience
- Results of modeling and simulation tools
- Special investigations/studies
- Subject-matter expert opinion

The scope, level of detail, and level of decomposition determines the level of credibility (uncertainty) in the results. While this sounds like an obvious statement, it carries a powerful feature of the probabilistic nature of the framework; since the uncertainty is quantified and shown to the decision-maker, the resources needed to perform an analysis are directly related to the level of acceptable risk. Lower-fidelity input knowledge can still be executed through the framework with a corresponding reduction in output fidelity.

Consequence Analysis

As discussed earlier, the consequences are the result of analyzing a sequence (path through the scenario) to

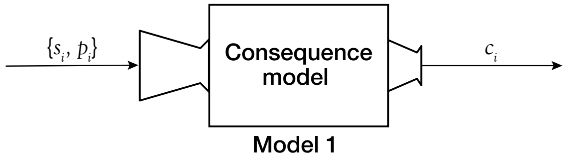

Figure 4. Consequence outputs from scenarios. System consequence is a function the scenario and its probability of occurrence.

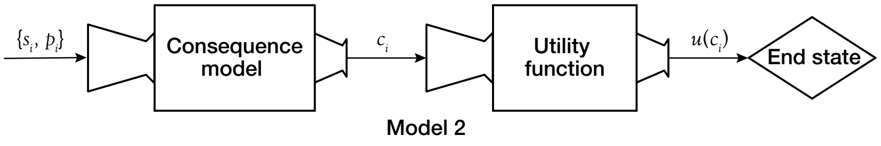

Figure 5. Utility value output from scenarios and decisions. End state consequence is a function of the decision-maker's risk preference.

determine the impact. Figure 4 shows the set of probabilities and sequences as inputs with the consequence value, c_i, as the output. Analysts must work closely with those in other disciplines to fully populate the resilience model. This is an intentional feature in that it forces communication and understanding of potential end-state contexts.

End states map the consequences from the impact graphs. Each MoX will have its own graph and nuance. They identify attribute data, so that when one of the attributes or utility functions changes, a new value can be recomputed. Current techniques often force analysts to choose to focus on only the worst-case attribute. With this implementation, all attribute information is kept and analyzed. Should the leadership's utility function focus on a particular attribute, so be it, but it is a traceable and defendable aspect of the entire model.

After computing the consequences, end states are "gathered," meaning that probabilities for identical end states are combined. Since the paths within a scenario are mutually exclusive, the probabilities are summed.

Using only raw consequence values as the basis for a resilience metric would be convenient and straightforward. However, it would be equivalent to computing an average consequence. This can lead to intuitively unpalatable decision recommendations because it does not account for the decision-maker's tolerance for risk or imbedded preferences. Decision analysts solve this problem by using expected utility theory. The application of this concept is succinctly characterized in Ref. 28: "If an appropriate utility is assigned to each possible consequence and the expected utility of each alternative is calculated, then the best course of action is the alternative with the highest expected utility." The next step in the end-state evaluation process incorporates utility functions. Since scenario models conform to the assumptions and constraints specified within the field of decision analysis, utility theory can be used as a driving analytic technique for examining decisions with uncertainty.[28–33]

Utility functions translate consequences (percent degradation and recovery time) to a unit-less number typically scaled so that the least preferred level equates to 0 and the most preferred is 1. It is this value that we use to quantify end states (see Fig. 5) so that profiles and metrics can be computed. For instance, there may be utility only if the MoX is above a threshold with none below, and therefore the utility function would reflect this binary view of the consequence.

One point to note is that utility is independent of probability and belongs to an individual decision-maker or analysts. Changing the decision-maker or his/her perspective (i.e., the utility function) can change results significantly. The practical implication here is that we can change the framework based on who is using it (project manager, acquisition authority, mission commander), but not the underlying data.

Resilience Quantification

Once quantified, scenarios provide probability and analysis of consequences for a given adverse action. A figure of merit is needed to distill all this information for use as a basis of comparison. With such a metric, it is possible to gain insights for a variety of aspects of project risk, including whether a mitigation is likely to increase resilience and by how much, or to derive a list of elements affecting resilience the most. Much like Kaplan and Garrick[8] championed the use of Farmer curves (consequence versus frequency), we create resilience profiles in the form of exceedance probability curves as a way to describe the resilience. An exceedance probability curve specifies the probabilities that a certain level of loss will be exceeded. In our case, loss is utility. These curves are referred to as resilience profiles.

Profiles are created starting with a set of ordered pairs of probability and utility values from a scenario. Once the consequences for each path through a scenario are computed, they are sorted in increasing order and plotted against the cumulative probability associated with each state. This construct is called a complementary cumulative distribution function or survival function and is useful to study how often the MoX is above a particular level. Figure 6 illustrates this profile. The resilience is now defined as the expectation (ℝ) of

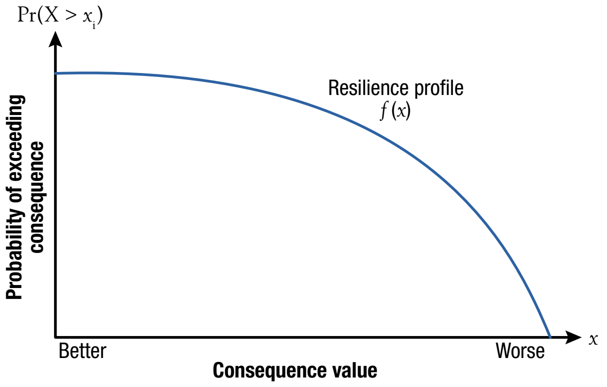

Figure 6. Resilience profile curve characterization. The resilience profile provides more information than a point estimate.

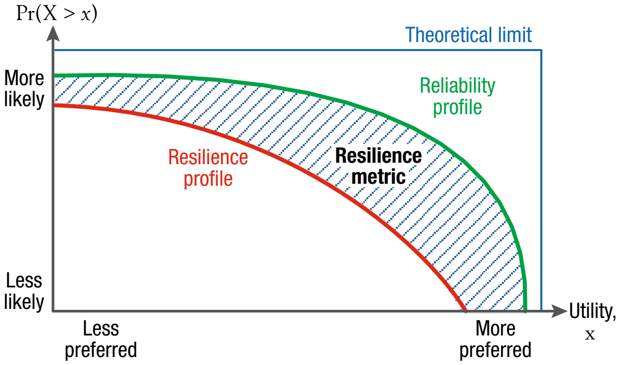

Figure 7. Resilience metric with respect to a reliability baseline. Resilience metric is the difference between reliability and resilience at each level of system consequence.

the resilience profile. There will be a resilience profile and ℝ for each threat or adverse action. Since there can be a multitude of attack types, there will be many ℝ values.

Additionally the area above the curve can be computed. A reference is needed for this value to have any meaning. One such reference could be to find the area above the curve and below a reliability curve or the theoretical limit, as shown by the blue line in Fig. 7. A better choice is to reference the area to the profile curve of the system reliability. Note that this line (green) is the best achievable due to random occurrences of failures and environmental variability. In other words, it is the how the system behaves in its nominal condition without the presence of an attack.

The end result, given all the threats, scenarios, and end states, is a vector of ℝ values, one for each threat. This is consistent with the notion that resilience is only meaningful with respect to a specific perturbation to the system's nominal operations.

Quantitative analyses of the phenomena occurring in many engineering applications are based on mathematical models that are then turned into operative computer codes for simulation. A model provides a representation of a real system dependent on several hypotheses and parameters. The model can be deterministic or stochastic. In practice, the system under analysis cannot be characterized exactly. This leads to uncertainty in both the values of the model parameters and the hypotheses supporting the model structure. Uncertainty is an unavoidable component affecting the behavior of systems and more so with respect to their limits of operation.[34] Despite how much dedicated effort is put into improving the understanding of systems, components, and processes through the collection of representative data, the appropriate characterization, representation, propagation, and interpretation of uncertainty remains a fundamental element of the risk analysis of any system. Following this view, uncertainty analysis is considered an integral part of resilience assessment, although it can also exist independently in the evaluation of unknown quantities.

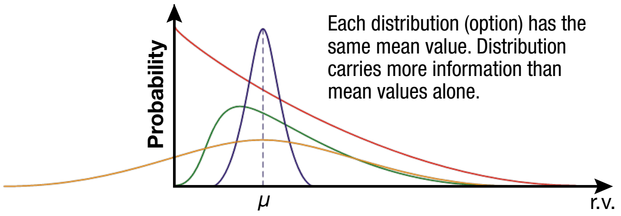

Figure 8. Effect of distributions. r.v., Random variable.

Resilience metrics are estimates and subject to the uncertainties discussed above, and therefore analysts must tell the decision-maker what uncertainties exist and how they affect the results. Probability distributions are the mathematically correct way to communicate to the decision-maker the assumptions and approximations and to give a sense of how reliable the numbers are. It is also easier to justify central tendency to reach consensus on a range and distribution than it is on a point value. Take, for example, the four probability distribution curves in Fig. 8 signifying the resilience of separate alternative architectures. Each has the same mean value. If only the point estimate were provided to a decision-maker, all four options would be identical. However, the uncertainty about them tells a different story both in the amount of spread and shape.

FRAMEWORK BENEFITS

This framework provides metrics for resilience consistent with current definitions and constructs within the space community and addresses the challenges laid out in policy documents. The framework is a cost-effective assessment tool in that it enables an analytical level of effort that is commensurate with the acceptable level of uncertainty and leverages existing modeling and simulation tools. The scenario models provide an intuitive communication vehicle and the flexibility to incorporate space-based, ground, and cyber threats. Since the scenario model and documentation are the same, leadership does not have to accept modeling completely on faith. Adding to the improved communication is that uncertainty is explicitly addressed, providing a level of credibility in the results and making uncertainty part of the overall architecture trade space. Analyses are scalable to spacecraft, architecture, or mission level using the same framework, taxonomy, and processes. Finally, the framework improves traceability and transparency of subject-matter expert evaluations and inputs. The ability to quantify resilience in context with mission allows planners and designers insight into system- and enterprise-level trades that heretofore could not be accomplished.

REFERENCES

[1] Kehler, C. R., "Implementing the National Security Space Strategy," *Strat. Stud. Q.* **6**(1), 18–26 (2012).

[2] Shelton, W. L., "Military Space: At a Strategic Crossroad," *Air Space Power J.* **27**(5), 4–10 (2013).

[3] Office of the Secretary of Defense, *Space Policy*, Directive 3100.10, U.S. Department of Defense, Washington DC (2012).

[4] Meink, T,. "Resilience Deep Dive," in *Proc. NDIA Space Forum on Space Resilience*, Colorado Springs, CO (2013).

[5] Smith, C., "Integrated Scenario-Based Methodology for Project Risk Management," Ph.D. dissertation, University of Maryland, College Park (2011).

[6] Office of the Secretary of Defense - Space Policy, *Fact Sheet: Resilience of Space Capabilities*, U.S. Department of Defense, Washington, DC (2011).

[7] Office of the Under Secretary of Defense for Acquisition Technology and Logistics, *Risk Management Guide for DoD Acquisition*, 6th Ed., U.S. Department of Defense, Washington DC (2006).

[8] Kaplan, S., and Garrick, J., "On the Quantitative Definition of Risk," *Risk Anal.* **1**(1), 11–27 (1981).

[9] Aven, T., *Foundations of Risk Analysis: A Knowledge and Decision-Oriented Perspective*, 1st Ed., Wiley, West Sussex, UK (2003).

[10] Bedford, T., and Cooke, R., *Probabilistic Risk Analysis: Foundations and Methods*, Cambridge University Press, Cambridge, UK, (2001).

[11] Cox, L., *Risk Analysis of Complex and Uncertain Systems*, Springer, New York (2009).

[12] Frank, M., *Choosing Safety: A Guide to Using Probabilistic Risk Assessment and Decision Analysis in Complex, High-Consequence Systems*, RFF Press, Washington, DC (2008).

[13] Henley, E. J., and Kumamoto, H., *Probabilistic Risk Assessment and Management for Engineers and Scientists*, IEEE, New York (1996).

[14] Modarres, M., *Risk Analysis in Engineering: Techniques, Tools, and Trends*, CRC Press, Boca Raton, FL (2006).

[15] NASA, *Probabilistic Risk Assessment Procedures Guide for NASA Managers and Practitioners*, NASA: Washington, DC (2004).

[16] U.S. Nuclear Regulatory Commission, *PRA Procedures Guide: A Guide to the Performance of Probabilistic Risk Assessments for Nuclear Power Plants*, Office of Nuclear Regulatory Research, Washington, DC (1983).

[17] Vose, D., *Risk Analysis: A Quantitative Guide*, 3rd Ed., John Wiley & Sons, West Sussex, UK (2008).

[18] Haimes, Y., "On the Definition of Resilience in Systems," *Risk Anal.* **29**(4), 498–501 (2009).

[19] Berkowitz, M., Kelsey, A. A., Swietek, G., Schierling, J. G., and Williard, L. D., "Space Mission Resilience," in *Proc. AIAA SPACE 2013 Conf. and Exposition*, San Diego, CA, pp. 1–8 (2013).

[20] Brtis, J., *How AFSPC Can Address Resilience*. MITRE Corp., McLean, VA (2013).

[21] Jackson, S., and Ferris, T., "Resilience Principles for Engineered Systems," *Sys. Eng.* **16**(2), 152–164 (2013).

[22] Bodeau, D., Graubart, R. D., Picciotto, J., and McQuaid, R., *Cyber Resiliency Engineering Framework*, MITRE Corp., McLean, VA (2012).

[23] Francis, R., and Bekera, B., "A Metric and Frameworks for Resilience Analysis of Engineered and Infrastructure Systems," *Reliab. Eng. Syst. Safe.* **121**, 90–103 (2014).

[24] Dalziell, E., and McManus, S., "Resilience, Vulnerability, and Adaptive Capacity: Implications for System Performance," in *Proc. International Forum for Engineering Decision Making (IFED)*, Stoos, Switzerland, pp. 1–17 (2004).

[25] Ayyub, B., "Systems Resilience for Multihazard Environments: Definition, Metrics, and Valuation for Decision Making," *Risk Anal.* **34**(2), 340–355 (2013).

[26] Burch, R., "A Method for Calculation of the Resilience of a Space System," in *Proc. IEEE Military Communications Conf. 2013 (MILCOM 2013)*, San Diego, CA, pp. 1002–1007 (2013).

[27] White, L. S., "Evaluating Resiliency," in *Proc. NDIA Space Forum on Space Resilience*, Colorado Springs, CO (2013).

[28] Keeney, R., and Raiffa, H., *Decision with Multiple Objectives: Preference and Value Tradeoffs*, Cambridge University Press, Cambridge, (1993).

[29] Clemen, R., and Reilly, T., *Making Hard Decisions with Decision Tools Suite*, 2nd Ed., Duxbury, Pacific Grove, CA (1999).

[30] Fishburn, P., "Foundations of Decision Analysis: Along the Way," *Manag. Sci.* **35**(4), 387–405 (1989).

[31] Raiffa, H., *Decision Analysis: Introductory Lectures on Choices Under Uncertainty*, Addison-Wesley, Reading, MA (1968).

[32] Kirkwood, C., "An Overview of Methods for Applied Decision Analysis," *Interfaces* **22**(6), 28–39 (1992).

[33] Howard, R., "Decision Analysis: Practice and Promise," *Manag. Sci.* **34**(6), 679–695 (1988).

[34] Aven, T., and Zio, E., "Some Considerations on the Treatment of Uncertainties in Risk Assessment for Practical Decision Making," *Reliab. Eng. Syst. Safe.* **96**(1), 64–74 (2011).

Clayton A. Smith, Space Exploration Sector, Johns Hopkins University Applied Physics Laboratory, Laurel, MD

Clayton Smith is a member of the Principal Professional Staff at APL and has more than 30 years of experience analyzing systems from risk, reliability, and safety perspectives. These systems included NASA and DoD missions, payloads, ground communication systems, air traffic control systems, and missile systems. He is developing approaches to assess intentional threats against space assets using probabilistic risk analysis and game theory techniques. He created and managed NASA's International Space Station Program probabilistic risk assessment specifically geared toward quantifying the safety risk during operations. Clayton is currently the reliability engineering lead for APL's PSP mission. He received a B.S. in aerospace engineering, an M.S. in engineering management, and a Ph.D. in reliability engineering all from the University of Maryland. His e-mail address is clay.smith@jhuapl.edu.