

Quantum Resilience

by

Thomas Wade Roberts

A Dissertation Presented in Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy

Approved April 2015 by the
Graduate Supervisory Committee:

Brad Allenby, Chair
Mikhail Chester
John Anderies

ARIZONA STATE UNIVERSITY

May 2015

ABSTRACT

Quantum resilience is a pragmatic theory that allows systems engineers to formally characterize the resilience of systems. As a generalized theory, it not only clarifies resilience in the literature, but also can be applied to all disciplines and domains of discourse. Operationalizing resilience in this manner permits decision-makers to compare and contrast system deployment options for suitability in a variety of environments and allows for consistent treatment of resilience across domains. Systems engineers, whether planning future infrastructures or managing ecosystems, are increasingly asked to deliver resilient systems. Quantum resilience provides a way forward that allows specific resilience requirements to be specified, validated, and verified.

Quantum resilience makes two very important claims. First, resilience cannot be characterized without recognizing both the system and the valued function it provides. Second, resilience is not about disturbances, insults, threats, or perturbations. To avoid crippling infinities, characterization of resilience must be accomplishable without disturbances in mind. In light of this, quantum resilience defines resilience as the extent to which a system delivers its valued functions, and characterizes resilience as a function of system productivity and complexity. System *productivity* vis-à-vis specified “valued functions” involves (1) the quanta of the valued function delivered, and (2) the number of systems (within the greater system) which deliver it. System *complexity* is defined structurally and relationally and is a function of a variety of items including (1) system-of-systems hierarchical decomposition, (2) interfaces and connections between systems, and (3) inter-system dependencies.

Among the important features of quantum resilience is that it can be implemented in any system engineering tool that provides sufficient design and specification rigor (i.e., one that supports standards like the Lifecycle and Systems Modeling languages and frameworks like the DoD Architecture Framework). Further, this can be accomplished with minimal software development and has been demonstrated in three model-based system engineering tools, two of which are commercially available, well-respected, and widely used. This pragmatic approach assures transparency and consistency in characterization of resilience in any discipline.

DEDICATION

To my longsuffering and lovely wife

Laura

without whom this would never have been possible.

ACKNOWLEDGEMENTS

I gratefully acknowledge the partial support provided for this dissertation by the Fulton Fellowship of Arizona State University.

I thank those who acted as sounding boards and kept me honest without throwing ice water on this endeavor. Specifically, I appreciate the insights and challenges offered by my committee members Mike Chester and Marty Anderies.

I would particularly like to thank my advisor Brad Allenby who played his roles of provocateur and supporter with equal alacrity. His encouragement to pursue something a bit anti-establishment was freeing, and his graciousness in adopting an outsider into the sustainable engineering fold was exemplary.

Special thanks go to my wife and family. I look forward to becoming reacquainted with you all!

TABLE OF CONTENTS

	Page
LIST OF TABLES	viii
LIST OF FIGURES	x
PREFACE	xii
INTRODUCTION	xiv
OVERVIEW AND PRÉCIS: QUANTUM RESILIENCE IN A NUTSHELL	1
CLARIFYING RESILIENCE	15
Resilience and Robustness	18
Resilience and Adaptation.....	29
Resilience is Resilience.....	41
Resilience and the Adaptive Cycle.....	53
Resilience and the Urban Space	64
Resilience and the Safety Industry	71
QUANTUM RESILIENCE	76
Valued Function – Systems “For All They’re Worth”	77
Scale – Systems “To a Certain Extent”	86
Identity – Systems “As They Really Are”	89
Redundancy – Systems “Pure and Degenerate”	94
Defining Resilience	101
Characterizing Resilience.....	109
Quantum Resilience, Reductionism, and Complexity	121
Resilience and the Time Domain	124

	Page
Management of Units in the Numerator.....	132
Resilient Design	136
INSTRUMENTING QUANTUM RESILIENCE	140
Genesys™ for Quantum Resilience	141
Innoslate™ for Quantum Resilience.....	147
FractalSys for Quantum Resilience.....	152
ILLUSTRATING QUANTUM RESILIENCE	163
Lake-Agriculture System	164
Simple Desk	171
Imbricated Theater	172
Electric Power Plant/Grid.....	174
Pipe System.....	179
Referent Organizations.....	182
Urban Foraging	185
RESILIENCE IN THE GOULBURN-BROKEN CATCHMENT.....	190
Introduction	190
System Analysis	195
Valued Function	200
Resilience	206
Discussion	217
RESILIENCE IN THE PHOENIX FUEL SUPPLY CHAIN.....	224
Introduction.....	224

	Page
System Analysis	226
Valued Function	229
Resilience	230
Discussion	233
CONCLUSION.....	235
REFERENCES	238
APPENDIX	
A FRACTALSYS OVERVIEW	245

LIST OF TABLES

Table	Page
1. R-characterization: View Only	11
2. R-characterization: Safe View	12
3. R-characterization: Safe Accessible View	12
4. R-characterization: Function-differentiated Viewing System	13
5. R-characterization: Tower Only Viewing System	14
6. Systems and Their Environments	26
7. R-characterization: Canonical S_0	117
8. R-characterization: S_0 without S_D	118
9. R-characterization: S_0 without S_B	118
10. R-characterization: S_0 with Reduced Complexity S_B	119
11. R-characterization: Full S_0 (No Decomposition)	120
12. R-characterization: S_0 without S_D (No Decomposition)	120
13. R-characterization: S_0 without S_B (No Decomposition)	120
14. R-characterization: Lake-Agriculture System	169
15. R-characterization: Lake-Agriculture System (Ag Only)	170
16. R-characterization: Desk	171
17. R-characterization: Simple Theater	174
18. R-characterization: Multi-operator Grid	178
19. R-characterization: Multi-operator/Multi-turbine Grid	178
20. R-characterization: Simple Pipe System	180
21. R-characterization: Simple Pipe System (Differentiated Functions)	181

Table	Page
22. R-characterization: Simple Pipe System (Combined Function)	181
23. R-characterization: Two Organizations with Referent Organization	184
24. R-characterization: Urban Foraging	188
25. R-Characterization: Urban Foraging Simple Progression	189
26. Partial List of SRA-2 GBC River Valley Assessments	204
27. Valued Functions of the GBC	206
28. R-characterization: GBC Agricultural Economic Production	209
29. R-characterization: GBC Complex Economy.....	211
30. GBC Resilience by Valued Function.....	213
31. R-characterization: GBC Moderated Economy	214
32. R-characterization: GBC Pristine	216
33. Biophysical Resilience in all three regimes	219
34. R-characterization: Phoenix Fuel Supply	231
35. R-characterization: Phoenix Local Distribution	232
36. R-characterization: Phoenix Long-haul Transport.....	232

LIST OF FIGURES

Figure	Page
1. Simplistic “Viewing” System	10
2. System-of-Systems Providing Functional Redundancy Across Scales	98
3. Canonical Example: System-of-Systems with One Valued Function	116
4. Genesys Representation of Canonical S_0 Decomposition	142
5. Parameters Added to "Component"	142
6. Parameters Added to "Function".....	143
7. Block Diagram Showing Links between S_0 Subsystems	144
8. Decomposition of SB (into S_3 and S_4) and of S_4 (into S_5 and S_6)	145
9. Innoslate Representation of Canonical S_0 Decomposition	148
10. Attributes Added to "Asset"	149
11. Attributes Added to "Action".....	150
12. I^2 Chart of Input/Output Items	151
13. FractalSys Representation of Canonical S_0 Decomposition	154
14. Clarification of “System Boundary” and “Parent System”.....	155
15. FractalSys State Panel Showing Provision of Valued Function	156
16. Example of Interface Modeling	159
17. Simple Lake-Agriculture System.....	168
18. Simple Theater System	173
19. Referent Organization as Productivity Enhancer.....	183
20. Simplistic Foraging Supply Chains	187
21. Notional Overview of Actors and Forces in the GBC	196

Figure	Page
22. Simple Decomposition of Agro-Economic Sector	198
23. Simple Decomposition of Biophysical and Social Sectors.....	199
24. Typical Farm Relationships	200
25. Phoenix Fuel Supply - West Line	227
26. Phoenix Fuel Supply - East Line	227
27. Phoenix Fuel Supply System Overview	229
28. “Fractal” Nature of Systems Modeled by FractalSys	247
29. Recursive Definition of Transition Objects	249
30. FractalSys User Interface.....	250
31. Direct and Indirect Connections Between Systems	252
32. FractalSys Variable Scoping Protocol	252
33. FractalSys Data Model.....	253

PREFACE

I confess to being surprised by the definitional latitude allotted to the idea of resilience. The literature is filled with metaphor and analogy that permits resilience to serve as an umbrella for a wide variety of concepts that already had distinct and growing research areas. It is with an admittedly broad brush that I am painting, but if the evolution of the “resilience industry” had avoided this expansiveness, it might have been more effective. Instead, a significant amount of clarity has been sacrificed in the name of definitional tolerance. This has, unfortunately, resulted in stagnation when it comes to actually *operationalizing* the concept of resilience. In spite of the fact that one might frequently read or hear about “designing for resilience” there has been little rigor directed at actually achieving an ability to do that and there is certainly nothing that permits a measure of whether or not such a design has actually been accomplished. Even literature purporting to propose metrics for resilience frequently falls short because it generally offers long lists of “things” that must be considered when creating metrics, instead of actually providing a usable formulation.

As a veteran of many large systems engineering projects, I can also testify that the concept of resilience is rarely addressed by systems engineers (though I’m willing to add “until lately”). As to why that is (was) the case, I can only guess that it has something to do with an intuitive understanding among systems engineers that resilience was being provided by the many functional redundancies built into the systems. Whether telecommunications, national defense, intelligence gathering, or healthcare claims processing, “more” function was always better, and the challenge was always cost and

schedule management. It was *never* a question of *how* to make the system more resilient—that was a routine part of the job.

Still, resilience is a systems concept, so I found myself qualified to speak into this discourse. On occasion I found it difficult to hear my discipline misunderstood in the literature, but this only reinforced my desire to not only clarify the concept of resilience, but provide some insights into the discipline of systems engineering. Further, it was vital for me to provide a way forward—one that was acceptable and useful in the interdisciplinary future we know to expect. The result is quantum resilience. As a first generalized theory of resilience, it both clarifies the concept and provides a way forward for resilience theory.

My background and pragmatism drove inexorably to providing mechanisms that enforce transparency and consistency in resilience analysis. Importantly, it was required that this be accomplished with commercially available tools and scant methodological impact. There is no need for custom tools, just simple extensions that are easily vetted in the marketplace of ideas and then adopted and incorporated into commercial tools as needed. There is nothing particularly specialized about resilience analysis since it flows directly out of good systems engineering practice. Importantly, quantum resilience does not create a new “resilience elite” it simply operationalizes resilience for everyone.

INTRODUCTION

Resilience is frequently noticed to be heir-apparent to the increasingly diffuse sustainability discourse (e.g., Cascio, 2009). Unfortunately, this exposes it to similar risk of being lost as a usable engineering concept. Its popularity as a topic has ensured its mention in a large percentage of recent publications as researchers strive to demonstrate their awareness of its currency. Frequently the references to resilience are very casual and most merely attempt to contribute incrementally to the definition of resilience instead of actually demonstrating how to quantify it and use it in sustainable engineering practice. To ensure resilience can be operationalized, a generalized and quantifiable theory is required. Quantum resilience provides this.

Founded in general systems theory (von Bertalanffy, 1968; Rapoport, 1985), quantum resilience recognizes that to be a meaningful engineering concept, resilience requires rigorous definition and must be quantifiable and embedded in existing systems engineering approaches and tools. Systems engineers are uniquely aware of how the concept must be bounded if it is ever to be useful in engineering projects (both forward- and re-engineering). This resulted in a first step of clarification and distinguishing resilience in the discourse. Disambiguation was required because resilience is frequently conflated with broad and often inchoate ideas like adaptation, transformation, recovery, and learning. Quantum resilience specifically acknowledges the importance of such research areas, but holds that they do not belong in a generalized theory of resilience.

Importantly, quantum resilience not only clarifies resilience and makes it quantifiable, but it positions other scholarly work, and the common terms they use (e.g., “adaptation”), in the space. That is, it allows other researchers to see the merits of their

work and suggests where it must fit into the clarified discourse. For example, many are speaking of robustness and casually adding the word resilience in order to meet publication and grant requirements. Quantum resilience frees them to distinguish their work and suggest specifically how it contributes (e.g., to system “productivity” or “complexity”). Those who speak of resilience metrics and list items that should be considered in calculations will be challenged to rethink their ideas in terms of system productivity and valued function. Those who find resilience to be an emergent characteristic of their systems will be challenged to model their systems and specifically allocate the functions they value to specific parts of their system, forcing actual definition and proper system scaling. Quantum resilience also supports any forthcoming theories of complexity while specifically recognizing the need to analyze systems in order to understand them better.

By actually modeling the systems in question, quantum resilience removes the ability to mask important and frequently normative details behind euphemism and analogy. Quantum resilience makes the systems of interest completely transparent and initiates the important dialogue over what the specific valued functions of a system are, and how they should be quantified. It is generalized and operationalized for all systems in all disciplines at all scales, and is a required step in the establishment of resilience as a meaningful science. Quantum resilience allows resilience to be formally characterized for all systems at a time when no other theory is available to accomplish that. Researchers who speak of metrics but do not formulate them are frequently hampered by their inability to either identify or quantify normative concepts. Quantum resilience provides a mechanism by which transparency of such measurements can be guaranteed and

consensus garnered. The vanishingly small numbers of researchers who actually present calculable metrics for resilience tend to focus on system efficiency as a proxy for resilience which leads them into optimization campaigns and Monte Carlo explorations. This is not wrong, but it is not resilience theory.

Vitality, quantum resilience is easily instrumented in model-based system engineering tools that support industry standard lifecycle, unified, and system modeling languages (LML, UML, SysML) and fit into important frameworks like the US DoD Architecture Framework. Deploying the resilience characterization formulation in commercially available tools ensures broad adoption and maintains the goal of transparency and consistency. Since, for the first time, engineers will be able to compare homologous systems and make decisions based on alternative deployment options, it becomes vital they are using the same calculations.

Quantum resilience establishes a new measure of fitness for resilience literature, one that demands resilience be characterized for any system to which it is attributed. That is, if a so-called resilience theory cannot provide a characterization of a system's resilience, it must be abandoned. Though quibbling over specific methods and characterization approaches is expected and applauded, scholars must force themselves to specifically quantify resilience in terms of valued function and system structure, or resilience will no longer be a useful and operational scientific concept.

OVERVIEW AND PRÉCIS: QUANTUM RESILIENCE IN A NUTSHELL

As a concept that has become more expansively defined over time, resilience has struggled to achieve operational value. Quantum resilience allows the concept to be operationalized while serving to clarify it in the discourse. While resilience remains a *systems* concept, quantum resilience changes the focus of resilience analysis to the functions and services of value and enforces proper analysis of the systems that provide them. Resilience is not about a system until it is first about a valued function. The system which provides the valued functions can be characterized for resilience, but only *after* the valued functions are acknowledged. Quantum resilience employs system function and structure (productivity and complexity) to properly characterize a system's resilience.

Systems within all disciplines exhibit quanta of resilience which define the level at which resilience can be observed and (potentially) managed or engineered. A quantum of resilience is a unitized output of a valued system function or service. In any analysis, quanta of resilience must be defined and related in units that are germane to the valued function. Admitting to such a quantum enforces the consideration of scale and identity in anything that might be called "design for resilience." Importantly, there may exist in each system potentially dissimilar "subsystems" which redundantly deliver the valued function (whether by pure redundancy or degeneracy). If specific functions are not considered, designing (or managing) for resilience makes no sense, since it is the functional delivery level at which resilience can be discussed and designed. The implication is that managers and engineers need expend no energy "designing for resilience" unless they are addressing particular functional quanta. When resilience is an important design or management requirement, quantum resilience permits a more appropriate focus.

Quantum resilience takes a pragmatic approach to resilience theory and focuses on operationalization. *Once characterized, resilience defines the extent to which a system delivers its valued function.* In short, a system’s resilience can only be characterized when it is understood in the context of delivering its valued functions.

Quantum resilience asserts that resilience cannot be reduced to a universally absolute value but can be characterized on a per-system basis which permits comparison of similar systems or alternative system configurations. Further, because of the vital focus on function, characterization of resilience demonstrates that system resilience can be bolstered *only in increments* as quanta of each valued function are incrementally delivered by redundancies added to the system. In this way, resilience is a quantum concept. Resilience is not binary—that is, systems are not either “resilient” or “not resilient” because this is meaningless without reference to function and context. Further, resilience cannot be said to be qualitatively “higher” or “lower” unless specific reference is made to identified quanta of valued function.

Resilience is characterized by (1) identifying valued functions and the systems implicated in their delivery, (2) specifying the quanta of the valued functions delivered, (3) calculating the apparent complexity of the valued function delivery system, and (4) combining these in such a way that resilience of similar systems can be compared and contrasted in design or management decision frameworks.

Each of the principles of quantum resilience is discussed at length herein, but to briefly set the stage, quantum resilience forces analysts to acknowledge:

1. *Valued Function*. Initially, quantum resilience assumes nothing about systems but acknowledges that there are important functions and services provided by them—functions that humans have decided should be enduring (frequently acting in proxy for Nature). Human values are expressly acknowledged because quantum resilience starts with the question “what functions or services are valued?” Ultimately, since it is a human-created concept, the meaning of resilience is wrapped up in observer-dependent superimposition of value. Hence, to operationalize it, these values must be acknowledged and it must be recognized how they are reflected in the functions and services provided by the systems humans refer to as resilient. Resilience is only observed in the delivery of valued functions and services. Identifying (admitting to) the valued functions is the first step in properly analyzing resilience.
2. *Scale*. By focusing first on function, there is consequent and automatic acknowledgement of scale (both physical and temporal) because the answer to the question of value defines a scope at which the valued functions are (or can be, or should be) facilitated. In practice, and as witnessed in the literature, the system that delivers the valued function has been observed to be *larger* in scale than is initially thought. Admitting to appropriate scale often serves to temper the goals of engineers and managers who wish to design (or manage) for resilience. Further, valued functions must have a temporal extent that allows them to be valued at the scale on which they occur.
3. *Identity*. Acknowledging the scale at which the valued function is delivered leads to definition of the structure of the system that delivers the valued functions.

Structure and function comprise system identity. If a system structure changes such that a valued function ceases to be provided, the system has lost its identity vis-à-vis the valued function and cannot be considered resilient. At such time, it is acceptable to *carefully* suggest that the system has evolved (or adapted, or transformed) into a new system, but this is not the same as resilience. Instead, a recognition of (or search for) new valued functions can begin (with the new system in mind), or a search for an alternative system that delivers the originally valued function can occur. In general, a system that adapts or evolves becomes a different system. It stretches credulity to suggest the system is “resilient” if it fails to maintain its identity.

4. *Redundancy*. Once it is acknowledged that it is the valued function that matters in discussions of resilience and that this sets the scale and establishes system identity, it is easily observed that resilient systems are those which incrementally and redundantly deliver their valued functions or services (via pure or degenerate redundancy). Operationalizing resilience, therefore, will pragmatically focus on redundant provisioning of the quantum of resilience. Once the quantum of resilience is identified, a redundant and diversified portfolio of service delivery can be planned. Functional redundancy remains the driving design principle for resilient systems.

As discussed at length below, quantum resilience suggests that resilience is the extent to which a system delivers its valued function. Characterizing resilience requires both the function and the structure of the system to be quantified. Assuming system **S**

performs some function of value \mathbf{V} , the resilience of \mathbf{S} (R_s) must be a function of the system and its function:

$$R_s = g(\mathbf{S}, \mathbf{V})$$

Quantum resilience employs the notion of nearly decomposable hierarchies (Simon, 1962) to characterize system structure. This implies that as systems are decomposed, their subsystems will have progressively less inter-system interaction and become effectively more isolated. This has proven to be a worthy model over many years of systems engineering experience and is almost universally witnessed in Nature. Further, quantum resilience refers to the structural and relational character of a system \mathbf{S} as *apparent complexity*, C , and calculates it based on hierarchical decomposition (number of subsystems (s), interfaces and connections (c), and inter-system dependencies (d)). Hierarchical decomposition provides a measure of structural complexity, while the interfaces and dependencies provide a measure of the relational complexity. Complexity is recursively summed over the entire system hierarchy as shown in the formulation below. Other complexity formulations could certainly be substituted though (in keeping with Gell-Mann & Lloyd, 1996) no alternatives have been discovered that better “measure” system identity.

$$\mathbf{S} \stackrel{\text{def}}{=} C = s + \sum_{i=1}^s (c_i + d_i)$$

For clarity, quantum resilience refers to the overall system functional value \mathbf{V} as system *productivity* and asserts that system productivity, P , is the sum of the product of the quanta of valued functions provided (q) and the number of subsystems that provide those functions (m). Multiplying the quanta of valued function by the number of systems

that provide it implements a sort of force multiplication factor that recognizes how more independent subsystems providing the function (however little of it) magnifies system productivity:

$$V \stackrel{\text{def}}{=} P = \sum_{i=1}^f \left(m_i \times \sum_{j=1}^m q_j \right)$$

The final resilience characterization formulation is based on the intuition that resilience is directly proportional to system productivity and inversely proportional to system complexity. Hence, resilience can be characterized as follows:

$$R_s = \frac{P}{C}$$

The terminology of quantum resilience can be exemplified and clarified by thinking of a typical boulder you might encounter while hiking in a typical mountain landscape. Boulders are fairly rugged entities. They have existed for millions of years, been stepped on by a variety of species, been subjected to wind and rain, fire and ice, and withstood the onslaught with aplomb. In fact, it is likely that you could airlift one and drop it on pretty much anything and expect it to inflict more damage than it would sustain. Boulders are arguably, if superficially, robust and enduring. *But we do not learn about resilience from boulders.* Why? Very simply, we do not value them because they do nothing for us.

That changes when you arrive at a particular mountain pass and need a better vantage point from which to view the scenery. If you climb the rock and are presented with a good view, the boulder has taken on value. It is no longer simply a rock but a “viewing system.” If the view is enrapturing enough to evoke expressions like, “Wow! I

could stay up here all day!” or “I really need to tell other people about this!” you might well start to worry about the *resilience* of the rock, but this is only because you have started to value it for the function it delivered. Now it is a “system” that you wish to be resilient. Importantly, however, your interest in resilience is not incited by the system (the rock). Instead, it is brought about by the valued function it provided. This is the way resilience must be addressed if we are to operationalize the concept. That is, however, not really the end of the rock story.

Recall that the rock became an important viewing system because of the treasured view it facilitated. But depending on your perspective it might be considered short-sighted to see the boulder as the full extent of the system. In fact, if *the view* is so beautiful, you might want *that* to be enduring and resilient as well. This greatly expands the scope of the implicated system. If the valued function is extended beyond provision of a “view” to providing a “beautiful view” you must be concerned not just about the rock, but also about what is being viewed. Now it becomes apparent how valued function establishes the *scale* of the systems for which resilience is a goal. Further it reminds us that both aspects must be included in any model of the system.

Assuming you only cared about getting a “view” and did not care about what you were looking at (and recognizing that this is a conscious engineering or management choice) the next logical step is to agree that a rock is not the *only* way to get a view. In fact, there are many redundant and diverse ways to enable a hiker to get a view and express wonder. For example, you could climb a tree as easily as a rock, or you could walk on stilts. This obviously introduces other social aspects to the system. For example, to satisfy Occupational Safety and Health Administration (OSHA) regulations, a tower

with a safety railing might be required, and to meet accessibility requirements, an access ramp and elevator must be installed. Notice that while all you need is a rock, you must admit to certain other values that creep into the analysis. The rock might provide a *view*, but the tower provides a *safe and accessible view*. In every case, resilience must be characterized in light of the valued function in the context of the full system.

Note as well that the many uses of a boulder extend beyond making it possible to get a view. I once saw a sharp curve in a rural road adjacent to which someone had placed a huge boulder to ensure that overly aggressive speeding cars (presumably driven by resilient teenagers) did not venture too far beyond the road and come to rest in his living room. It is obvious that this person—perhaps from experience—valued the rock in a way that most do not. As with providing a view, there are many other ways to keep cars out of living rooms. Roads can be rerouted, guardrails can be installed, teenage driving can be restricted, etc. If the valued function was keeping cars out of living rooms, it must be seen that there are many redundant means by which this can be accomplished. The system need not include a rock, and it is the valued function that lends clarity and drives the analysis.

What if the boulder-with-a-view was crushed into a million small stones? Obviously, this is a significant insult to the system and one which might inspire some to lament that system's "lack of resilience" in the face of a particular perturbation. While this gets more into the idea of robustness (to be addressed at length later), recall that such lament would only occur if you lost some valued function. In this case, the resulting stones *could* be piled and climbed in order to view the treasured landscape, but it might be more difficult to climb, provide unstable footing, and not enable as far a view in the

end—a severely degraded function. Alternatively, you could argue the boulder has entered a new stability regime—one in which the services offered are once again different than providing a view. For example, now the boulder could be used in xeriscaping a Phoenix yard. But it is important to recognize that the valued functions are entirely different for this new system. Such a transformation disallows characterization of resilience to proceed as it did previously. This is new function and new system identity requiring new characterization of resilience.

This simple example is used in this brief introductory overview of quantum resilience to elucidate the formulation of resilience characterization. Later sections go into intricate details, but this will serve to set the stage and demonstrate the operational value of quantum resilience. It illustrates the need to focus on valued function, system scale, system identity, and redundancy when characterizing resilience for a system. Perhaps we *can* learn about resilience from a rock.

Extrapolation from this hypothetical rock story suggests a simple “combined viewing system” based on a rock and a tower. What is shown in Figure 1 adequately highlights the difference in system complexity between a rock and a tower (though clearly the tower can be further decomposed), and shows that both provide a view. Additionally, it should be pointed out that for simplicity sake, only hierarchical complexity is shown. There is no relational complexity in this simple example though it should be clear that connectivity and dependencies exist in even the simplistic tower model. The only relational complexity counted in this simple illustration is from the “output” of the valued function.

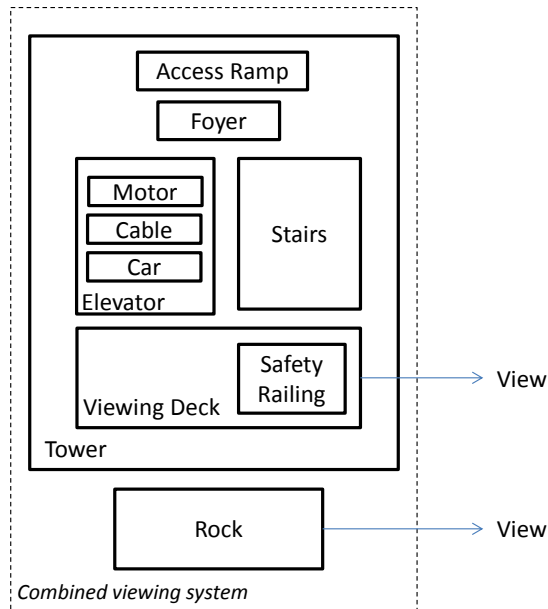


Figure 1. Simplistic “Viewing” System

Importantly, as I alluded above, a “view” is seldom enough for a system that is to be used by humans, so the example will present a comparative analysis of resilience based on a progressively complicated, but *integrated* function in which the “view” is first expanded to a “safe view” and then ultimately to a “safe and accessible view”. Then the same models will be presented with *differentiated* functions (i.e., view, safety, accessibility). This will allow some discussion of the differences in the resilience characterizations calculated by the model.

Note that for the system where “view” is the sole valued function, the tower is configured with *only* stairs and a viewing deck, dramatically reducing its complexity. For the simple example, assume the tower holds 20 people on the viewing deck and that the quanta of the valued function, “view”, is therefore 20 views. Further assume the rock holds 10 people and therefore provides 10 views. This results in the following calculations (note that the output of a valued function constitutes a connection):

Tower Productivity = 20 quanta of the valued function
 Tower Complexity = 3 systems (tower, deck, stairs) + 1 connection = 4
 Rock Productivity = 10 quanta of the valued function
 Rock Complexity = 1 system + 1 connection = 2

$$\text{Resilience (R)} = \text{productivity/complexity}$$

$$= (2 \text{ systems provide valued function}) * (20+10) / (4+2) = 10$$

Table 1 summarizes the outcome of the resilience characterization for the combined system where “view” is the only valued function.

Table 1
R-characterization: View Only

System (Function)	Productivity
Tower (View)	20
Rock (View)	10
Productivity	60.00
Complexity	6
R-characterization	10.0000

If the valued function is expanded to “safe view” the tower structure will require a railing, increasing tower complexity to 5. Note that the rock is not considered safe so it is calculated as if “view” and “safe” contribute equally to the function output. This means instead of 10 quanta when the function was “view” it becomes 5 because the rock can only provide *half* of a “safe view”. For this simple example, this is a fair approach, but recall that system experts would need to come to consensus and thereafter be consistent when such situations arise. It would be equally valid to determine the rock simply can no longer participate in providing the valued function (as, for example, is demonstrated in the function-differentiated example to follow).

Tower Productivity = 20 quanta of the valued function
 Tower Complexity = 4 systems (tower, deck, stairs, railing) + 1 connection = 5

Rock Productivity = 5 quanta of the valued function (just “view”, not “safe”)
 Rock Complexity = 1 system + 1 connection = 2

$$R = \text{productivity/complexity} \\ = (2 \text{ systems provide valued function}) * (20+5)/(5+2) = 7.1$$

Table 2 summarizes the outcome and the lower resilience characterization. Note the increase in complexity due to adding the railing and the decrease in productivity since the rock can only provide a view that is not safe.

Table 2
R-characterization: Safe View

System (Function)	Productivity
Viewing Deck (Safe View)	20
Rock (Safe View)	5
Productivity	50.00
Complexity	7
R-characterization	7.1429

To provide a “safe and accessible view,” all the complexity shown in Figure 1 is required. Here, the rock is neither accessible nor safe, so following the previous approach, it can only contribute one third of the valued function (10/3). Table 3 demonstrates how the decrease in productivity and the increase in complexity further erode the resilience of the system.

Table 3
R-characterization: Safe Accessible View

System (Function)	Productivity
Tower (Safe and Accessible View)	20
Rock (Safe and Accessible View)	3.33
Productivity	46.66
Complexity	13
R-characterization	3.5892

In general, it should be obvious that as complexity increases, resilience is diminished. But also note that increases in productivity can often offset increases in complexity and maintain resilience. Importantly, when the valued function is integrated into one function (i.e., “safe and accessible view”), the rock has limited ability to provide it. Hence, it might be interesting to determine if isolating the valued functions results in a different outcome.

The same modeling can be done using differentiated valued functions: view, safety, accessibility. Assuming experts are happy with the way the “view” is quantified, decisions must be made about how to quantify safety and accessibility. OSHA may publish guidelines or regulations that are used by inspectors, but importantly, it matters only that consensus is gained among the experts and that consistency is employed if such models should ever be compared to others. For the purpose of this simple example, it is arbitrarily determined that one full unit of “safety” is provided by the railing, and the ramp and the elevator each provide one half a unit of “accessibility.” Table 4 shows the outcome of the resilience characterization.

Table 4
R-characterization: Function-differentiated Viewing System

System (Function)	Productivity
Access Ramp (Accessibility)	0.5
Elevator (Accessibility)	0.5
Railing (Safety)	1
Viewing Deck (View)	20
Rock (View)	10
Productivity	63.00
Complexity	16
R-characterization	3.9375

Note that in all cases modeled, the rock and the tower contribute redundant “viewing” capacity. If safety is a paramount concern, it could be argued that the rock must be removed from the analysis. This can easily be done and Table 5 summarizes the outcome.

Table 5
R-characterization: Tower Only Viewing System

System (Function)	Productivity
Access Ramp (Accessibility)	0.5
Elevator (Accessibility)	0.5
Railing (Safety)	1
Viewing Deck (View)	20
Productivity	23.00
Complexity	13
R-characterization	1.7692

Note the dramatic decrease in resilience due to two factors. First, the rock’s incremental delivery of valued function “view” has been removed. Second, there is no longer redundancy in the provision of “view”.

Though this example has all the risks associated with any oversimplification, it is clearly useful in demonstrating the specifics of quantum resilience characterization. Valued function is figured predominately and is used to establish the scale at which the analysis proceeds. Identity of the system (structure and function, see later discussion) is clearly tracked through the analysis providing adequate basis for comparative analysis of alternative system deployments. Finally, the contribution of redundancy is easily observed.

CLARIFYING RESILIENCE

The resilience literature has enjoyed a several-decades-long season of unrestrained expansion. During this period many titillating ideas have found safe haven under the resilience umbrella, but it is vital that resilience be understood for what it is without these expansions. These other ideas certainly deserve to be researched and pursued, but they should not be conflated with resilience. To introduce this chapter, I offer a simple taxonomy to position resilience in a highly confused literature. Each of the concepts is discussed at length in the following sections.

As defined herein, once properly characterized, resilience is the extent to which a system delivers its valued function. Resilience must take into consideration the functions provided by a specifically scoped system and therefore it implicates system function (what it does) *and* system structure (what does it) in its characterization. Note well that resilience is defined on a scale or spectrum, so it cannot be said to have opposites. Suggesting an opposite for resilience is like suggesting an opposite for a certain wavelength of light, or a temperature of 60 degrees. Since all systems are resilient to a degree which can be characterized, there is no antonym for resilience.

Robustness is the extent to which a system is protected from the environment in which it must operate. Jen (2005, p. 17) understands this protectionist idea, but fails to understand that it is not simply unanticipated disturbances from which systems must be protected. Robust design includes protection against threats and disturbances which are expected (as with evolutionary biology) as well as those that can be reasonably predicted (e.g., think of the Mars Lander design). Though caution must be used with this terminology (more below), opposites of robustness include *specific* vulnerabilities.

Robustness cannot be said to contribute to resilience unless a proper characterization of resilience has demonstrated that particular protections enhance resilience. That is, sometimes features designed (or evolved) to contribute to a system's robustness do not directly contribute to valued function delivery and will only be represented as part of the system's complexity (in the denominator of the quantum resilience characterization).

Though glaringly absent from the resilience literature, *anticipatory systems* (Rosen, 1985) are a breed of system that contains a model of the system and its environment. Importantly, the model operates faster than real time providing a limited look-ahead into possible futures. Depending on the way they are instrumented and the sophistication of the model, anticipatory systems can plan a future based on what is deemed an optimal path through a variety of projected landscapes. Humans are anticipatory (for example, we experience what some refer to as a "theater of the mind") so systems with humans in the loop are sometimes observed to be anticipatory (but again, this depends on how they are instrumented). In fact, though they go unidentified as such, *it is anticipatory systems that are most frequently confused with resilient systems*. Sometimes such systems are referred to as adaptive, but since adaptation can only be seen after the fact it is best to avoid that word. Biological organisms are anticipatory to varying extents depending on their complexity and how they are instrumented.

Adaptive systems are generally anticipatory though their internal models may be very simple. Living organisms from the smallest bacteria are adaptive because they satisfy two criteria. First, adaptive systems are homeostats that can regulate themselves (within tolerances) for operation in a given environment. Given the definition above, this obviously means they are, to a specific extent, robust. Second, they also have some

manner in which to “remember” phenotypic “lessons-learned” in their genotype. For many millennia, biological life has been adaptive, but the aforementioned “remembering” mechanism by which the genotype was enriched was managed by completely arbitrary and accidental processes. If a bacterium (phenotype) happened to get the right random mutation in a gene that made it somehow survive longer and pass that gene on to its progeny, it is (retrospectively) considered adaptive. Humans are adaptive systems and, having mastered language and memetic evolution, are no longer at the mercy of accidents of nature to transfer lessons-learned to their progeny—effectively extending the genotype to include culture, society, and its tools. Today, humankind determines what is adaptive and to a great extent what is selected. Because humans are adaptive, some of the systems in which we participate are also adaptive, starting first as anticipatory systems and surviving long enough for us to look back upon and notice they have been selected. Adaptive systems can certainly be resilient, but that has very little to do with them being adaptive. In fact, a system that must be adaptive to be resilient is taking a great risk.

Learning systems are special in that they have “solved” what scholars refer to as the stability-plasticity dilemma (cf. Grossberg, 1980). Not only do such systems have a repository of learned behavior (usually a real or artificial neural network), but such systems have “figured out” how to remember new things without forgetting what they already know (a phenomenon known as “catastrophic forgetting”). There are many theories on how this is best done, including those which have a filter that suggests what *should* be learned (e.g., a “resonance”) and those that implement short term memories that become learned once reinforced and promoted. Again, when learning systems are conflated with resilience, it is because humans are in the loop lending use of their neural

networks in an attempt to improve their future by not repeating negative past experiences, or, as is frequently the case, normatively dictating what *should* occur in the future. A system can be (and is) resilient *without* being a learning system.

Transformation beggars the idea of resilience, anticipation, adaptation, and learning since it resets all systems back to experimenting with accidental evolution, spontaneous generation, punctuated equilibria, or flashes of imagination. Transformation is not related to resilience because transformation necessarily results in a *new* system. Transformation is a wonderfully aggressive and revolutionary-sounding word that permits humans to pursue their hopes and dreams based on projected futures which may or may not be supported by historical trends or anticipatory models. It is also a dangerous word since it usually assumes facts not in evidence and relieves us of the need to do the hard work of properly characterizing resilience.

Resilience and Robustness

As clearly indicated in the definition and characterization formulas, resilience is *not* about perturbations, disturbances, environmental threats, or other insults that might be suffered by our systems. This section discusses why, and proposes the resolution is the tried-and-true idea of robustness as a unique and independent concept. Understanding the environment in which our systems operate teaches engineers to expect and protect against perturbations, but including perturbations in characterization of resilience is not a way forward for resilience theory. Instead, perturbations fall under the purview of robustness, forcing the distinction of resilience and robustness as separate concepts. Very simply, robustness is environment specific, resilience is not.

The rationale for excluding perturbations from resilience analysis is contained in the following (rather long) sentence: If systems can be said to be resilient, and such a moniker applies in some degree to *all* systems, yet individual systems are exposed to only a subset of all possible perturbations by virtue of their specific operating environments, and if this subset of perturbations is different from, or not entirely a subset of the expected perturbations on another system operating in another environment, then the idea of perturbations can neither contribute to any generalized definition of resilience nor be included in any generalized characterization of resilience without introducing infinities into the formulation.

Assume E is the universal set of possible environments in which systems operate. Assume P is the set of all possible perturbations, imaginable or unimaginable, infinite in scope, a universal set of perturbations collected from the universal set of environments. Assume system S_1 operates in a target environment e_1 (a subset of E) wherein some subset p_1 of P defines the perturbations on S_1 . Further assume system S_2 operates in a target environment e_2 wherein some subset p_2 of P defines the perturbations on S_2 . Assume that p_1 is not equal to p_2 and that $p_1 - p_2$ is not the null set. That is, perturbation sets are free to intersect but are not proper subsets of each other.

$$\begin{aligned}
 E &= \text{universal set of environments, } E = \sum e_i \\
 e_1 &\subseteq E \\
 e_2 &\subseteq E \\
 P &= \text{universal set of perturbations, } P = \sum p_i \\
 p_1 &\subseteq P \\
 p_2 &\subseteq P \\
 p_1 &\neq p_2 \\
 p_1 - p_2 &\neq \emptyset
 \end{aligned}$$

If resilience characterization were to be generalized *to include perturbations* for *only* systems S_1 and S_2 , such a generalized characterization can easily be seen to require acknowledgement of all perturbations (p_1 and p_2) in the union of e_1 and e_2 . This is obviously larger than smaller subsets p_1 and p_2 , but perhaps still a tenable formulation. Extending this, however, to *all* systems S_i and their environments e_i and associated perturbations, p_i , in order to achieve a completely generalized resilience characterization approach would require acknowledgement of the entire set P , the universal set of perturbations, and the entire set E of environments. It is clear, then, that any generalization of resilience that attempts to acknowledge perturbations is intractable.

Instead, quantum resilience asserts that specific defenses against specific sets of perturbations can be considered under the concept of robustness. Robustness can be defined as the non-infinite set of protections required for S_i to operate in e_i . Since a robust system exists within a specific environment, there is no generalized formulation of robustness, just a list of specific disturbances and defenses. In fact, it is fair to say that systems are robust against specific insults on a case-by-case basis to a certain degree. Even if this list is large, it is finite. For example, a satellite might be able to operate at temperatures down to -50°C so it is robust to that degree, but it would fail if exposed to *any* lava flow, so it is not robust against that.

Carpenter et al. (2001) suggest that resilience analysis must answer the question “of what to what?” In light of this question, they define two kinds of resilience: specified and generalized. They suggest specified resilience has to do with responding to perturbations we know and expect, and generalized resilience is about responding to threats we predict might occur. In fact, what these authors have proposed is exactly the

process employed by engineers as they prepare a system to operate in the environment for which it is designed. Most engineers would agree that their systems must be robust enough to perform their duties within their operational environment (including some wiggle room or, tolerances that are sometimes summarized in a system specification or service level agreement). If systems fail to perform in their target environment, engineers are willing to agree they are poorly designed. What Carpenter et al. are promoting is the simple concept of robustness. In fact, if resilience is about either “specified” or “generalized” perturbations, its characterization can never occur because as demonstrated above the calculations are intractable. No matter how long the list of perturbations, it is always possible to suggest another. Such infinities are the bane of science and must be eradicated, there is no “to what” in resilience analysis. *Resilience is about the system, not possible perturbations on the system.*

Quantum resilience asserts that resilience is the extent to which a system delivers its valued function. It goes on to assert that this characterization of resilience must be accomplished without taking perturbations into account. That is, a system is resilient to some degree *before* any perturbations strike. If we can only characterize resilience after perturbations strike, it is not a useful concept for forward engineering projects since an infinite progression of destructive testing would be required to verify a system’s resilience. Robustness, on the other hand, deals specifically with expected and predicted perturbations by employing solid engineering practice to instrument protections against them and to define the acceptable degree of defense against them. In this manner resilience is distinguished from robustness.

Robustness is most helpfully thought of as a constellation concept that applies to specific systems and their specific environments. In keeping with Jen (2005, p. 17), robustness can be defined as the extent to which a system protects itself from known and projected threats. As a concept distinct from resilience, robustness can be validated and verified with solid test efforts that can be specified as needed. Robustness allows for a discipline-specific short-list (or, constellation) of “protections” that adds no infinities because each perturbation and the associated protection is largely isolatable from the others. Robustness in the face of specific perturbations can be managed and implemented on a case-by-case basis. Each threat can be evaluated for the extent of system exposure (vulnerability) and appropriate defenses created to protect the system from the threat, or to respond to the threat. This is typical engineering design. When threats are only postulated (as for example they were for missions like the Mars Lander where the environment could not be known in full), they can still be managed on a case-by-case basis with our best available vulnerability and failure modes analyses.

Robust systems are so-called because they satisfy criteria which have accreted around particular classes of systems (e.g., bridges, satellites, governments, mammalian cells, etc.). That is, there is a constellation of protections that are associated with a system by virtue of its inclusion in a particular class of system. If a system in that class has the expected protections, it might earn the moniker “robust”, but it is important to be specific about the specific areas of defense and the degree to which the system is protected. These lists of protections can be short or long, rough or detailed, depending on one’s familiarity with the system class. For example, most people can understand that a bridge is (or is not) “robust” in a way that is different from a satellite. Bridge engineers know more

specifically about gravity, load, erosion, and concrete strength, and that discipline developed a long list of standard procedures to ensure robustness of their products. Similarly, satellite engineers know more specifically about vacuum, zero-gravity, attitude and orbit maintenance, high radiation, and vibration at launch, and have evolved a discipline that consistently delivers robust systems for operation in space. In effect, it is malpractice to build a bridge without solid footings and to launch a satellite into orbit without taking the vacuum of space into consideration.

Further, the constellation of protections may specifically exclude certain perturbations because of the class of system under discussion. For example, no carpenter who produces hand-carved mahogany desks requires that they be protected against the onslaught of a chainsaw—he might, but if he did not he would not be accused of malpractice. To earn the moniker robust, a wooden desk need not be protected against such an insult.

These domain-specific lists of protections that make a specific system robust are well-understood in the given discipline and frequently find themselves in designs. Importantly, to a significant degree, the perturbations and their associated protections are *independent*. For example, a satellite engineer can adequately protect against the unique thermal issues of space but forget to use radiation-hardened electronic parts. It can be easily argued that though such a satellite might function, it is “less robust” because of such a failing, but note well, it is less robust in that specific defense and could be very robust in other areas. In general, each of the protections can be measured in isolation from other protections on the list. This removes the crippling infinities.

In light of this, it is important to manage the vocabulary. Like resilience, robustness really has no antonym because it is a constellation concept. The word that comes closest to being an antonym is “vulnerable” but care must be taken to identify the specific vulnerability (e.g., a satellite that is invulnerable to vacuum, may still be vulnerable to mudslides). The word “resistance” is also sometimes seen in the literature. Resistance can be considered the outcome of a specific “protection” that is implemented in order to make the system robust against a specific threat or expected disturbance. In this regard, resistance can be viewed as defining specific operating ranges or a service level agreement for the system (e.g., paint that should not be applied at temperatures lower than 40F, or a wristwatch that is waterproof to depths of 50 feet).

Robust design

It would be ludicrous to assume a system is *not* resilient simply because it cannot accomplish something for which it was not designed, but this is the logical outcome of including perturbations in definitions of resilience. For example, it would be silly to suggest an electrical power plant is not resilient simply because it cannot serve sandwiches and coffee to the public. While such an extreme example easily makes the point, researchers sometimes fail to see the similarities of that example with demanding that a power plant continue to function after being submerged by a tsunami. Somehow recovery after a natural disaster prompts us to implicate the idea of resilience, while demanding silly bistro functions does not. Instead, after admitting that a power plant was not *robust* against the onslaught of sandwich demand, we can consider whether it must be

redesigned to support these new bistro requirements, but we cannot suggest it is not resilient.

The same argument applies when we expect pristine Nature to survive intact through the exigencies of human development. Instead of realizing that Nature is not robust against the human challenge, many argue instead that its resilience has been diminished—but this can only be demonstrated after valued functions are identified and the system is properly characterized. That Nature has been so incredibly durable sometimes leads us to assume it is “resilient” when in fact, human development amply demonstrates Nature’s inability to meet expectations for which it was not “designed.” In such cases, engineers must admit that we are privileging Nature as we currently know it (or as we think it can be), and are expecting that Nature wants to remain that way. Obviously this is as ludicrous as a power plant serving sandwiches. Nature has never remained in any particular state, and assuming a “reference condition” from which it has decayed merely reflects normative beliefs.

Simply because a system cannot meet expectations for which it was *not* designed does not make it a candidate for a resilience analysis. Instead, it is either a candidate for redesign, or simply a silly notion. As a less extreme example, ask whether or not a teenager’s wardrobe consisting of jeans and T-shirts is not resilient simply because it cannot deliver an appropriate outfit for a Whitehouse dinner invitation. It is not legitimate to discuss this deficiency in terms of resilience. Instead, it is valid to say that a Whitehouse dinner invitation is a “disturbance” the wardrobe was never designed to handle (because it was not designed for the Whitehouse environment), and to consider a redesign, but this does not enter the resilience space. It is better to suggest that the

wardrobe is insufficiently robust, or inadequately designed to be responsive to such an invitation. Could a wardrobe be augmented to support such an environment once the need arises? Absolutely, but note that such an augmentation must come from *outside* the original system. It would require expanding the wardrobe system to include, for example, a shopper and a nearby tuxedo store. This is valid system expansion, but it would be wrong to suggest the wardrobe is “adaptive” because of this system expansion. Instead, as a part of a human-in-the-loop larger system, one could argue that someone was able to upgrade the wardrobe system to be useful in yet another environment. In fact, if one was *anticipating* a Whitehouse dinner invitation, a tuxedo may have already been added to the wardrobe in order to make it robust enough to withstand such an environmental insult. But again, this is not the same as the original wardrobe system being adaptive. Similar thought experiments could be staged for the remaining examples in Table 6. Such examples serve to illustrate that frequently there is confusion over the idea of resilience.

Table 6
Systems and Their Environments

System (function)	Normal “designed-for” Environment	Environment <i>outside</i> design parameters
Wardrobe (appropriate outfits)	Work, recreation, home outfits	Whitehouse dinner
Bridge (vehicle conveyance)	Gravity, vehicle load, erosion, thermal expansion	Zero-gravity of space
Satellite (science, or telecommunications, or surveillance, etc.)	Vacuum, thermal environment, zero-gravity, attitude and orbit control	Under water
Commercial Aircraft (transport people and cargo)	Controlled civil airspace	Active jamming, live fire
Desk (work surface and locked storage)	Level office floor, indoor setting, drawer use	Chainsaw, stairwell
Office building (office space)	Desk work, meetings	Automobile repair, junkyard

Erica Jen (2005) suggests “systems that are robust often are required to maintain their functions while exploring new functionality” (p. 16). Taken baldly this statement is clearly overreaching. For example, should a bridge “explore” serving coffee or teaching kindergarten while it continues to convey vehicles over a ravine? What would exploration of new functions resemble for a bridge? Does this mean a bridge cannot be robust? Obviously, such a caricature is not her intent, but her specific example illustrates it is not far from her meaning: she suggests that the Internet must be able to be upgraded “without interrupting functionality.” There are several things to comment on here. First, though perhaps a quibble, her example is not really about “exploring new functionality.” That is, her example does not leave the Internet performing new functions after the change. Instead, she describes a protocol upgrade that *alters the manner in which a current function is delivered*. Second, and importantly, her proposed change is *certainly not* at the Internet scale, which is the scale she defined as her system. That is, it is not the Internet that is “exploring” the new function. She has casually expanded her sense of the system from protocols and data delivery between computers to include the engineers and organizations that build and maintain those protocols. My discomfort with this approach is that it is simply too casual. In the case of Jen’s Internet example, the proposed “robustness” of being able to upgrade while still delivering function is facilitated by the *redundancy* in the network. No network engineer would argue that the software implementing a protocol can change while it is operating. There must be a discrete point in time when one approach stops and another starts (and for protocols that generally involves both ends of the interface). Third, her example is not really about robustness.

Instead, she is suggesting that the designers and builders of the system are actually changing it and deploying a new system. It is tantamount to saying that a factory floor producing widgets is “robust” because the assembly line can be shut down, reconfigured, and restarted to produce a different kind of widget. While this is certainly a testament to good planning and engineering, it beggars the idea of robustness. A better example of robustness in the Internet is that you specifically *cannot* change a protocol while it is operating. This illustrates the ability to defend against a known perturbation and ensuring the integrity of the communications interface.

Why does it matter?

The distinction between resilience and robustness is vitally important because, first, there must be clarity in the discourse. Conflating terms will never lead us to quantifiable results. Second, allowing perturbations and robustness (the “solution” to perturbations) to be implicated in the quantification of resilience leads to infinities that cannot be tolerated. Third, recall that most features that protect a system (i.e., augment robustness) arguably exist simply to enable the system to deliver its valued function in the first place. The satellite in orbit *may* protect itself from the unique thermal environment presented by space, but thermal management is a critical part of the system without which it could not perform its mission. Similarly, a factory built at the bend of a frequently flooding river might have a levy to protect it from the inevitable flood, but this must first be considered good design that contributes to system robustness and makes it possible for the factory to perform its purpose. Only after taking the entire system into

account and characterizing resilience properly can we determine if the levy contributes to or detracts from the factory's resilience.

The individual protectionist features may actually contribute to the complexity of a system without contributing to its productivity. This is why it is so vitally important to fully document and model the system and to appropriately allocate functions to subsystems. Only when the full complexity of the system and the full productivity of the system are calculated can we characterize the resilience of the system.

Resilience and Adaptation

Among the more problematic trends in the literature is the tendency to conflate adaptability with resilience, even suggesting that “adaptability is part of resilience” (Folke et al., 2010). In the process, the words have been redefined to expand the scope of the resilience research project to an extent that makes such a project intractable. The temptation to conflate these concepts is irresistible because of the simple equation that emerges when it is noticed that Nature is *both* adaptive *and* resilient. While quantum resilience provides a testable theory about why and how these concepts are *connected* (discussed later), it also recognizes them to be separate research areas. Resilience must be a useful scientific concept without conflating it with adaptation.

Though speaking in a different context, Ahl & Allen (1996) refer to “adaptation” as a “weasel word,” one that “confuses law-like processes that move a system forward, and observer-based rules that recognize significance and purpose after the fact” (p. 190). They are correct. Our ability to predict the future is fairly limited, but such an effort generally starts with anticipating future environments for our systems based on current

trends established with short histories. Once these future environments are projected, engineers begin to focus on how systems can be protected against or adjusted to survive the expected future landscapes. Frequently so-called adaptation is simply an exercise in good planning, including brute force methods like stockpiling resources in order to meet future maintenance or operational demands. Future system evolution becomes simply a matter of humans executing their plans on a schedule that has been determined by the cognoscenti. References to resilience are then simply self-congratulatory because if a new landscape was adequately predicted and successfully navigated, we have proven our own resilience. This brings us no closer to understanding resilience.

Typically, since we can anticipate only a few things we tend to privilege the present state and target that (or an only slightly modified version of that) in our adaptive management plans. There is no consideration that a completely different future might actually be better. This feeling persists despite the fact that throughout history we have consistently proven to ourselves that radically different and unimaginable futures are *always* better. Unfortunately, since real adaptation is an extremely risky approach to resilience, we will tend to fall back on limited projections of incrementally alternative futures and revert to good planning and safety reserves to see us through. Instead, real adaptation requires definition of future landscapes, definition of fitness on those future landscapes, exploration of evolutionary pathways for the phenotypic system, feedback with measures of fitness, etc. Not only does this become intractable very quickly, ultimately our success is only observable after the fact.

Still, the resilience literature has begun to define resilience in terms of adaptability. Since resilient systems can superficially be said to demonstrate higher

fitness, it has proven to be an irresistible equation and the Resilience Alliance (cf. Folke et al., 2010) has struggled to secure resilience as an independent research area.

Unfortunately, they extended their umbrella framework into the realm of adaptive and transformative behavior before resilience was even codified and operationalized as an individual and important scientific concept. Because of this, resilience has become another word like “sustainability” which is defined as needed to meet the needs of the moment.

While on the surface adaptation *might occasionally* be construed as a mechanism that prolongs the existence of some systems, it can equally well be a mechanism that radically changes the system enough to kill it, make it unrecognizable, or unable to deliver its valued function. Dramatic changes in function (and associated changes in structure) *may* make a system more *fit* on a particular landscape, but this does not mean the system is more resilient. If function and structure change, how will the system be recognized and how would one argue for its resilience? Even if function remains, but structure changes, it is a new system that is providing the function. If the old system is gone, how can it have been resilient? Adaptation can only be seen as an extremely risky effort at targeting resilience—and the outcome is only known after the fact. This term has introduced even more confusion into an expansive literature that continues to miss the opportunity to operationalize the concept of resilience.

Toward the end of their piece, the team from the Resilience Alliance says:

Confusion arises when resilience is interpreted as backward looking, assumed to prevent novelty, innovation and transitions to new development pathways. This interpretation seems to be more about robustness to change and not about resilience for transformation (Folke et al., 2010, p. 25).

This is troublesome, but it is not just because phrases like “transitions to new development pathways” are vague, unexemplified, and devoid of actionable content. Neither is it troubling because “resilience for transformation” sounds grandiose and promising, but provides no operational target. This statement is troublesome because it disparagingly refers to other definitions of resilience as “backward looking” (whatever that means) and makes the assumption that other definitions of resilience somehow seek to “prevent novelty” and stop innovation. Not only is there no basis for this charge, it is highly self-serving. In fact, “confusion arises” when researchers continue to expand their definitions, open doors to all manner of alternative approaches, and include them under their research umbrella. Note well that systems *can* be resilient *and* adaptive *and* transformative (and many other cool-sounding words *whatever* they may mean), but these are all unique and independent ideas—and must remain so. Instead, such comments establish their definition of resilience as a bright and shining star that pretends to guide us to novel, innovative, and transformative salvation, while providing no mechanism to instrument our safe passage.

The consensus work of Folke et al. (2010) leverages the definition proffered by Walker et al. (2004, p. 4) where resilience is suggested to include the capacity to “reorganize while undergoing change.” As the proverbial camel’s nose under the tent, this idea of “reorganization” has been intermittently employed over time and has finally found purchase to the extent that ideas like adaptation and transformation are cited as goals for “resilience thinking.” It is troubling that there is an unqualified pretense of some ability to track the “resilience” of a system *before* and *after* a transformation given that

(1) words like “transformation” suggest that *entirely* different systems are being discussed (making their resilience incomparable), and (2) absolutely no mechanism for quantified characterization of resilience is provided. Perhaps it is possible that some kind of “transformation” can occur without impacting delivery of valued function, but it is very risky and it is absolutely required that this caveat be mentioned. It is meaningless to talk about resilience in the context of changes to systems that alter the functions they deliver and change their interfaces to the rest of the world. Simply “defining it away” by including phrases like “change to maintain” (Folke et al., 2010), is not sufficient.

For adaptation to be a useful design goal, we must have some knowledge of the future landscape. Rosen (1985) points out:

It must be emphasized that it is in fact meaningless to characterize a phenotype, or behavior, as adaptive apart from a postulated measure of fitness.... It must be stressed that, until a measure of fitness is introduced, no simple state can be meaningfully characterized as adaptive or maladaptive; and this includes all the cybernetic mechanisms which have been proposed as examples of adaptation. This is a subtle point, but it cannot be emphasized too strongly.... The concepts of adaptation, fitness, selection and evolution are themselves linked; none of them can be completely understood unless all of the others are taken into account (p. 374ff).

As a theoretical biologist and polymath, Rosen contributed the definitive work on anticipatory systems. While reviewing his mathematics on adaptation and selection is beyond the scope of this work, it should be pointed out that in most cases when adaptation is discussed in the resilience literature the adaptive systems referenced are really anticipatory systems. It behooves us to understand that adaptation is really something we can only discern after the fact—and that makes it difficult to sell as a design goal. Further, it is generally forgotten that *any projection of future fitness will*

necessarily be a normative projection. A future we pursue is always one we feel *should* be pursued.

Still, given our human foresight, it seems we should be able to set ourselves up for future success and that somehow this should become a matter of good design. To accomplish this, we must leave behind the idea that we can build systems that can “change” or transform. Recall that any system with humans in the loop qualifies as such. Instead, we must focus first on building systems that are robust in the environments in which they operate, and that have spare capacity to deliver their valued functions. Such spare capacity can legitimately be termed *adaptive* capacity, not because it does any kind of adapting, but because it provides extra system structure which can be repurposed without impacting overall system productivity. Such an idea flies in the face of human engineering efficiency (especially when we are being told that resources are dwindling), but this is the way nature works as can be easily illustrated.

Once upon a time there was a widget factory on the river with 100 highly productive widget manufacturing machines. One sad day, one of the widget machines broke in a way that resulted in production of defective widgets (which became affectionately known as wonkets). The quality control group recognized the defects, and because widgets are difficult to recycle, simply discarded them behind the factory, forming a pile. Because widgets were selling so well, management barely noticed the 1% decrease in manufacturing efficiency and failed to investigate the waste of energy and materials. Over time the pile of wonkets grew, forming a large pile and creating a barrier between the factory and the river. Some creative employees even created a picnic area on the pile and would regularly enjoy the view while eating lunch. Then, on another sad day,

the river flooded unexpectedly and the pile of wonkets, serving as a levy, protected the factory by diverting the water and saving it from a disastrous flood. Management and employees alike rejoiced and cheered at their good fortune.

This is how *degeneracy* works in nature (more below). This is how nature builds adaptive capacity. Nature is completely oblivious, extraordinarily wasteful, and *sometimes* successful. As you might expect, having been saved from the flood by their pile of wonkets, management immediately formed a team to investigate where they originated. In the course of the investigation, the team found the broken machine and repaired it so it once again was manufacturing quality widgets. Then, management invested a large sum of money to have the wonkets removed to a landfill and to design and install a proper levy to protect against future floods. They have yet to establish a serviceable picnic area. That is how purposeful and efficient humans work. The differences are dramatic and serve to illustrate why we struggle with the idea of adaptation.

Active functional redundancy (i.e., degeneracy) can be said to provide adaptive capacity in the sense that delivery of more function allows risky and potentially deleterious experimentation with system parts while not falling below some necessary service delivery level. Assuming the system continues to deliver its valued function, this degeneracy can be viewed as a mechanism conferring greater resilience, but it is due to redundancy, not because of any as yet unproven “adaptations.” In fact, adaptation, as illustrated by the widget factory, comes through coopted redundancy (cf. Gould, 1997; Gould & Vrba, 1982). The factory and its management had no idea it was setting itself up for success on a different (flooded) landscape. Likewise, we have no way of planning for

adaptation other than providing redundant capacity that might eventually be tweaked, coopted, and employed for new functions.

As nature demonstrates and as illustrated by the widget factory, adaptation seems to be best orchestrated as opportunistic exploitation of redundancy through cooption. Kauffman (1995, p. 154) outlines how evolution by natural selection cannot really operate on a fully optimized “program” (or, organism) since optimization means there is no room for error. In many respects this disqualifies engineered systems from adaptation since they are generally designed to minimize extra parts and cost. Hence, while the extra parts in Nature may be added through extravagant wastefulness, there is an inherent conservatism because if it expects to survive (so to speak) Nature is only free to tinker with *extra* stuff. Nature must obey what has been referred to as a (sort of) law of “interim viability” (cf. Page, 2011, p. 124). So while redundant systems are instilling resilience, they are also providing this space for experimentation. Research demonstrates that “leaps” in evolution tend to be traceable to accidental copies of genetic material that eventually mutates, gets activated, then expressed. As a clever moniker for this idea, Gould and Lewontin (1979) coined “spandrels” for what Gould (1997) later suggests “arise nonadaptively as secondary consequences... but then become available for later cooptation to useful function in the subsequent history of an evolutionary lineage” (Gould, 1997, p. 10750). This is very likely the only way evolution can “leap.” The conservative exploitation of (near) redundancy is a way so-called adaptation can be facilitated in systems that are *already* resilient (note well the reversal in the equation: adaptation is *not* “part of” resilience, instead resilience is necessary for adaptation). Since

resilience is instrumented through functional redundancy, it should not surprise us that “adaptive capacity” is also instrumented through redundancy.

Quantum resilience asserts that resilience analysis must focus on the function (productivity) of the system. Nature is a good teacher in this case. It is not the complexity of the system that augments its resilience. Instead, it is the vast amount of functional redundancy that is contained (sometimes hidden) within all that complexity that augments resilience. Recall that since there is no absolute “high mark” for resilience (it is only comparative among homologous systems), resilience characterization numbers need not be large. Instead, think of “incremental contribution” and “incremental loss” of function (in the numerator) as the driver of resilience. Complexity must be accurately modeled no matter what function it supports.

For example, if (among other ways) a mammalian cell manufactures adenosine triphosphate (ATP) from glucose molecules during glycolysis, and manufactures more ATP during oxidative phosphorylation, then the full output of ATP must be calculated and placed in the numerator, while the full complexity of the glycolysis and phosphorylation systems must be modeled and calculated for the denominator. Obviously, if the cell can “figure out” another way to manufacture ATP, those additional productivity and complexity values would contribute respectively to numerator and denominator and then resilience could be re-characterized accordingly. If the new way is comparatively more productive and comparatively less complex than the others, this would disproportionately increase resilience. Likewise if the new approach is similarly productive and similarly complex as the others, it would still (very likely) improve resilience since it adds another redundant pathway.

Reiterating the above, using biological examples as teachers, I can conclude that opportunistic exploitation of functional redundancy is how adaptive capacity is built. For example, when a gene in a germline cell is accidentally duplicated and not expressed, it is free to mutate (or not) without damage to the organism and with little additional energetic cost for its survival. This clearly adds complexity without function. This reduces resilience but, given the vast complexity of the organism, the energetic cost of maintaining the extra gene is small, so the impact to resilience would likely be unmeasurable.¹ As we know, eventually another mutation might cause the extra gene to be expressed, resulting in either deleterious functional impact (e.g., if the gene had mutated) which (in the extreme negative case) kills the phenotype, or (with a positive mutation) perhaps adds new or additional function to the phenotype. This is important because now the additional complexity has function that contributes to the resilience.

Though gaps in knowledge are closing, biology still tends to have far more complexity than known function (e.g., there exists a significant amount of so-called “junk DNA”, there are metabolic pathways that are still not understood, etc.). Again, this is because biology wastefully exploits all the resources at its disposal and is unconcerned with efficiency. As a robust theory, quantum resilience can lead us to look for function within complexity that is already seen, but not understood. If function cannot be found in the complexity, however, it still must be properly quantified in the denominator.

Generally, though without planning, if Nature has allowed complexity without function,

¹ Though recent estimates suggest *Alu* exists in over 1 million copies and comprises over 10% of the human genome, according to Maynard Smith & Szathmary (1999, p. 97): “there is an element known as *Alu*, 282 bases long, present in 300,000-500,000 copies distributed throughout the genome, and accounting for some 5 percent of the genomic DNA.... So far as is known, these *Alu* elements do nothing useful for the organism. They are only one of many kinds of repeated elements in the human genome.”

it is very likely not hurting the organism (e.g., has only a slightly negative energetic drain, as in accidental replication of genes in DNA). Systems of human design, on the other hand, are seldom so wasteful since deadweight is noticed and pruned. Still, quantum resilience drives the analysis to identify all complexity to ensure proper calculation of the denominator, so if significant complexity is detected in such a system without a corresponding provision of function, it might drive the analyst to identify forgotten function, or to acknowledge an expansion in system scope.

Since expansion in scope leads to a much larger analysis, it provides an important segue to a final comment about adaptation. Frequently the literature mistakenly assigns the term “adaptation” to systems that are experiencing upper-directed behavior (*sensu* McShea, 2012) in Simon (1962) hierarchies. That is, so-called adaptive behavior is actually being facilitated by the larger “containing” system when it has something “extra” to contribute to a child system by way of redundant capacity. This is almost certainly the intent of most of the literature that conflates resilience with adaptation. For example, with the goal of creating “resilient infrastructure”, humans plan for exigencies and stockpile resources (sometimes human resources) to enable swift recovery and restoration of services. Typically, this will be extolled as resilient infrastructure when in reality it is resilient human society that has learned to expect certain things and plan ahead. As usual, Nature is a good example of this approach (though without the foresight). Though it is not usually stated this way, Nature’s adaptive capacity stems from its access to (1) effectively infinite resources in the environment (at least until recently!), (2) effectively infinite time to employ them in experimentation, and importantly, (3) effectively no constraints on how trial and error adaptation is implemented. That is, Nature does not

care if it breaks itself while trying something new because nothing in Nature whines about failures in service provisioning and demands amelioration. Nature simply migrates or dies. In general, humans have learned from Nature, and only recently have we started encountering constraints that force us to rethink our previously adaptive behaviors (i.e., over-exploiting and then moving on). In this regard, our consciences make us far less profligate wasters of resources than is Nature.

To resist the trend of merging ideas like resilience and adaptation it is helpful to return to the principle that suggests resilience is about a function before it is about a system. Remember that all our observations of valued function are value laden. This implies that any idea of *future* goals would also be value laden. Recall that *value laden* is not necessarily *wrong* but we must admit to it and address it. Planning or designing “adaptive behavior” has proven difficult not only because we do not know the shape of the future landscape, but also because we do not necessarily know what future generations may need. What we *can* do is develop redundant and excess capacity that delivers the functions we currently view as important so that in the future we might have some excess capacity to exploit as values change and new systems become required. Forays into problem solving with artificial neural networks (ANNs) have demonstrated at least one thing: organization, programming, and planning only go so far in problem solving; *more* neurons and *more* connections contribute disproportionately to resolving such problems. Until we quantify valued functions, we cannot hope to deploy the redundancy we need in order to establish real adaptive capacity. Unfortunately, when the literature refers to adaptive capacity it is highly euphemistic and fails to provide examples of how it might work. Fortunately, quantum resilience provides a way forward.

Quantum resilience suggests adaptive capacity is instrumented through functional redundancy. Migrating excess capacity into alternative uses is a matter for differentiated research, but should be maintained as a separate discipline.

Resilience is Resilience

Gunderson and Holling have supervised an exceptional body of work on resilience (e.g., Gunderson, Holling & Light, 1995; Gunderson & Holling, 2002; Gunderson & Pritchard, 2002; Gunderson, Allen & Holling, 2010). In general, these edited volumes target ecological systems and avoid focus on systems that involve significant levels of human engineering. To support this differentiation, Holling has coined the terms “ecological resilience” and “engineering resilience” and suggested they are fundamentally different. In-depth discussion of what Holling (1996) originally termed “engineering resilience” does not generally occur since few engineers are involved in the production of these edited volumes and related resilience literature. Engineers, however, might have difficulty appreciating the “fundamental” difference expressed by Gunderson, Holling, Pritchard, and Peterson as “essentially between a focus on maintaining *efficiency* of function (engineering resilience) and a focus on maintaining *existence* of function (ecological resilience)” (2002, p. 5).

While I applaud this mention of function in a discussion of resilience (it receives too little), it would be difficult for an engineer to separate these and think that somehow a necessary function could continue to be efficient if it were to disappear. Non-existent functions constitute system failure. Efficiency seems to demand existence. In fact, engineers generally go to great lengths to instrument functions with a variety of

performance levels since degradation is expected over time and under a variety of environmental insults and operational loads. Gunderson et al. go on to describe the difference between ecological and engineering resilience as “so fundamental that they can become alternative paradigms in which subscribers *dwell on received wisdom rather than the reality of nature*” (Gunderson et al., 2002, p.5, emphasis added).

The implication here is that engineers operate in a rote manner, ignorantly following the received wisdom in their handbooks, never being required to interact with the exigencies of the natural world, and blindly ignoring reality while they deploy their solutions. It is admittedly difficult to swallow such an allegation. While “received wisdom” *can* blind engineers to important nuance required in novel designs, it more often allows many years of hard-won experience to be leveraged. It is cavalier to assume that “received wisdom” is not at least in part based on the “reality of nature.” After all, these ecologists are themselves documenting their observations in order that future ecologists can receive and benefit from their wisdom. The caricature of the ignorant engineer with his head down banging out products for some alternative reality is unwarranted.

When “engineering resilience” is used in the literature, it is generally assumed that the explanation given by Holling (1996) was complete and correct and needed no further explication or defense. It is also apparent that “engineering resilience” is viewed as incapable of contributing to the management of the current array of problems facing humankind—those that are usually associated with complex socio-ecological systems (SESS). With little overstatement, it is not difficult to conclude that “engineering resilience” has been relegated to step-child status—it is that *other* resilience that must be

“bettered” if we are to successfully deal with hard problems. This is demonstrated in the following passage:

The existence, or at least the importance, of multiple or single stable states determines the appropriateness of an engineering or ecological approach to resilience. If it is assumed that only one stable state exists or can be designed to exist, then the only possible definition and measures for resilience are near-equilibrium ones—such as characteristic return time. And that is certainly consistent with the engineer’s desire to make things work—and not to intentionally make things that break down or suddenly shift their behavior. But nature and human society are different (Gunderson, Holling, Pritchard & Peterson, 2002, p. 6).

Beyond the overly simplistic rule that begins the paragraph (can resilience really boil down to whether or not there is one or several stable states?), there are at least two implications in this paragraph that are worthy of comment. The first is that engineers are never required to manage multiple stable states in their systems. This view might simply reflect the scales at which ecologists think engineers work (e.g., they might think engineers build “bridges” but not “transportation systems,” or they think engineers develop “microprocessors” but not “telephony networks”, or engineers build “radios” but not “global military command and control infrastructures”). It also disingenuously implies that ecologists *do* seek to simultaneously manage multiple stability domains. However, even if ecologists *recognize* multiple stable states like “clear” and “turbid” in a lake ecosystem, they must admit they only seek to manage *toward* the one with higher social utility. This would be analogous to engineers recognizing stability domains of “working” and “broken” and attempting to keep the system “working”. While they recognize the “turbid” state, it is never the intent of ecologists to ensure a lake remains “turbid.” Understanding the turbid state (just as engineers want to understand the

“broken” state) is important, but management efforts must proceed in a certain direction and that is usually toward the state of higher social utility.

The second implication is more interesting because it hints that while employing the (allegedly more powerful) tools of “ecological resilience,” the SES manager’s goal is *not* to “make things work”, or to prevent sudden shifts in system behavior. Apparently, they want to avoid that since after all, that is what *engineers* do! This is also disingenuous. Nature and human society may well be “different” as these scholars suggest, but when ecologists step into SES management roles they must admit that their clear intent is to “make things work.” They may have normative conservationist goals, but they obviously have a “working” outcome in mind. Resilience practitioners should not lose sight of their observer-dependent and value-laden goals.

But the denigration of the engineering approach continues as they attempt to further differentiate the ecological resilience approach as a novelty:

If there are multiple equilibria, in which direction should the finger on the invisible hand of Adam Smith point? If there is more than one objective function, where does the engineer search for optimal designs? In such a context, a near equilibrium approach is myopic. Attention should shift to determining the constructive role of instability in maintaining diversity and persistence and to management designs that maintain ecosystem function despite unexpected disturbances. Such designs maintain or expand the ecological resilience of those ecological “services” that invisibly provide the foundations for sustaining economic activity and human society (Gunderson et al., 2002, p. 8).

Once again, while I applaud the reference to maintaining ecosystem function (and lament that it is so quickly forgotten), there is much that must be rebutted in this paragraph. First, “myopic” engineers have long worked with multiple equilibria (Rapoport, 1986, p. 67ff) and multiple objective functions (Rapoport, 1986, pp. 195,

202ff; ReVelle, Whitlatch & Wright, 2004, pp. 121ff), and multi-criteria decision analysis (MCDA) is a field with its own journal (cf. Wiley InterScience *Journal of Multi-Criteria Decision Analysis*), so whether or not fault can be found with their approaches, the complexity of conflicting goals is not new to engineers. Literature about multiple objective functions does not imply that the problems are easily managed, but engineers have a long history of wrestling with multiple objectives and finding the non-inferior operational “sweet spots” that can be afforded or accomplished with the limited resources available.

More importantly, what can it possibly mean to determine “the constructive role of instability in maintaining diversity and persistence”? Even speaking in evolutionary terms, the “role of instability” is *not* constructive. At best, such language could be construed to be reminiscent of Schumpeter’s (1950) “creative destruction” as long as we understand that even his use of that idea was as a rhetorical device. Instability neither *maintains* nor *leads to* diversity. Even if the authors are envisioning “random mutation” followed by “natural selection” as a kind of “instability” that can lead to diversity, then they have forgotten two things. First, it must be remembered that evolution is largely conservative or ecosystems and the life they support would be too fragile to be sustained. It is difficult to imagine a successful ecosystem that “persists” in “instability” as they suggest. The terms are opposing—and it is hoped these authors are attempting more than clever rhetoric. While it is true that some systems “flip,” the unstable period during the “flip” is not a desired “state” and neither is it enduring (or “flip” would not be used to describe it). Second, the goal of pursuing resilience at all is to manage an SES toward some higher social utility while balancing Nature’s needs for the long term benefit of

both humans and Nature (maximizing *enviro-social utility*). Since that is the case, it must not be forgotten that humans (and other species) desire *stable* systems that deliver their valued functions over appreciable timeframes. Most species cannot survive a mostly “broken” system while natural selection takes its time coming to a “resilient” solution. (See the later discussion of Holling’s adaptive cycle).

Gunderson et al.’s (2002) internal conflict is highlighted when one page later they write: “Ecosystems are resilient when ecological interactions reinforce one another and *dampen disruptions*” (p. 9, emphasis added). Here, the authors negate their previous statement about the “constructive role of instability in maintaining diversity” and it must be observed that we cannot have it both ways. Does *instability* lead to resilience as suggested on page 8? Or does resilience require such disruptions to be *dampened* as indicated on page 9? The authors know the answer, but in an attempt to define a new kind of resilience, they have lost sight of the value of a cohesive approach—one which *includes* so-called engineering resilience.

The resilience literature generally focuses on what authors imply are the more complex socio-ecological systems and the management problems associated with such systems. They suggest such systems demand treatment with *ecological resilience* in mind. While most examples are fairly tightly defined ecosystems, these volumes do occasionally foray into social organization, politics, governance, and economics. Unfortunately, they fail to acknowledge that these are *all* human-engineered systems which quite feasibly can be addressed using so-called *engineering resilience*. To know that economies and governments are engineered by humans and to not at least demonstrate why so-called engineering resilience does *not* apply seems a significant

oversight. Blanket and broad introductory remarks like those cited above are insufficient to convince practicing engineers of the need to relinquish their “received wisdom” and abandon engineering resilience as a viable tool. Fortunately, I will demonstrate that the confusion spawned by the unnecessary distinction between engineering and ecological resilience is quite avoidable. There is only *one* “resilience.”

Abandoning multiple kinds of resilience

Since humankind has only been environmentally conscious for a couple generations, it is only recently that we have recognized the need to manage the ecological systems on which we depend. It is perhaps obvious that, since they are unconscious of their own existence, natural systems are extravagant and wasteful, unconcerned by the resource cost required to maintain themselves. Before the advent of the environmental discourse, humans themselves tended to assume the earth was big enough to not be impacted by a few “wasters” and we (perhaps excusably) adopted Nature’s approach of using everything at our disposal with little concern for dwindling resources. Of course, now that population and technology have multiplied our impact, the story is changing. We have noticed and have started feeling the resource crunch—as many species before us have, only to migrate or die. We have recognized our more negative impacts and realize that a natural response (i.e., migrate or die) is less feasible now that we have filled all the easily accessed parts of the Earth. In important attempts to come to terms with these changes, would-be managers of socio-ecological systems are feeling constraints that engineers have faced for a long time. Ironically, these SES managers have lately realized that “efficiency”—understood as appropriate allocation and consumption of resources,

and thought to be a hallmark of “engineering resilience”—might actually be an important step beyond focusing merely on the “existence” of the functions they value.

Unfortunately, in isolating “ecological resilience” from “engineering resilience,” the literature implies that “engineering problems” are somehow simpler, better bounded, and better scoped, than so-called socio-ecological problems. In fact, engineers have long been forced to deliver reliably functioning systems in the face of significant constraints, and engineering managers faced economic “realities” far earlier than their ecological counterparts. To manage such constraints they have adopted design patterns and principles (“received wisdom”) that have proven to be effective over time. For example, redundancy and distribution are employed when system function is critical enough for it to be cost-effectively implemented. For this reason, it is unfortunate that “engineering resilience” is not being explored as part of the solution space for SES management.

Interestingly just a few pages after they imply that engineering resilience is inadequate to the task, Gunderson et al. add:

We propose that the resilience of ecological processes, and therefore the ecosystems they maintain, depends upon the distribution of functional groups within and across scales.... Across-scale resilience is produced by the replication of process at different scales. The apparent redundancy of similar functions replicated at different scales adds resilience to an ecosystem (Gunderson, Holling, Pritchard & Peterson, 2002, p. 10).

Such functional redundancy at different scales has been a hallmark of robust engineered systems for many years—largely because engineers noticed that it was effective in Nature. Several years prior to this Peterson, Allen and Holling (2010 [1998]) had suggested:

It is difficult to envision how ecosystems without redundancy could continue to persist in the face of disturbance. We assume that since no

species are identical, redundancy does not reside in groups of species, but rather it emerges from interactions of species (p. 175).

Though the second sentence is somewhat unclear (and by using the term “emerges” even implies some mystery) they later clarify that “overlap in ecological function leads to ecological redundancy” (p. 176). The “interactions of species” they specify is instrumented specifically by the redundant delivery of important and valued functions. In these passages, the authors properly notice that resilience is instrumented through redundancy and diversity (what biologists call “degeneracy”, or, redundancy of function through different structure), and in doing so, provide a strong testimony in support of the power of engineering resilience which employs these very tactics.

Ultimately, the *equivalence* of engineering and ecological resilience is made completely clear when Holling, Gunderson, and Peterson (2002) speak of “imbricated redundancy” (p. 85). This within-scale and between-scale (sometimes termed “across-scale”) redundancy is exactly the kind of redundancy that system engineers use to ensure robust systems from telecommunications, to health care claims processing, to space-based surveillance. In fact, redundancy that is provided by diversity or portfolio approaches would be done more if it were not so expensive. For example, telecommunications networks have more than one way to connect a call; there is not just one kind of surveillance satellite (and neither do they all observe the same regions in the electromagnetic spectrum); and health plans do not have just one way of processing claims, but do them with a wide-variety of rules, in batches, one-by-one, manually, etc. Interestingly, these three exemplar domains are particularly rich with examples of resilience (through redundancy) because they are deemed vitally important and are well-

funded. It appears that what ecologists have discovered is exactly what engineers have known for a long time: that it requires a lot of money to solve hard problems.

In engineered systems, the lowest level (smallest scale) sensors like thermistors or current monitors are frequently redundant because they are important to system health and status monitoring as well as being relatively inexpensive. Larger scale systems that contain these redundant sensors (e.g., power supplies, processors, telecommunication switches) are frequently redundant because they are deemed mission-critical enough to invest extra resources (e.g., mass, parts, labor, development and test schedule) and because at this scale they are still *relatively* inexpensive. At the largest scales (e.g., computers, disk farms, satellites) the systems are frequently redundant when replacement or repair time would significantly impact the mission or the revenue stream. For example, in addition to the full constellation of 66 satellites, the Iridium global telephony network (*iridium.com*) was deployed with six on-orbit spare satellites, not to mention redundant ground-based tracking stations, mission control centers, and telephone network gateways—recall, telecommunications systems are well-funded and customers demand they function, so resilience is ensured through engineering redundancy.

Unfortunately, Gunderson et al. (2002) appear to be unaware that engineers frequently provide resilience through such portfolio approaches. Rightfully, they say “within-scale resilience complements cross-scale resilience” (p. 11), but then they specifically deny the existence of the engineer’s portfolio approach:

The consequence of all that variety is that the species combine to form an overlapping set of reinforcing influences that are less like the redundancy of engineered devices and more like portfolio diversity strategies of investors (p. 11).

It is apparent that at least part of their problem stems from thinking that engineers design “devices” but not systems. This is a debilitating and inexcusable oversight. By reinforcing the idea that resilience can be operationalized through redundancy, they have *equated* ecological resilience and engineering resilience even while denigrating the latter. This is an important lesson that future SES managers must learn—and can avoid learning the hard way. So the question is why was a resilience dichotomy ever created?

It is possible that some fear so-called engineering resilience because it may lead to “deconstructionist science,” a command-control approach to management, or simple reductionism. However, there is simply no need to fear that since it was Nature that taught engineers how redundancy is the path to resilience—and Nature did it without command and control hierarchies. As a particular way to implement a management plan, command-control can be effective and has its place. It is certainly a way to provide humans a sense that they are in control—often an important outcome in itself. Since the eco-resilience literature specifically repudiates command and control approaches (Holling & Meffe, 1996), ecologists call themselves “managers” of SESs, but the feint is obvious. Their goal is control as they decide on appropriate outcomes and shepherd the system in what they perceive is the right direction. But it matters very little what humans call themselves, as Nature has modeled, and as engineering has followed, resilience will be implemented through redundancy. It is in the “doing” not the “defining” that a real definition emerges.

Peterson, Allen and Holling (2010[1998]) suggest:

Ecological and engineering resilience reflect different properties. Ecological resilience concentrates on the ability of a set of mutually reinforcing structures and processes to persist. It allows ecologists or managers to focus

upon transitions between definable states, defined by sets of organizing processes and structures, and the likelihood of such occurrence. Engineering resilience on the other hand concentrates on conditions near a steady state where transient measurements of rate of return are made following small disturbances. Engineering resilience focuses upon small portions of a system's stability landscape, whereas ecological resilience focuses upon its contours. Engineering resilience does not help assess either the response of a system to large perturbations or when gradual changes in a system's stability landscape may cause the system to move from one stability domain to another, for these reasons we concentrate on ecological resilience (p. 179).

Unfortunately, the comparative statements made in this passage are not mutually exclusive and are largely meaningless, so it is difficult to understand the launching point of their comparison. That one kind of resilience focuses on "stability landscape" and the other on "contours" seems to be arbitrary employment of metaphor. Until the poetry is replaced with meaningful and actionable ideas, it is unhelpful. It also seems strange that engineering resilience helps with neither "large perturbations" nor "gradual changes" and forces one to wonder how engineers have accomplished anything at all with their meager tools. This is odd, especially since engineers use the same terminology (systems, structures, states, processes, transitions, stability, etc.) and very likely used it first.

It is for these reasons that I find the unnecessary distinction between engineering resilience and ecological resilience to be confusing and harmful to the discourse. This is especially true when tacit denigration of engineering resilience is nearly immediately followed by glowing testament to the power and importance of what is easily demonstrated to be the tools of engineering resilience. Engineers (perhaps indirectly or by osmosis) have learned the power of redundancy from the natural systems they observe. There is no reason that "engineering resilience" cannot be a powerful contributor to managing the complex socio-ecological systems we face today. It would be

far better to focus on operationalizing “resilience” than to perpetuate a confusing and meaningless distinction. Fortunately, when arbitrary distinctions and misunderstandings are clarified, resilience is revealed as a unified concept. Ecosystems are resilient due to their many functional redundancies and engineers have learned from this. Redundancy in the delivery of valued functions is exactly the engineering approach.

Resilience and the Adaptive Cycle

Leveraging a traditional view of ecosystems which are observed to have periods of exploitation followed by conservation, Holling proposed the adaptive cycle as a model of how natural systems continually reformulate themselves (cf. Gunderson & Holling, 2002, p. 25ff.). His model has four phases. In their mature conservative phase (K), ecosystems employ the resources available to them and deliver useful functions and services. During this phase, Holling suggests the resilience of the system is in decline. Ultimately, resources dwindle and resilience wanes leading to release and system collapse (Ω). Thereafter, Nature recovers and experiences a reconfiguration phase (α) which results in a new system. This leads to an exploitation phase (r) where the system builds to maturity while resilience is high. The cycle, Holling suggests, repeats indefinitely.

Holling’s adaptive cycle may be an important contribution, but the way it is employed in SES literature is sometimes forced to the extent that Holling himself might quibble about the way his model has been assumed to be a proven and working theory of the natural world instead of being merely an analogy. It is my purpose here to demonstrate that it must not be uncritically employed—and certainly not for systems with

significant human design and involvement. In fact, while some of its tenets are applicable, the entire model must be employed only with extreme caution. Note well, *resilience characterization should not start with analogy, but with system analysis.*

As a model of how natural systems might adapt over time (or “flip”, or “enter a new stability regime”, or “find a new basin of attraction”, the euphemisms abound), Holling’s adaptive cycle has been frequently employed in the literature and mapped to the behavior of some ecosystems. As pointed out earlier, when an ecosystem “flips,” organisms in Nature do not start whining about failures in the provisioning of services or valued functions, they simply migrate or die. If, however, the adaptive cycle were to be representative of systems engineered by humans, we would call such systems “broken” (and there would be significant whining). In fact, human-engineered systems that experience “collapse” and “reorganization” are repudiated. The goal of human engineered systems is to arrive quickly at productive stability (Holling’s *K* phase) and remain there for as long as possible.

Since humans have trouble using constantly changing systems, systems of human engineering (including SESs) must exhibit *K* phase resilience. *K* phase systems are mature and conservative, but are, unfortunately, where Holling predicts resilience will be declining (Gunderson & Holling, 2002, p. 44). If “bouncing back” (to employ a frequently seen euphemism) from perturbations is important, it is to the well-understood stability of the *K* phase humans wish to return, not to the frenetic and still-coalescing system of the *r* phase. Humans cannot effectively operate in the constantly changing but (allegedly) highly resilient space of the α and *r* phases. Humans need stable systems that provide stable services through which work can be accomplished. Picture, for example,

humans trying to flourish in the face of constantly changing economic or political regimes—history demonstrates our inability to do this. So it appears that humans need high resilience in the *K* phase at the very time Holling suggests it is waning. Hence, the adaptive cycle might reflect the operation of pristine Nature, but is clearly not a model of the way human systems are most effectively managed. Unfortunately, overemphasis of the adaptive cycle analogy has frequently engendered adherence to it as if it were an inexorable law of the universe.

The conflict between this presumed law of the universe and human needs is recognized by Baskerville (1995):

The important point here is that the management of a natural system such as a forest is an attempt to *prevent* the cycle of exploitation/conservation/creative destruction/ mobilization from *operating normally*. Indeed, if the cycle did *not* operate, there would be no need for society to manage those systems. In general, management is invoked to *prevent* the creative destruction step... (p. 92, emphasis added).

Later, he suggests humankind is on a “treadmill” trying to prevent the impending creative destruction. The tendency to see the cycle as inevitable—even in the face of management—is clear. Walker and Salt (2006) agree with Baskerville and brace themselves for the inevitable:

No system can stay in, or be kept in, a late conservation phase indefinitely.... A significant back loop [i.e., collapse and reorganization] of one form or another is inevitable (p. 85).

This self-conscious adherence to the model is odd in light of the many equivocations Walker and Salt (2006) make. They allow for “many variations” (p. 82) and draw the model in a way that suggests a system can find its way around the four identifiable phases in a fairly uncontrolled manner (p. 83). Engineers would conclude that

a model so fluid and free has lost its ability to be useful—let alone inevitable—but ecologists cling to this model despite its inability to model reality and make predictions. Unfortunately, this resignation has prevented movement toward operationalizing resilience and instead has put us on Baskerville’s treadmill, bracing ourselves for the inevitable.

Because of this, it is difficult to envision how the adaptive cycle can be used as a model for socio-ecological systems where human influence and technology is dominant. In these cases, perhaps the “adaptive cycle” is best employed as a non-cycle or a one-time process that illustrates system growth from r to K to Ω and does not attempt to suggest systems arise from their ashes like a mythical Phoenix. While it is an incorrect application of the model to suggest a system’s “resilience” is demonstrated by repeated passes through the α - r - K - Ω cycle (recall that according to Holling resilience waxes and wanes *within* the cycle), it is equally problematic to suggest the *same* system emerges after the Ω phase. In fact, in any collapse and reorganization, it becomes clear that a new system emerges. While this is hinted at by Walker and Salt (2006, p. 75), the best expression that it is really a new system is provided by Holling, Gunderson and Peterson (2002, p. 75) when they demonstrate how the adaptive cycle can be stacked and used to reflect up and down the hierarchy in a Panarchy of systems. The very introduction of the Panarchy idea exposes the deficiencies in the adaptive cycle because it demonstrates how a system either escapes from the cycle or is really just part of a larger system. Since systems-of-systems ideas are well-worn in the engineering space, a new and confusing concept like Panarchy is not really necessary.

What *is* possible (and is at least implied by the Panarchy idea) is that the entity that is progressing from r to Ω is a *subsystem* of a potentially resilient larger system. This is vital but not particularly clear in the literature. For example, when a human-engineered system like a government collapses, if the residual human society creates a new government, it is clearly *not* the same government system as the one that collapsed. That is, it is not the “government” that is resilient, but the human society which is the parent system. This is vital, because the human society draws on far more resources to recreate the government than the previous government left behind. This cannot be effectively shown with the adaptive cycle and is not sufficiently clear in the presentation of the Panarchy concept.

The adaptive cycle and Panarchies are tricky things since they tend to imply the same system is (re)cycling. Students struggling with the concept say they want their business to be resilient like the adaptive cycle, but are startled when asked how frequently they plan to schedule a business collapse. As expected, nobody wants their business to ever collapse. Instead, they eventually agree that they want to have a business model that supports a variety of projects that come and go as necessary, redundantly generating income to keep the business viable as the market evolves. When we talk instead about specific business thrusts that come and go, they can make better application of a (partial) adaptive cycle analogy (specifically the r and K phases). This also drives home the important fact that such “rebirth” must be greatly subsidized by the parent system. The assumption that a “system” simply reorganizes its mass and is reborn trivializes the important concept of system identity. There is more to identity than mass

and constituent particles (or people and buildings, etc.). As Ahl and Allen (1996) put it “from the perspective of the old system, the new system is a mystery” (p. 183).

This is illustrated in the conflict between resilience and the adaptive cycle exhibited in Carpenter et al. (2001, pp. 770-771). The authors employ Holling’s adaptive cycle to discuss the resilience of a lake-agriculture system. They suggest there are two parts of their SES, the social (which is comprised of human activities like agricultural production and regulatory action) and the biophysical (which consists of the lake’s water quality). This particular SES is managed to maximize (1) the net social utility of the agricultural production and (2) the freshwater use in recreation. Carpenter et al. suggest that “resilience of the clear-water state [is] high during the *r* phase” which follows Holling’s suggestion, but the claim is difficult to justify since they are describing a build-up of Phosphorous in the surrounding soil. In fact, in specifically physical system terms, what is occurring is the *degradation of the water quality* as P increasingly leaches into the lake. By forcing use of the analogy they have forced themselves to suggest something that has not yet been demonstrated. That is, they have done nothing to quantify resilience. There are at least two issues with this approach. First, they have not adequately defined the system in the *r* phase. In fact, it appears they want to map the increase in P with the upward trajectory of the adaptive cycle in the *r* phase. This mapping might be irresistible, but unfortunately it leads to the second problem of making a single variable a proxy for resilience. Since P is increasing (and is, to them, inversely proportional to resilience), they cannot conclude the resilience of the clear water state is high. In fact, what they must say is that water quality (their valued function) is *decreasing* along with resilience. For their mapping to work, resilience must be decreasing in the *r* phase. Unfortunately, the

adaptive cycle forces them to await the *K* phase to allow for decreasing resilience. So, either they must better define what their resilient *r* phase consists of, or they need to acknowledge they are attempting to conflate increasing P with decreasing resilience and force-fit an analogy.

There is danger in using one measurable value as a proxy for resilience of a system. Carpenter et al. (2001) suggest “*indicators of resilience* that are appropriate for the current regime may become useless as ecological structures and social expectations shift” (p. 779, emphasis added). While this is superficially true (e.g., a hypothetical new regime may contain *no* Phosphorus and this would eliminate their ability to use P level as a proxy for resilience) it is also problematic since it ignores the overall system (and contradicts any acknowledgement of its complexity). The most graphic illustration of the confusion is in the chart on page 773. Figure 4 indicates it is plotting the resilience of the clear water state. But this simply cannot be—especially if the graphical portrayal of four Holling cycles depicted beneath it is accurate. This would require that resilience should wax and wane *at least four times* in this figure. Instead the chart *may* depict the *turbidity* of the water (or a notional representation of overall water *quality*), but it has nothing to do with its resilience—and the cycles shown are largely meaningless. All we can glean from the discussion is that higher P levels cause water turbidity which is bad for the recreation business on the water. Overall system resilience has not been addressed and it certainly has not been quantified.

In discussing one of the cycles, Carpenter et al. (2001, p. 773) say “soil P levels in the watershed began to stabilize or decline.... This change has increased the resilience of the clear water state.” But this is not completely accurate. Instead, the change may have

reduced the turbidity of the water, but no quantification of resilience was done.

Obviously, their system includes more than the water and this must be acknowledged in the characterization of resilience. It should also be clear that we cannot change (or plot) the *resilience* of a “*state*” as these authors indicate. Instead we can only speak of the resilience of a system. Even if it is granted that one could characterize the resilience of a system *operating in* a particular state, it still requires that *all* the state variables be employed in identifying that configuration. Most systems can exist in many states and as a system concept resilience should apply to any and all states in which a system may exist. Later they conclude:

In the case of lake eutrophication, such indicators include soil P concentration, animal stocking densities, and land area under construction, which are inversely related to the resilience of the clear water state (Carpenter et al., 2001, p. 774).

This belated recognition of other contributors to “resilience” indicates the oversimplification previously discussed (and points to other, unquantified, valued function). Further, as “indicators” in the system, each of these contributes to the definition and identity of the overall system. They are not, however, indicators of resilience and since there has been no formulation of resilience, they cannot be assumed to be “inversely related” to resilience. Instead, these must contribute to or detract from whatever valued function has been defined for the system. What this simple statement demonstrates is the need for a system-oriented analysis. No conclusions can be drawn until such an analysis occurs. As they allude when they mention “the clear water state”, the items they list are ultimately related to the *clarity of the water*, not the resilience of the system. These are very different things.

Ultimately, Carpenter et al. contend that the increasing P-concentration leaves the lake in a turbid state, triggering the Ω phase collapse that spurs “strong regulation of agriculture and a tumult of confusion, debate, and evaluation of the problem” (p. 770). They suggest, the outcome of regulatory activities serves to clear the water sufficiently to restore positive net social utility. Without even getting into the way they have casually expanded their system to include regulatory bodies, recall that this activity only superficially increases resilience of the system (because water clarity is the only measure). Afterward, Carpenter et al. suggest, the regulatory system relaxes and the lake returns to the turbid state forcing another reactionary regulatory cycle. Even if hypothetical, this is a poorly managed system, not an example of resilience. At best, according to *their* definition of resilience, it is a reminder that the lake system is highly sensitive to perturbations and hence *not* resilient. Obviously, a more effective regulatory regime would not simply relax and allow the same problem to occur again. Carpenter et al. hint at this by closing the discussion with:

Active adaptive management leads to low-amplitude cycles that can keep the system in the clear water–high utility attractor for longer than laissez-faire management (p. 771).

In fact, their concession that “low-amplitude cycles” might be achievable, does not mean they are necessarily desirable. As I contend above, humans do not want obvious cycles in their managed socio-ecological systems. How much “tumult of confusion” can humans tolerate? We desire our systems to be mature *and* stable. Unfortunately this occurs (according to Holling) where resilience is waning. That cannot be said to be successful management.

In an attempt to employ the adaptive cycle, Carpenter et al. have assumed single measurable variables (P) are direct indicators of resilience and they have attempted to force the analogy to fit the gradual increases and decreases in this single measure. At no time is the resilience of the system characterized. In arguing for a management plan that allows “low-amplitude cycles” they seem to be suggesting that Holling’s cycle can operate without its Ω and α phases and hover in a shallow but continuous r-K cycle. As previously stated, this is closer to what humans need, but it does not match the cycle they have tried to force. This clearly exemplifies why the adaptive cycle should not be overused for SES management. Holling himself warned about its adoption as a theory or rigid model (Gunderson & Holling, 2002, p. 399). We must heed that warning.

Assuming there is an adaptive cycle frequently causes us to forget about the system and leads us to oversimplify and analogize our system in order to fit the analogy of the cycle. As illustrated, this too quickly leads to finding proxies for resilience and leads to incomplete analysis. This is not a way forward for systems research or resilience. This is especially true since the adaptive cycle can only predict future collapse and mysterious reorganization—something that is not useful when humans require stable function from their systems.

Given that the adaptive cycle is based on Schumpeterian “creative destruction,” it seems strange we so quickly forget the very reason for Schumpeter’s original idea. In fact, the seemingly random contributions of technology are what drove Schumpeter’s observation of the trend to randomly experience a regime shift that changes everything (cf. Allenby, 2012, p. 172). Technology is based on the human contribution and humans are the only species to supersede genetic evolution with memetic evolution. Because of

this (and with a few important examples), it seems strange that we would expect any kind of observable cycle. In fact, most agree that the earth system is engineered and sometimes recursively—or ironically—refer to it as “terraformed” as inspired by science fiction author Williamson as early as 1942 (cf. Williamson, 1951, 2001). Hence, it seems disingenuous to strictly expect a Holling adaptive cycle to apply even in the situations where the SES is *mostly* natural. Given *any* human input (engineering, design, management, participation, use, harvesting, etc.) we must concede that the system is “human” and thus subject to our manipulations and meme-based evolution instead of a purely natural gene-based evolution. This would imply that at any time in the so-called cycle, everything could change. This is a far more realistic model—and is supported by history. Unfortunately, it is not very illuminating. Given that, it is better to stick to actual system analysis instead of analogical reasoning.

Enthusiasm for the adaptive cycle requires moderation. In fact, it is likely fair to ask if it only applies after the fact. It seems that we only notice cycles retrospectively, and if its explanatory power is only in its ability to either model the past or predict eventual collapse, it seems limited. For example, it seems unhelpful to be forced to view an operational SES as currently in some kind of *K* phase on a trajectory toward collapse, but have no way to measure resilience either now or at some point in the future. Hence, as an analogy it might work, but it cannot really approach “theory” status.

Fortunately, an approach like quantum resilience provides a way forward and allows significantly moderated use of the adaptive cycle. Metaphors and analogies can be helpful, but they should not substitute for system analysis. A start at a proper assessment of the Carpenter et al. lake-agriculture system is provided as a brief example below.

Resilience and the Urban Space

Vale and Campanella (2005, p. 3) remind us that throughout history few cities have disappeared. Even after facing ruinous natural disasters, the vast majority are rebuilt, returning to viability and then vitality. Does this imply cities are simply *de facto* resilient? If this is the case, then cities must be studied so humankind can learn the lessons of resilience. Instead, however, it seems that much scholarship is invested in answering the question of how we can ensure our cities are resilient. Despite humankind's history of near total success in establishing resilient cities, it seems our approach to urbanization has come under question and can no longer be trusted. Are we afraid that we have reached a stage in human evolution at which cities will no longer be resilient? Is it because we specifically recognize that "sustainability" demands we adopt a new approach for which we have no practice? Has the promise of future sustainability clouded our judgment to the extent that we fail to look at past successes for clues about future potential?

There are at least several reasons this may have become an important research topic. First, it is to be expected given the pervasiveness of the sustainability discourse and the importance of cities. At significant cost and risk, humankind is planning for a sustainable future—and it seems fair that cities *must* be involved in contributing their part. Alternatively, fears have arisen as we have witnessed recent bankruptcies of long-important cities that have left their futures uncertain. The increasing threat and incidence of terrorism has also caused much concern, promoting discussion of protection and recovery. For whatever reason, urban resilience is vigorously trumpeted in the literature.

It must be considered, however, that resilience might be meaningless at the level of a city. Certainly the principles of quantum resilience call into question the rationality of such an exploration. When considering, for example, the *function* of a city, it taxes the imagination to see it as anything but a huge list, and certainly smacks of impossibility to collect it in a single analysis model. Instead, it seems more likely that urban resilience is most meaningful if targeted at smaller “urban infrastructure systems” and the functions they provide.

Specifically, the “function” of a city stands at a different level than its robustness against, for example, terror threats, or hurricanes can impact. A city tends to frame itself in cultural and symbolic terms. For example, New York is “the economic capital of the world” or, “a symbol of freedom for huddled masses yearning to breathe free”, or “the city that never sleeps”, etc. Such cultural and symbolic constructs are “protected” by their redundancy throughout human society, throughout the world. Such constructs are obviously not limited to NYC or US cities, but includes many cities which are cultural icons like Paris, Rome, etc.

Some discussions of resilience are more or less meaningful from a scientific or operational standpoint and the resilience of cities *qua* cities is among the less meaningful. They are, as we know, already pervasive and have long histories. In fact, cities are resilient because humans make them so, and will continue to be because humans live in and invest in them. It is a “human” thing, and, very simply, sometimes humans decide to perpetuate something. Since New York City’s value is wrapped up in the fact that it is a cultural and economic icon, *of course* we will ensure it persists. But we must not forget that everything we do there can be done somewhere else. If NYC goes away we need not

lose any important functions of society. If we value New Orleans similarly, then the same will be true. It will be made permanent by our investment in it and immigration to it.

Note, however, that as with anything, a city's value is observer-dependent. This is highly important and it is the very reason I use the word "value" in "valued function". For example, the Chinese are not likely to value New Orleans the same way American culture does, and if it were up to the Chinese, New Orleans (or particular iconic aspects of it) may cease to exist because they will neither invest in them nor move there. This illustrates how quantum resilience is more meaningful from an operational standpoint. There are functions provided by New Orleans that the Chinese *may* value, but that is not the same as valuing New Orleans for our current cultural and aesthetic reasons. And, since all those non-iconic functions can and are provided by other cities, the loss would be inconsequential when viewed through their eyes.

At some point we must admit to *valuing* NYC or New Orleans for some specific reason (like an environmentalist might value a salt marsh for purely aesthetic reasons). And the reasons can constitute a long or short list. There are many "functions" in and of a city that one could value or be interested in preserving. Quantum resilience can start wherever anyone wants in defining valued functions. If the valued function boils down to its status as a cultural icon (e.g., "the joy of New Orleans"), that is a perfectly acceptable starting point for quantum resilience. The challenge of course is finding others who agree with you and coming to some consensus on how to quantify that joy.

The fact that a "system" like a "city" makes one think of "cultural icon" as a valued function is a perfect example of how the quantum resilience approach appropriately manages the analysis. With a valued function like "city as cultural icon,"

many would enter the analysis thinking “the city” (e.g. NYC or New Orleans) is the system we want to be resilient. But quantum resilience makes it clear that is the wrong starting point. Instead, it becomes glaringly obvious that the city is *not* the system because the collection of *systems* that delivers the *function* of “cultural icon” is *vastly* different from common ideas about what a city is and does. Instead, quantum resilience makes it clear that the entire *society* might be the system in this case. For example, when disaster strikes the city, the entire society sends rebuilding teams and invests in reparations. The entire society donates to the National Red Cross and permits its government to provide disaster relief funding. The entire society tracks the progress on the national news and does not stop caring until viability and vitality return to the city. We will invest until we are reasonably assured that the iconic status has been restored. In this way, quantum resilience has revealed the real system. Resilience is not about a system until it is about a valued function.

It is also unclear that building resilient cities—however it is defined—actually saves money for a culture. For example, there is recent concern about urban resilience in the face of terrorist threats. First, quantum resilience reminds us that resilience is not disturbance-focused. Second, it is arguable that a staggeringly large investment in security, and a similarly staggering erosion of personal freedoms could prevent such attacks, but it is completely unclear that such investment would contribute to the resilience of a city. In fact, any willingness to invest in such a large project already attests to the resilience of the city because it illustrates the human interest in protecting the cultural icon. Some of the media-reported reactions to 9/11 and the Boston Marathon

bombing are informative. They have the flavor of “are you *kidding* me!? You want a piece of *me*!?” That *attitude* defines the resilience of the city.

As regrettable as the loss of human life is at any level and at any time, complete prevention of such incidents also serves to remove any measurable way of assigning dollars-per-lives-saved (i.e., if you never lose a life, the expression becomes infinite). Further it would eliminate any idea of assessing whether or not the city was actually resilient in the face of the threat (not that quantum resilience agrees with such terminology, but assuming that terminology, if there is never a test, you have no way of knowing). In the absence of measurable progress, investment is likely to cease. Further, it is abundantly clear that there is no way to stop such crime. Since the sole intent of terrorism is spreading terror, the only way to prevent it is to stop being terrified by it. American reaction to 9/11 and Boston is indicative of this posture. This leads to defining the system at a scale that is not city-oriented. Again, this is a cultural thing.

But to focus on the terrorist threat is the wrong approach when resilience is under consideration. Instead, there are arguably two approaches that can be taken. A robustness approach can start by addressing the projected impacts of a bomb and providing defenses (physical or otherwise). A resilience approach might consider the actual impacts to valued function and propose that redundancies be implemented. Recall, there is no “of what, to what” (*sensu* Carpenter et al., 2001) in resilience analysis; it must focus on function. Quantum resilience understands that many perturbations result in similar functional impact. We do not specifically care, for example, that it was a bomb, a flood, or an operator error that ruined, say, the communications infrastructure. What we care about is that the function of communication must be redundantly provisioned in order to

increase resilience. The same operational approach applies if we are worried about buildings being lost, or infrastructure damaged, or transportation disturbed, etc. Such impacts can only be mitigated through functional redundancy (additional buildings or alternative workplaces, alternative water and sewer lines, alternative transportation and communications infrastructure, etc.). The answer does not change. Resilience can be provided through expensive but purposeful redundancy.

In their edited volume *Resilient City: How Modern Cities Recover from Disaster*, Vale and Campanella (2005) offer an interesting look at the narratives, symbols, and politics of urban restoration. Though their definition of resilience differs from mine, in their conclusion (pp. 339-351), they suggest twelve “axioms” of resilience:

1. narratives of resilience are a political necessity,
2. disasters reveal the resilience of governments,
3. narratives of resilience are always contested,
4. local resilience is linked to national renewal,
5. resilience is underwritten by outsiders,
6. urban rebuilding symbolizes human resilience,
7. remembrance drives resilience,
8. resilience benefits from the inertia of prior investment,
9. resilience exploits the power of place,
10. resilience casts opportunism as opportunity,
11. resilience, like disaster, is site-specific, and
12. resilience entails more than rebuilding.

Phrases like “political,” “governments,” “national renewal,” “human resilience,” “remembrance,” and “power of place” reinforce the point that cities are cultural icons more than they are factories of function. Urban resilience is more a matter of *public will* than anything else. Perhaps the words “narrative” and “symbol” (#3 and #6) are the best indicators of this because they remind us that icons are more about *story* than actual production. Though the stories and symbols may be exaggerated, they are generally based

in fact, so after the *will* is found, there is a generally a direction in which it can be exerted.

While urban resilience is far more about *symbol* than anything else, there *are* material aspects that must ultimately be brought to fruition in order to keep the stories alive. My insistence that we focus on the *valued functions* is because that is the only *practical* way to roll-out an icon! Consider for example a culturally rich city like New Orleans with no external access (transportation infrastructure), no ability to make a hotel reservation (communications infrastructure), and no means by which to freshen-up if you ever did arrive (water and sewage infrastructure). Without these “amenities”, important culture would be wasted and quickly fade from memory.

Interestingly, Vale and Campanella did not neglect to mention the power of investment (“underwritten” in #5 and “investment” in #8). In fact, these two axioms are the only ones that can be operationalized beyond promoting a cultural *esprit de corps*! Very importantly, they mention “outsiders” (#5), which further illustrates that resilience is very often assumed at a smaller level than it is actually instrumented and operationalized (neglecting the “real” system extent). Further, when the idea of “cultural icon” extends beyond nation-state boundaries (as it does, for example, with Paris, Rome, New York City, etc.), resilience and permanence are fairly guaranteed, because now there is significantly more redundancy in those “outsiders” who are underwriting the recovery and contributing to the rebuilding.

Though they are large, there are only two requirements for a resilient city: “attraction” and “infrastructure”. First, there must be some iconic status, some mythos, some symbol that will serve to keep a city alive, first in the hearts and minds of humans,

and only then in the physical dimensions. Note well that such mythos can be simple nostalgia or a strong media campaign. Second, there must be infrastructure. Water, power, transportation, waste management, communications, businesses, buildings, homes, etc. are “life blood” items that support any and all who find themselves in the urban space. Attraction draws, infrastructure keeps. If either aspect wanes, the resilience of the city is diminished accordingly. Quantum resilience can certainly model both aspects of urban resilience, but it should be obvious that consensus among experts is far more likely to be gained when concentrating on the specifics of smaller portions of urban infrastructure than when attempting to establish the relative merits of particular cultural items.

For this reason, though an important topic, urban resilience is not one of the more meaningful scientific analyses. Instead, because we know physical and organizational infrastructure contributes directly to urban resilience, it seems better to focus system analysis and modeling efforts in these areas.

Resilience and the Safety Industry

Though this is a rather broad definition, there is a particular “angle” on resilience that is taken by researchers I position in the “safety” industry. Though their work is not particularly relevant to what is discussed herein, their literature is frequently cited and takes similar approaches to merging and conflating the idea of resilience with other ideas. I discuss this here very briefly because it presents another angle on the “time domain” questions (see further discussion below). The “safety” industry is a broad moniker I apply to systems where humans-in-the-loop and emergent events lead to active decision-making

(e.g., air traffic control, healthcare, etc.). Hollnagel, Woods and Leveson (2006) demonstrate how they are clearly positioned in the “emergency management planning and response” domain where protection of human lives is paramount (see also Hollnagel et al., 2011). Note well that where decision-making humans are in-the-loop *at a level that impacts the dynamics of the valued function delivery* (e.g., landing an airplane), just about *anything* can be called adaptive, or transformative, or learning, because humans can change their mind and “direct” the system in a different way—usually toward a *safer* outcome that prolongs the “life” of the system (so you can see why “resilience” is somewhat cognate). There are, however, important caveats and distinctions that must be discussed.

Transient “structures” (like decisions) cannot be considered “resilient” because they are ethereal and passing. There is no expectation of durability or permanence for a spur-of-the-moment decision no matter how well-reasoned and hard won. This does not mean that we want such decisions to be wrong or poorly considered, but there is no intention that the decision last any longer than it takes to be recorded in a log book (or digital equivalent)—and note that the *record* of the decision is distinct from the decision. Note that while we might legitimately want to learn from the decision-making process in order to improve our future success, this contributes more to the robustness of the system and may not contribute to its resilience. In fact, such “learning” contributes to the model that exists within the anticipatory system (see above), making the system with humans in-the-loop potentially better able to anticipate a future event and plan accordingly.

Recall that resilience characterization requires that *function* be delivered over a time frame long enough to actually value it. So, for example, do we value and

characterize resilience for an airplane? Yes. Do we characterize resilience for its particular flight path? No, because a flight path is just an option, a choice, a transient that is based on many other factors that will never be repeated. Can we value and characterize resilience for the air traffic controllers and their organizations and equipment? Yes. Can we characterize resilience for the specific directions they give that are contingent on weather and wind and other aircraft? No. We *can* characterize resilience for the ATC system (including the operators) at a given time, but we cannot (and it makes no sense to) characterize resilience for the thousands of transient events that exist “in” the system at that specific time. Recall that the goal is a resilient decision-making system, not resilient decisions. This should be obvious because decisions are transient and they are not systems.

Perhaps an easy way to distinguish these is to ask “where is that flight path now?” Or, “where are the decisions that determined those choices now?” If the answer is that they are “history,” that is, at best recorded in a log along with all the rationale that contributed to them, then they are not in the class of items for which we can characterize resilience. Sure, we can characterize the resilience of the system that does the “logging” but that is an entirely different thing. We can certainly ensure the log is “persistent,” and we might do this because we want to learn from it so we can make good, repeatable decisions in given contexts. Repeatability, however, is not permanence.

As humans we find it difficult to distinguish our “agency” from our “system-ness”. For example, humans are resilient, but it is *not* because they can change their mind. Humans are resilient because they are chock-full of redundant structures (some of which even comprise this idea of “mind”). Humans are resilient not because they can

make their human subsystems do specific things like mind-changes, but because of the way their systems are deployed. Consider that “changing one’s mind” is a *valued function* performed by the human system. Now that it is clear it is something the “human system” *does*, quantum resilience forces you to model the actual system that delivers that function. Once proper consensus has been achieved, the system can have its resilience characterized.

Woods provides a definition of resilience that demonstrates how the “safety industry” works and thinks:

Resilience/brittleness is a parameter of a system that captures how well that system can adapt to handle events that challenge the boundary conditions for that system’s operation.... The capacity to respond to challenge events resides in the expertise, strategies, tools, and plans that people in various roles can deploy to prepare for and respond to specific classes of challenge (Woods, 2009, p. 499-500).

First, it is fair to ask why “brittleness” is inserted. Presumably, this defines a logical opposite for resilience and establishes a spectrum for its values. It has already been noted, however, that such opposites can only be notional (and should probably be avoided). Second, resilience is defined as a “parameter” of a system, very casually employing a word that has a clear meaning in the systems engineering world. Again, calling resilience a parameter almost makes it sound as though it is a *setting* that can be adjusted, or at least something that can be measured. Of course, Wood offers no measures. Third, adaptation is assumed, and is later redefined as “capacity to respond” (note the positioning of the repeated phrase “events that challenge” and “challenge events”) which conflates resilience with both robustness and simple if-then response protocols. Fourth, it is unclear what “boundary conditions” actually means in this context.

Does it mean the analysis boundary? Does it mean that a system should be able to perform in an environment other than that for which it was designed? Does it mean that its performance criteria (e.g., operating tolerances) or service level agreement must be expanded? Fifth, it is made obvious that humans must be in the loop since a “capacity to respond” specifically implicates things that “people in various roles” do. Sixth, if it “resides” in human activities and expertise, “resilience” has clearly become attributable *only* to systems with humans in the loop—a very narrow definition of resilience. Finally, it is clear that “specific classes of challenge” position this idea in the robustness space and that the “parameter” Woods refers to is the humans in the loop.

Obviously, those contributing in the domain I refer to as the “safety industry” are doing fine work in improving the systems in that discipline, but it is only marginally related to resilience and would benefit from consideration of anticipatory systems theory (Rosen, 1985).

QUANTUM RESILIENCE

Since I argue strenuously that “resilience is resilience” and call attention to authors who append terms to the front of the word, it may seem odd that I would do the same. If resilience is not general, specified, ecological or engineering, why is it “quantum”? First, note that quantum resilience is a title and not a “kind” of resilience. Second, it demonstrates what I believe to be a more fundamental idea of resilience around which most of the research hovers while not specifically discovering it. Instead of the ongoing expansion of the idea, it suggests contraction to a more focused idea is required. It makes resilience something that can be understood, managed, and characterized instead of simply talked about. Quantum is a useful word to employ in expressing that fundamental idea. Third, the idea of a “quantum” suggests there is a level at which we can productively design and operationalize resilience. If left diffuse and expansively defined, there is little hope the idea can actually be put into specific use. Fourth, this forces engineers, analysts, and managers to be less casual about the concept. Quantum resilience forces *quantification* of the concept and demands rigor in its characterization. Once observed at the levels we can actually impact it, resilience is easily seen as a quantum concept. Resilience increases or decreases by the quanta of valued function delivered by participating systems.

First and foremost, quantum resilience forces analysis of the system. Engineers and managers must not assume the system is understood. Instead they must focus on the valued functions it delivers and allow those functions to scope and scale the system. Second, focus on known or projected disturbances is ultimately distracting and must be avoided (see the section on robustness). Instead, focus must be on the parts of the system

that contribute to the valued function and their interfaces and dependencies. Third, engineers and managers must avoid creating proxies for resilience. It is frequently observed that poorly scoped systems have led to discussions where resilience is simply another word for what is perceived to be pristine Nature, or measures of something specific in the environment (e.g., Phosphorous). Instead, focus must be placed on quantifying the productivity of the system vis-à-vis the valued functions identified. Identifying valued functions reminds us why we care about the system in the first place. Finally, analogies and oversimplifications of the system must be avoided. If it is a real system, there is no need to analogize. Instead, the system must be faithfully documented through analysis in order to codify its structural and relational complexity. Only then can its resilience be characterized.

The following sections provide some philosophical underpinnings for quantum resilience. Throughout, the rationale for my operationalized definition is revealed and then the characterization formulation is discussed. The practical nature of the approach should be evident throughout (and has already been presaged in the précis).

Valued Function – Systems “For All They’re Worth”

While quantum resilience emphasizes system *function*, it is important that *valued function* has been chosen as a defining phrase (note well, however, it is *not* “value function” or “values function”). *Valued function* is used because it directly calls attention to the observer-based position of all science, and reminds each analyst of their position within the system (Ahl & Allen, 1996, p. 71, 196; Allenby, 2012, p. 361). Varela makes this quite clear:

...the presence of the observer (of the observer-community, the tradition) becomes more and more tangible to the extent that we have to build upon a style of thinking where the description reveals the properties of the observer rather than obscuring them (Varela, 1979, p. xvi).

Though there is likely to be significant overlap, different observers may value different functions and quantum resilience forces these to be specifically identified and included in the analysis. Once identified, even conflicting valued functions can be employed in characterizing the resilience of a system. Focusing on valued functions and appropriate quanta is a vital first step to operationalizing resilience.

Scholars have suggested resilience is a scientifically defined and measurable quantity: “the size of the basin of attraction” (Holling, 1973, p. 20; Carpenter et al., 2001) or the “distance to a threshold” (Walker & Salt, 2006, p. 63) but the terminology remains metaphorical and has yet to be demonstrated with empirical measures. In fact, Carpenter et al. (2001) confess to the difficulty of empirically measuring resilience and remind us of the different approaches and variety of system-specific outcomes. Resilience, then, is a notion we find desirable, but not something we can specifically measure. Interestingly, Rapoport points out that “terms derived from introspection often defy operationalization” (Rapoport, 1986, pp. 31-32) while Gould reminds us that “metaphors can be liberating and enlightening, but new scientific theories must supply new statements about causality” (Gould, 1991, p. 339). Unfortunately, resilience research remains fairly solidly locked in introspection, metaphor and analogy. Resilience has defied operationalization because we have not properly characterized it by first recognizing and quantifying the functions we value. Quantum resilience resolves that by specifically requiring valued functions to be quantified.

There are certainly hints in the literature that suggest “values” should be appreciated. For example, Gunderson speaks of “social benefits” and “desirable stability domain” (Gunderson, 2000, p. 432). This language is clearly value-laden. The Resilience Alliance (2010) speaks of “issues,” and issues only arise when expectations based on values are not being achieved. While their efforts fall short of acknowledging valued *functions*, codification of issues can be a good first step toward identifying the underlying values and outlining what services are important in resolving the issues. Once these services and functions are more clearly defined, they can be quantified.

Observer-dependent valuation can be illustrated by the way Carpenter et al. (2001) speak of a lake’s “turbid” state as having “low utility” (p. 770). Their point is that humans value clear lakes for recreational use and real estate assessment. Turbid lakes are generally thought “ugly” or less appealing. This is a clear superimposition of human values. Later they specifically talk about “*socially preferred* ecosystem state” and suggest “indicators of resilience that are appropriate for the current regime may become useless as ecological structures and *social expectations* shift” (p. 779, emphasis added). For this reason, they define “net social benefit” as a goal that establishes what they value. What is required next is that they quantify this “net social benefit”.

Interestingly, when Carpenter et al. (2001) speak of a lake’s “turbid” state as having “low utility” they continue by suggesting “low utility *but high resilience*” (p. 770, emphasis added). This comment requires some unpacking. How can such language be justified within a quantum resilience framework which suggests resilience starts with definition of the function or service we *value*? That is, when would we ever value the “low utility” of a turbid lake system? To answer this we must first, and perhaps

superficially, ask: “low utility” according to what or whom? Recall, since values are observer dependent, for all we know, there is some person or species which values turbidity and prefers the cloudy water. The overall thrust of their discussion, however, is to specifically identify a turbid lake as qualitatively bad. Referring to a turbid lake system as “resilient” reflects their (problematic) definition and is simply a way these authors acknowledge the difficulty that Nature and/or humans would have in recovering a clear system from it. The real implication of their statement is that apparently the system can become *stuck* in the turbid state. They conflate this stuck-ness with resilience because it seems like a permanent and enduring state. Note that this can only be superficially referred to as resilient—and only if dynamical “stability” is conflated with resilience (see discussion of robustness above). If the “design goal” of that lake system was to be turbid (which they would vehemently deny) then it could be suggested that the lake was *robust* in its ability to maintain that state (for a while), but this is not resilience unless the valued function of the lake was turbidity.

Even if, in a backhanded sort of way, a turbid lake system is valued because of the way it anchors the opposite end of the spectrum from “clear” in someone’s notional conceptualization, it is the “clear” system toward which all management activity is focused. Not only does this clarify the human values, but it clarifies how we stand in proxy for Nature with those values—and this is important to the goal of operationalizing resilience. The functions we (and, presumably, Nature) value from a lake system can be delivered when the lake is clear, not when it is turbid. We would (most likely) never purposefully manage a lake system into turbidity because we simply do not value that condition—but we absolutely could do it (if we valued it). So, while herein I will not

refer to systems in undesirable states as being resilient, it must be understood that such terminology is casual and accommodating and is not in conflict with the principles of quantum resilience. For example, similar comments were made about a “resilient” dictatorship as being “undesirable” (p. 766). But again, this ignores the fact that all such value assignments are observer dependent and that, at the very least, the *dictator* desires and values his dictatorship.

Using the same dictator example which, despite near-universal repudiation, very definitely remains in the realm of observer dependent valuation, Anderies et al. (2013, p. 11) suggest resilience is “not normative” and makes no “specific choices” about what is right. Technically this is correct. That is, if we can truly “measure” resilience, it can become one among many empirical measurements used in decision-making processes. But it must be acknowledged that the intent of most of the resilience literature is focused on enabling management toward what humans have determined is a *better* state. This is *highly* normative. So, while Anderies et al. might never make the choice to manage *toward* a more enduring dictatorship, they certainly understand that others might make that choice, and that dictatorships can therefore be resilient.

Quantum resilience holds that resilience is not about a system until it is first about a valued function. This may seem a clever feint, or it might be suggested that a focus on “function” instead of “system” belies a sort of “wave-particle duality” for resilience, but this is not the case. In fact, to a great extent it is *complex* systems which garner our interest, and these are generally beyond our ken to completely understand and fully characterize (cf. Flood, 1988). Conversely, functions we care about and value are generally well understood. While most talk about “resilient systems,” (see, for example,

Walker & Salt, 2006, p. 11) it is usually a very soft idea of the particular system of interest; unbounded, ill-defined, and amorphous (this, despite pretense that such ambiguities must be resolved). But it is important that while we acknowledge complexity, we do not hide behind it. In the absence of specifically identified valued functions, it is too easy to develop a romanticized idea of the system of interest that results in unhelpful generalizations. In many cases, this leads to a failure to see that the functions we value are delivered by systems other than we are envisioning. Starting with identification of valued functions is not just *another* way to approach resilience analysis; it is the *only* way to appropriately address system scale and identity on the way to characterizing resilience.

It can help to consider “valued function” as a way to elucidate Aristotle’s “final cause” for a system. For example, it is perfectly valid to value “wind resistance” for an office building or “immune response” in a human, but it is not valid to suggest these are the *reason* buildings and humans exist. That is, they are not “final” causes. We do not, for example, build office buildings in order to resist wind and humans do not exist in order to have immune responses. Both of these are examples of protections against the environment in which the systems operate. Instead, for example, buildings exist to provide climate-controlled workplaces for humans. Thus, the idea of *final cause* can help in distinguishing and defining valued function. This can also assist in determining scale. That is, it is valid to “decompose” the human system into a variety of subsystems that include an immune system. So while at the *human scale* it would be incorrect to model immune response as a valued function, in the case of an *immune system*, immune response might indeed qualify as a valued function. Similarly, if there was a “wind resistance system” for a building, its valued functions would necessarily be different than

those of the building itself. Note that while it is possible to model any function at any level, it is vital that consensus be established among analysts to ensure rational results.

Note as well that the “system” of interest must provide the function over a time scale in which it can be usefully valued (whether by humans or by those for whom we stand in proxy). For example, a chair that transforms into a table and back into a chair over a one minute period is an interesting novelty, but it cannot provide a function that could be valued in any sense that relates to the way we think of a chair or a table. In fact, in this case, such system behavior would be considered *only* a novelty (or perhaps “flaky” or “broken”). If for some reason the valued function was specifically “to provide novelty,” then the analysis of resilience can proceed on that basis and quanta can be identified for that valued function. If, however, the valued function was to “reliably provide seating for one” (or some such), then the described chimeric system would be inadequate and its resilience characterized accordingly (e.g., as “zero” or a similarly low number).

With this in mind, Holling’s proposed adaptive cycle (cf. Gunderson & Holling, 2002, p. 25ff.) *might* be viewed as an important analogy and model for some *biotic* systems, but the concepts of collapse and reconfiguration for a system (phases Ω to α to r) must be applied with care to socio-ecological systems. Recall that, according to Holling, resilience waxes and wanes *during* the cycle—it is not that going *through* the cycle demonstrates resilience. Especially for systems on which humans rely, functions provided in the productive and conservative K phase are those which are most likely to be valued. Hence, humans would appear to demand “high resilience” for such systems while they operate in the K phase where Holling suggests resilience is waning. For this reason

(and *contra* Walker & Salt, 2006) assumption of the inevitability of the cycle must be carefully considered, and application of the adaptive cycle as an explanatory tool must be moderated. The section above provides more complete coverage of this topic.

Part of the problem addressed by quantum resilience is that resilience is defined with words like “absorb disturbance and maintain function” (Walker & Salt, 2006, cover) but then immediately the focus is diverted by considerations of adaptive cycles and other metaphors. Specific functions are largely forgotten or given only passing attention. This has resulted in an inability to characterize and operationalize resilience as researchers become distracted with important but ultimately unresolved discussions about complex systems, adaptive behaviors, and basins of attraction. Once resilience has been conflated and confused with ideas like adaptation and transformation, the very reasons for caring about system resilience are forgotten. Adaptation, evolution, and learning, are powerful concepts worthy of exploration, but these concepts must remain distinct from resilience. The section above provides more complete coverage of this topic.

The literature has metaphorically described resilience as a domain or basin surrounding an attractor (Holling, 1973). For this analogy to be successfully employed, the size of the basin must somehow become a measurement of resilience for that system. If this is the case, then it is important that there be a way to characterize the basin. Unfortunately, most literature neglects this requirement. Quantum resilience fills this gap by reminding us that the functions the system provides while it operates “within” that desired basin are what are important in characterizing that system’s resilience. Obviously, any and all “basins of attraction” into which a system can “flip” can be characterized by their respective valued functions. That basins with lower social utility can also be

characterized simply reinforces the point that quantum resilience properly addresses complex systems by first identifying the valued functions they provide. Obviously, the “basin” metaphor must also be reified by the specific system structure out of which the valued functions are provided. This is discussed in the section on system identity.

It might be wondered if parts of a system that actively detract from valued function delivery should also be considered and included in the resilience characterization formulation. In short, the answer is no. Recall first that detractors from valued function are specifically not valued. Second, remember that if the detractors are a bona fide part of the system, they must be considered in the analysis and the valued function output will be attenuated by whatever extent the detractors are successful. Third, again remembering that the detractors must be considered part of the system, it becomes clear that they contribute to the complexity calculation in the denominator. Fourth, if “detractors” are determined to be outside the system, the system probably contains robustness features that protect against such detraction. These robustness features may or may not contribute to valued function, but they will certainly be included in the system complexity calculation (if they are found to be within the system). So, “detractors” are already taken into consideration in a thorough system analysis.

A specific example of “detractors” might be hard to find, but consider the possibility that fiery debate or filibuster might be thought to “detract” from a productive legislative progress. In this instance, valued function might be progress as measured in bills passed per year (or whatever) and disagreement will clearly impact that throughput. Recall, however, that debate is an integral part of the parliamentary procedure and is built into the system to allow for all perspectives to be represented. An accurate model of the

system will include this relational complexity and its (negative) contribution will be expressed in the denominator of the resilience characterization.

Scale – Systems “To a Certain Extent”

Since the quantum of resilience is a unitized output of a valued function and *not* a system, it leads to a very clear definition of how resilience can be operationalized.

Further, it necessarily identifies the actual extent of the system that is required to deliver the valued function. Though a system provides the function, the system cannot be the starting point for the analysis because that would *a priori* bound the system. Instead, quantum resilience recognizes that resilience is not about a system until it is first about a valued function. Discovering the system that provides the valued function is part of the analysis effort.

Many scholars address the need to properly scope a system for analysis. The system boundary selection process is often driven by the particular question you are asking and the assumptions you are making (Allenby, 2012, p. 183). Holling (1995, p. 8) outlines the potential for failure by reminding us that as managers of SESs we sometimes pretend to be outside the system, and hence forget that we are changing along with the system as we manage it. This implies that frequently we must consider the system at a larger scale that includes its managers. Allenby makes this a specific principle of Earth Systems Engineering and Management (ESEM) by insisting that we remember we are “part of the system” (Allenby, 2005, p. 187). Quantum resilience takes this into consideration by specifically acknowledging that delivery of valued functions is what ultimately *defines* boundaries for a particular system. Quantum resilience does not require

(or suggest) that any system be *a priori* bounded or understandable. In fact, it makes a step toward acknowledging complexity by focusing first on the valued function and only then on the system that delivers it. Identifying the valued function frequently leads to a discovery of system scale that is unexpected or even surprising.

In forcing the issue of system scale, quantum resilience analysis does not over-focus on thinking physically. Thinking functionally permits appropriate scoping of system size. If a particular system, subsystem, component, etc. is implicated in delivery of the valued function, it must necessarily be included in system scope—whether or not it initially came to mind. For example, it is tempting to think we want a power generation plant to be resilient when the actual function we care about is provision of electrical power to the grid. This forces specific notice of the boundary at the grid and results in a quantum of resilience that is perhaps some incremental delivery of electricity. If the valued function is properly specified, quantum resilience demonstrates that instead of designing robust and resilient power plants, the incremental provision of electricity comes from a turbine and generator combination and suggests a different place to focus effort at “design for resilience”. Depending on the ultimate scale under consideration, this might imply that there is little point in suggesting a power plant be resilient. Obviously, if the goal is a resilient grid, the valued functions might include “available energy to homes and businesses”. This leads to a different system scale and a different analysis.

Note that the same approach applies even if someone insists that the system really is a *specific* power plant. We still must, however, be honest about the valued function that is important to us. Since no plant exists for its own purposes it would be disingenuous to suggest “electricity output” is the valued function for a standalone, unconnected power

plant. Eventually someone might suggest we value the *employment* offered by a specific plant. Now that a valued function has been honestly specified, we are free to notice that a power plant is just one more redundant means to provide jobs for people. That is, the “employment system” is more extensive than the power plant alone (likewise for the auto industry, or computer manufacturing, etc.). Addressing the valued function is what ensures proper system scale is analyzed. Obviously, we can value *both* electrical output and jobs and quanta can be specified for each of these during a more complete analysis.

Most agree that complex systems emerge from both purposeful and accidental networking of other systems, both complex and simple. If this is the case, then it is a mistake to think we can characterize the resilience of a *whole* complex system. As the discussion of urban resilience revealed, however, nothing is lost because it would not likely be useful in any sense. In fact, completely defining and characterizing the complex system would seem to remove the system’s complex standing. Instead, focus must be placed on the system as it is revealed through the analysis (that is, the scale determined by valued function identification) as opposed to the system as it was envisioned or romanticized. By thinking first of valued function (certainly, *all* of them), the system scale can be formally defined. Once the scale at which important operations occur has been defined, it is less important that the identified system exists within a system of greater complexity. Solid systems engineering practice allows focus to be placed on the parts that impact delivery of valued function. Resilience must be pragmatic enough to not allow us to hide behind complexity.

Identity – Systems “As They Really Are”

Holling’s seminal discussion of resilience suffices to clarify the need to speak of system identity:

Resilience determines the *persistence* of relationships within a system and is a measure of the ability of these systems to absorb changes of state variables, driving variables, and parameters, and still *persist* (Holling, 1973, p. 17, emphasis added).

Though Holling did not take the necessary step and include valued functions in his definition, he makes it abundantly clear that the persistence of a system and its relationships is vital if systems are to be considered resilient. Failure to acknowledge the requirement for continuity in system identity is where much discussion of resilience enters non-productive territory. When a system is allowed to become anything and be termed resilient, it neglects the important requirement of persistence.

Rapoport discusses system identity from two angles: *preservation* and *recognition* (Rapoport, 1986, p. 79). The implication, to which I will return later, is that system-ness is captured in both internal structures and external interfaces. Even when considering systems as black boxes (i.e., where external interfaces alone are important), there must be something inside the box that is creating a recognizable output. Though *preservation* of identity is vital as Rapoport suggests, this need not imply that all systems must be active homeostats or living. In fact, at the timeframes for which the functions of a system are valued, the black box may be considered static or unchanging. Certainly for resilience considerations, the idea of preservation need not be active. For example, bridges and buildings can be resilient without active homeostasis, but their identity is clearly preserved over the periods they deliver their valued function.

This is important in discussions of resilience because (if it is thought of at all) analysts might think preservation refers to how a system reacts to protect itself in the face of perturbation. Recall from above, however, that system functions that permit the system to operate in its intended environment are part of a robustness analysis. Certainly such protectionist features serve to preserve system identity and are also part of its identity since they are part of the system. Further, it could be thought that identity preservation involves how a system adapts over time, but care must be taken in employing the word “adaptation” since that might mean a new system has emerged. This is vital in discussing resilience because there is a cavalier inclusion of the word “identity” in popularized “adaptationist” definitions of resilience where adaptation clearly means different systems (cf. Folke et al., 2010; Walker & Salt, 2006, p. 113). The inclusion of the word identity seems to serve only as a preemptive defense against those who find system identity important. The inclusion is, however, disingenuous given the phrases employed, i.e., “capacity to *change* in order to *maintain* the same identity” (Folke et al., 2010, Table 1). If restraints are placed on what *change* is allowed, this definition *might* work, but it is forced and unnecessary. Inserting a non-specific “capacity to change” in a definition requires further definitional work which is never done. There is, however, no reason to include “capacity to change” in a definition of resilience. As discussed above, significant care must be used in applying the idea of adaptation to resilience.

Cumming et al. (2005) highlight system identity in their definition of resilience. They suggest resilience is “the ability of the system to maintain its identity in the face of internal change and external shocks and disturbances” (p. 976). Despite a squirrely allowance for some undefined degree of “internal change,” this is a step in the right

direction because identity is vital to resilience. Later, they even suggest that if a system has “no scope for the maintenance of system identity, the system clearly lacks resilience” (p. 982). However, the definition falls short because they fail to highlight that a large part of maintaining identity is remaining *recognizable* to other systems. Fortunately, this may simply be an oversight because Cumming et al.’s characterization of identity involves what they refer to as “four *essential* system attributes (structural components, functional relationships, innovation, and continuity)” (p. 980, emphasis added). Again, it is difficult to understand why they list “innovation” among the *essential* system attributes. The requirement that *all* systems be innovative is unrealistic and unnecessarily limiting for a definition of system. Still, acknowledging “functional relationships” as a part of identity is a vital step toward understanding that resilient systems must continue to provide their valued functions.

How then is identity best defined? Varela (1979) suggests identity consists of *structure* and *organization*. To him, system *structure* is a current configuration of system parts and the connections that represent relationships between those parts. *Organization* consists of the rules that define and maintain the parts and their relationships. This is roughly equivalent to Rapoport’s ideas of recognition (structure) and preservation (organization), though there is no specific need to press that equivalence. Varela’s conceptualization of identity clearly maps to *phenotype* and *genotype* as might be expected of a biologist. In a nearly identical approach, Simon (1962) suggests that “state descriptions” and “process descriptions” constitute system identity. State descriptions describe what is, and process descriptions describe how to create and maintain it. Ahl and Allen (1996) poetically refer to the identity of the system as “context” (i.e., configuration,

p. 101ff) and stress the need to identify system “surfaces” (i.e., interfaces, p. 139ff).

While Ahl and Allen may deemphasize how a system arrives at its configuration, importantly they highlight the external “shape” of a system as critical to its identity.

Quantum resilience adopts Varela’s notion of identity (though emphasizing “structure” over “organization”) but clarifies and extends it by including the idea of function as a way to acknowledge Rapoport’s *recognition* idea. That is, identity is not simply what the system *is*, but includes what it *does*. Importantly, quantum resilience always stresses that identity is as much about structure as it is about function, since it is in the delivery of functions and services that a system remains recognizable to other systems. Vitality, this makes resilience *phenotypic* and not *genotypic*. This important contrast will be elucidated later because it serves to further distinguish resilience from adaptation and fitness which clearly involve ideas related to system genotype (see also discussion above).

Importantly, the identity principle also reinforces *redundancy* concepts (see next section). In a system of many parts, there may indeed be different parts that provide the same valued function. These parts need not be structured similarly or even operate the same way (see discussion of redundancy below). For example, in mammalian cells, adenosine triphosphate (ATP) is the primary energy currency used as fuel for biotic processes (and, incidentally, this molecule is an exemplary “quantum of resilience”). Oxidative phosphorylation, glycolysis, and the Krebs (or, citric acid) cycle provide most of the ATP production in the cell (Champe & Harvey, 1994; Campbell & Farrell, 2003). These structures and processes are wildly different, interdependent, and show tremendous diversity, but for the valued function of extracting energy from sugars and turning it into

usable fuel, they are ultimately redundant. Further, in delivering the function, they provide a unified interface in the form of ATP molecules.

Quantum resilience supports the idea that system structure and function can be changed, but stresses the requirement that the new system have its resilience re-characterized. This is because the altered system is a new phenotype. This allows a manager or engineer to identify targets of opportunity for system maintenance, augmentation, or abandonment and experiment with alternative deployments. Note that a proposed alternative deployment might be configured to provide incrementally more delivery of the valued *function*. Superficially it is tempting to conclude that such a change confers higher resilience, but this conclusion would be premature. Since identity also involves *structure*, the altered system's resilience must be re-characterized, because the change may have introduced significantly different structure (higher complexity) that results in *decreased* resilience. In either case, however, system identity must be maintained. Bridges that morph into coffee makers and salt marshes that morph into coral reefs cannot be termed resilient, notwithstanding the importance and utility of the outcomes. If a system's identity has changed dramatically, the resilience of the "before configuration" and the "after configuration" may not even be comparable.

Carpenter et al. (2001) suggest "resistance" is important in assessing long-term system persistence. Resistance, they suggest, is both a "complementary aspect of persistence" and "the complementary attribute of resilience" (p. 766). Though the conflation of terms makes it difficult to understand their goal, it is certainly reasonable to assume these could be viewed as contributing to identity maintenance. As described above, this is just another term for robustness. Learning *how* a system defends itself

against its environment is an important pursuit, but it should be remembered that such defenses contribute to the system structure and only in that way do they impact resilience characterization (in the denominator). If there are functions performed by the system that render it “resistant” to certain perturbations (like an immune system), then it might be thought that these functions can be valued and serve to augment the resilience of the system. This is certainly possible, but in most cases defensive features should be acknowledged as elements of good design that allow the system to do what it was designed to do in its intended environment. For the most part they should not be considered as separately valued functions.

The identity of the system, its organization and structure is vital. This is the key contributor to the complexity of the system. *Demonstrating* or *witnessing* resilience is a simple matter of observing function delivery, but *characterizing* resilience (so systems can be compared) requires some insight into the structure. This will be made clear in the characterization formulation discussed below.

Redundancy – Systems “Pure and Degenerate”

Quantum resilience specifically employs the term “redundancy” instead of “diversity” since diversity does not necessarily always imply duplicated function. Because function is paramount, this selection of terminology is vital. Redundancy is sometimes disparaged as simplistic, wasteful, or reductionist, but it is actually a rich concept. Quantum resilience recognizes several deployment options for redundancy that are sometimes confused and misunderstood. The first is *pure redundancy* where both structure and function are replicated. This kind of redundancy—sometimes mislabeled

“engineering redundancy”—is the easiest to track when characterizing resilience. It is certainly an approach employed by engineers, but is actually something learned from nature (e.g., humans have two carotid arteries). The second is *degeneracy* where structure may differ but function is duplicated (Tononi, Sporns & Edelman, 1999; Edelman & Gally, 2001). This kind of redundancy is ubiquitous in nature and is also frequently seen in human-engineered systems from surveillance to telecommunications to healthcare. A third kind of redundancy has been termed *imbricated redundancy* but is simply an umbrella term for the preceding two options. It has been suggested that imbricated redundancy of function is provided at different system “scales” (Holling, Gunderson & Peterson, 2002, p. 84ff). This is certainly possible, but is not specifically required.

Since quantum resilience supports an array of redundancy options, diversity can be viewed as just another way of saying redundancy when the focus is on the outcome: the delivered function or service. If diversity is doing something a different way, then redundancy simply reminds us that it is not the “different way” that is really important. It is the “something” being done that is important. The vital requirement is that the valued function be delivered. It can certainly be provided in a diverse and variety-filled manner. For example, if the valued function is electricity to a residential outlet, then assuming all the proper mechanisms are in place, it can be provided from the grid, directly from solar panels on the roof, from a gasoline-powered generator, or even by someone pumping pedals on a bicycle. This is marvelous diversity, but the output is identical.

Peterson, Allen and Holling (2010[1998]) suggest that ecological resilience “derives from overlapping function within scales and reinforcement of function across scales” (p. 182). They further clarify that “when a functional group consists of species

that operate at different scales, that group provides cross-scale functional reinforcement that greatly increases the resilience of its function” (p. 185). And finally they summarize with “ecological resilience is generated by diverse, but overlapping function within a scale and by apparently redundant species that operate at different scales” (p. 189). Later, Gunderson and Holling (2002) refer to this concept as “imbricated redundancy” (Holling, Gunderson & Peterson, 2002, p. 84ff). They discuss the idea of imbricated redundancy with the analogy of a theater where actors are “waiting in the wings” ready to go on stage when “change is required.” Their theory is that extra actors (spares) can be used to “change the pace” or add variety to the plot. Though obviously Nature would not “orchestrate” such a change (it cannot detect that “change is required” and insert new actors to “change the pace”), it is clear that—if there *was* a plot—such extras might contribute unexpected twists. Still, their example is insufficient since the proposed spare parts “waiting in the wings” are *actors*, which means all redundancy is *pure* and occurs at the *same* system scale (e.g., the play or production). This illustration can be improved.

Consider instead a theater where an actor falls ill and is replaced by an understudy. This demonstrates how pure redundancy perfectly (and expensively) allows the show to go on (some may refer to this as “engineering resilience”). Now consider, given the diversity of jobs at a theater, that instead of the actor, a ticket salesperson, usher, or concessions attendant gets sick and fails to report for duty. These roles could *also* be filled by the understudy, though less perfectly (and likely with a great deal of whining). This is *degeneracy* since it emphasizes that while identical *function* is ultimately delivered (ushering or selling) the *structure* is different (i.e., it is an *actor* not an *usher*). Note as well that this is closer to *imbricated redundancy* since the function

replacement is occurring at a *theater* scale instead of the *production* scale. To push the example a bit farther, assume there is no understudy. Now, if an actor falls ill, the actor could be replaced by a ticket salesperson that has seen the play a few times. This is an interesting example of degeneracy *and* imbricated redundancy because it illustrates high variation role-filling given that the ticket salesperson has not rehearsed the role. Since this is not complete functional redundancy, this vignette could result in a tragedy evolving into a comedy (or vice-versa) or larger system impacts like production failure leading to bad reviews, refunds, unemployment, and other social implications. Obviously a ticket salesperson confers little resilience to the *production*, while the redundant actor leads to increased *theater* resilience at some cost.

Figure 2 illustrates this idea but with a very important addition: functions have been identified to demonstrate how this can be operationalized by quantum resilience. Following Holling, multiple scales are shown (S_1 - S_4), though the number of scales in the figure is only illustrative and will be system-specific. Whether or not the scales shown become larger or smaller as the identifying numbers increase is also merely illustrative and (eco)system specific (that is, even “scales” can overlap and need not be construed as implying strict containment). At each scale, a certain number of species (subsystems) are found. Each performs certain functions and it can be seen that there is overlap in the functions provided by species (subsystems) within and across scales. As depicted here, each species (subsystem) at scale S_1 performs fewer (eco)system functions than those at S_2 and this trend is consistent in this figure, though it need not be in a real (eco)system where distribution of function and degree of overlap are unplanned (and, in fact, may be disputed by different experts). At the right, it is notionally shown that overall (eco)system

function is the sum of all the functions performed within the (eco)system. The choice of “sum,” though logical, is also a simplifying feature. In reality, quantifying overall function is likely more complicated, but importantly may not be necessary.

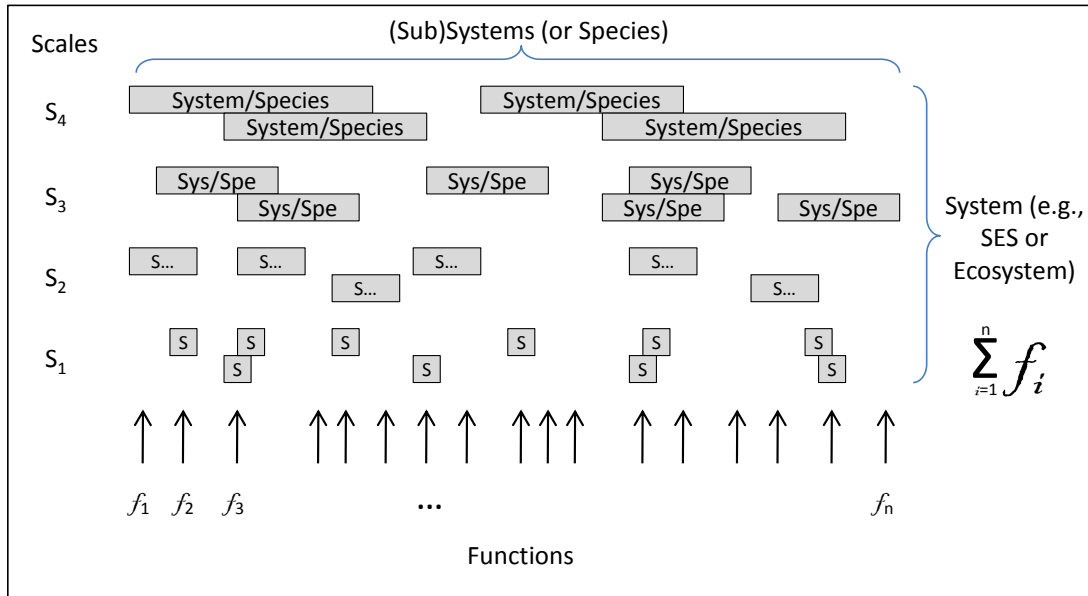


Figure 2. System-of-Systems Providing Functional Redundancy Across Scales

While Holling et al. (2002) allow for “imbricated redundancy” in natural ecosystems, it must be recognized that such a concept applies to—and is evidenced in—most socio-ecological systems and large-scale systems of human design. Of course, superimposing such functions on nature is a completely human activity which reflects human values and is usually derived *ex post facto* simply because “it is there” (as opposed to because “it is required” for achieving some “better” state). Humans will tend to preserve the functions they “observe” (or superimpose) and “value” these because “nature made it that way.” This is not necessarily wrong, but it must be admitted to be a fabrication. We know for example that nature does not care if a species providing some function f_i disappears. For all we know, by eliminating that species or function, nature is

progressing toward an even “better” future state and that, by preserving that function and its attendant species, humankind is risking a qualitatively diminished end game.

Despite the fact that human values determine the functions, it is clear from Figure 2 that resilience is simply a matter of redundancy (pure and degenerate). Resilience is redundancy writ large—reflecting the portfolio approach adopted by nature as its mechanism of endurance. In fact, pure redundancy and degeneracy are exactly the kinds of redundancy that system engineers use to ensure robust systems in many domains from law enforcement to surveillance to telecommunications to health claims processing. If cost and schedule were not an issue, such resilience could easily be exhibited in many more human-engineered systems. Engineers, however, are taught to deliver verifiable function, reduce costs, and balance the benefit of redundancy with careful scoping of system requirements as they are held in delicate tension by cost/benefit trades.

Importantly, among the reasons natural ecosystems have settled on redundancy as an approach to resilience is because it was accidental and Nature had no intention (or driving reason) to be frugal. Among the reasons researchers have (until recently) denigrated redundancy is because humans are interested in being wise with resource use (usually driven by economic considerations) and redundancy just seems extravagant and wasteful. Even the concept of “diversified financial portfolio” demonstrates and reinforces that redundancy (in this case, through degeneracy) leads to resilience. If the valued function we seek from an investment strategy is income, growth, or preservation of capital in a capricious market, then a diversified portfolio is what provides this (it is, after all, a probability game). However, if our goal is growth and we have insider trading information that guarantees a significant return on investment, it is counter-productive to

diversify. In this case, it makes sense to focus our strategy on what we know will be a high payoff investment. In general, since neither we nor Nature has such insider trading information, we find degeneracy to be the winning ticket.

Gunderson (2000) is one of only a few scholars who will go on record with specific recommendations on how to improve resilience. Despite later offering some highly conflicting conclusions (cf. p. 435), he suggests “In order to add resilience to managed systems, at least three strategies are employed: increasing the buffering capacity of the system, managing for processes at multiple scales, and nurturing sources of renewal” (p.434). Each of these emphasizes redundancy. Buffering capacity is increased through redundant structures (p. 434), processes are duplicated at a variety of scales, and multiple sources of capital and skills (p. 436) are employed in nurturing renewal.²

As shown, the idea that redundancy is a primary mechanism for operationalizing resilience has not been missed in the literature, but to date, it has not generally been exploited by practitioners. Lately, there is a cautious change in the scholarship though it has still not reached a level that allows resilience to be operationalized. With a goal of elucidating specific principles for enhancing the resilience of ecosystem services, Biggs et al. (2012) have unsurprisingly suggested that redundancy must head the list. Setting redundancy as principle #1 is a nod toward finally understanding the power of redundancy.

² Gunderson’s conclusions are repeated in Gunderson, Pritchard, Holling, et al. (2002) with equally problematic surrounding discussion in which resilience is said to be “generated by destroying and renewing systems at smaller, faster scales” (p. 264). In this conclusion alone, language like “resilience is *defined* as...”, “resilience is *reestablished* by...”, “resilience is *maintained* by...”, “resilience is *generated* by...”, and “strategies... *contribute to* resilience” is employed with a significant amount of euphemism. While it is beyond the scope of this work to address this in full, it must be pointed out that such casual language—even when presented in summary conclusions—is not helpful to the discourse. Much more precision is required if resilience is to be operationalized.

Low et al. (2002) call for the development of a “grounded theoretical approach to the study of redundancy, for efficient and responsive management depends on matching optimal levels of redundancy to the appropriate conditions” (p. 108). Quantum resilience provides this grounded theoretical approach and through its focus on valued function and incremental delivery of quanta of resilience demonstrates that value of redundancy in designing for resilience.

Defining Resilience

The quantum resilience approach stems from pragmatism. Resilience is only a useful idea if it can be characterized and operationalized. My focus is pragmatic because I am endeavoring to operationalize a concept so that management can occur. Ultimately, we care about resilience because we want systems that are useful and enduring. Hence, resilience is *not an end in itself*. Unfortunately it is treated as such in much of the literature. Resilience is frequently the goal. Resilience has become the Holy Grail *du jour*; elevated and idolized as if we forget there is a reason we pursue it. Instead, resilience must be viewed as an idea which, when properly operationalized, defines the character of systems that can deliver what we *really* want: functions we value. After reviewing some alternative definitions, this section defines what resilience is and operationalizes it for effective use across all disciplines.

Scholars argue resilience is a well-defined mathematical concept (Holling, 1973; Fiering & Holling, 1974), but there has been little progress in calculating resilience. Using the language of dynamical systems theory, Holling (1973) suggests it is like a domain or basin of attraction around an attractor. Carpenter et al. (2001) suggest

resilience is the size of that domain or basin of attraction. To date, however, little has been done to make inroads into the difficult effort of characterizing or quantifying the idea for specific systems. Hence, while resilience is acknowledged as a desirable notion, it has little operational value and the definitions remain broad and soft—perpetuating this dilemma.

Being able to compare such deployment options is vital since for decades Holling (1986, 1995, 2002, 2010) has been reminding us that when we try to “manage” ecosystems we tend to make them more fragile and less resilient:

Because of the initial success in reducing the variability of the target variable, features of the biophysical environment which were implicitly viewed as constants began to change to produce a system that was structurally different and more fragile.... The biophysical environment became more fragile and more dependent on vigilance and error-free management.... the ecosystems simplified into less resilient ones as a consequent of man’s success in reducing variability (Holling, 2010[1986], p. 102).

More recently Holling refers to it as a “puzzle” of success and reminds us:

Any attempt to manage ecological variables (e.g. fish, trees, water, cattle) inexorably led to less resilient ecosystems, more rigid management institutions, and more dependent societies (Holling, 1995, p.6).

Even if such assessments are notional and supported by scant empirical data, it is not my intent to argue the conclusions. Instead, I want to suggest that these failures might have led us to more fruitful ways of looking at resilience. Unfortunately, a huge literature has evolved that defines resilience with analogies and examples that have led to its dilution and conflation with adaptation, transformation, reliability, robustness, etc. (see extended discussion above). At best, we know we want resilience, but seem only barely to “know it when we see it.” At worst, operationalizing resilience leads to

recommendations to decrease human involvement because by reducing human impacts we can “increase” or “restore” resilience. “All natural” seems to be the recipe for “highly resilient” systems. In many cases “resilience thinking” has become the face of a new conservationism. Despite many good intentions, there are no definitive marching orders on operationalizing resilience. Quantum resilience is intended to provide some clarity in this expansive discourse and to operationalize resilience for use in system study.

Sometimes resilience is simply used as a euphemism for “environmental quality” where such quality was assumed to be maximized in pristine Nature that is untouched by humans. As such it stands as a new science behind conservationism and the normative values espoused therein. For example, by suggesting that “resilience has declined since colonization,” Walker et al. (2009, p. 17) assume that prior to colonization, the Goulburn-Broken catchment in Victoria Australia had “higher” resilience. This, however, is impossible to assert since the catchment today is an entirely different system than it was in pre-colonial times. This tendency to assume untouched nature is the exemplar of high resilience is evident throughout the literature (see also, Walker & Salt, 2006) and it is usually stated as a matter of acknowledged fact with no supporting quantification. In light of the fact that scholars have not produced a characterization of resilience that can be used to compare across (or even within) systems, this can only be recognized as an assumption that what Nature had done was somehow “best” and that if we want high resilience we must restore it. Quantum resilience provides a mechanism whereby such assumptions can be avoided and real comparisons can be made.

Especially in the ecological literature, resilience sometimes becomes a proxy for “fitness” of a system. So, a push for resilience in socio-ecological systems becomes a

push for optimum fitness of these systems. We want certain systems in perpetuity, so they must be resilient in order to survive through “thick and thin” (or however it is euphemized). This seems like a simple observation but usually the ramifications are not considered. Resilience *may* contribute to fitness, but fitness can only be observed after the fact, for a given system (phenotype), on a given landscape. Fitness is an after the fact measure of adaptation (or, adapted-ness). The trouble with this casual equation of fitness and resilience is that we have limited knowledge of how the landscape might change in the future and cannot wait that long to determine if our important systems are actually resilient. Resilience, however, must actually mean something here and now. By forcing focus on delivery of valued function, quantum resilience ensures this. Perhaps the best way to navigate the issue of resilience conflation with the idea of fitness is to think of resilience as phenotypic “fitness” while “real fitness” is genotypic. When it comes to valued functions and resilience, it is a phenotype (a particular system) that is implicated. As described at length above, adaptation and fitness are best left as research areas distinct from resilience.

Recall as well that a system’s resilience must be able to be characterized in the absence of perturbations and disturbances. Most authors do this tacitly by suggesting that resilience is *degraded* after a disturbance of some sort (e.g., a reef’s resilience is degraded by years of pollution, or a catchment’s resilience is degraded by years of industrial agriculture, etc.). This casual language implies that resilience stood at some *higher* value *prior* to the disturbance. Obviously this implies that resilience can be characterized without reference to specific perturbations so it is very odd that most definitions of resilience have perturbations or disturbances in them. Thinking specifically

of resilience “to what” as suggested by Carpenter et al. (2001) is *not* a mechanism by which we can extend resilience theory—it is not a way forward.

Fiksel (2003) defines resilience as the capacity of a system to tolerate disturbances while retaining its structure and function. This is very close to being a useful definition, as long as “tolerating” and “retaining” are given operational value.

Unfortunately, Fiksel does not provide an operational way forward. Worse, he teases:

Traditional systems-engineering practices have tried to anticipate and resist disruptions, but may be vulnerable to unforeseen factors. An alternative is to design systems with inherent resilience by taking advantage of fundamental properties such as diversity, efficiency, adaptability, and cohesion (Fiksel, 2003, p. 5330; Fiksel, 2006, p. 16).

While it is easy to casually state the need to design “inherent resilience” it is unhelpful to actual working engineers and SES managers. Comments like these clearly exemplify the manner in which resilience has become diluted and ill-defined. The trouble with his comment is that it is impossible to operationalize. For example, both diversity (if he means functional redundancy) and efficiency are what engineers already do and are frequently reminded how it will not work for the more complex systems of today (cf. Walker & Salt, 2006, pp. 7-8). To recommend engineers take advantage of “adaptability” after accosting them for trying to anticipate disruptions and leaving their systems vulnerable is contradictory. Further, since adaptability can only be observed after the fact, Fiksel is likely talking about designed-in system features that allow responses to known or projected environment changes, which are again, what engineers already do. Finally, proposing that “cohesion” (however defined) be employed just seems uninformed since (1) no system engineer wants his system to fall apart, (2) too much cohesion increases vulnerability, and (3) it flies in the face of the requested efficiency. Instead, engineers

must be taught to target something else: a quantum of resilience that *is* meaningful, measurable, and productive.

Zolli (2012) aggressively proposes that resilience requires a system's "core purpose and integrity" to be maintained "in the face of dramatically changed circumstances" (p. 7). He stresses two essential aspects: continuity and recovery in the face of *dramatic* change. Unfortunately, great concepts like "purpose" and "integrity" are left for the reader to define and operationalize. Importantly, quantum resilience provides the solution in both cases. When identity is specifically codified, and purpose is defined according to valued function, operationalizing resilience becomes an accomplishable task.

Walker et al. (2004, p. 4) define resilience as "the capacity of a system to absorb disturbance and reorganize while undergoing change so as to still retain essentially the same function, structure, identity, and feedbacks." But if a system is to *retain* its "function, structure, identity and feedbacks" it is very difficult to understand what it means that it can also "reorganize while undergoing change." This definition contains an assortment of interesting sounding words with enough equivocation to allow them to mean just about anything. This is problematic and indicative of the expanding softness of the concept in the literature. Adding a word like "essentially" does not eliminate the fact that words have meanings and that meanings are important. Obviously more must be said for this kind of definition to be operationalized.

As if recognizing this, Walker tightens the previous definition and suggests resilience is "the ability of a system to absorb disturbance and still retain its basic function and structure" (Walker & Salt, 2006, p. 1). With the exception of the part about

absorbing disturbance, this is *much* better. Unfortunately it is difficult to know which definition is real since Walker and Salt define resilience in at least *four* different ways in their little book (cf. pp. xiii, 1, 12, 32, 37, 62, 63, etc.). One of the definitions (p. 32) very nearly reverts to the originally problematic Walker et al. (2004) definition. Additionally, resilience is still defined in terms of an “ability” which is never characterized.

Surprisingly, toward the end of the book (p. 113) they employ the word “identity” and offer their best definition. Unfortunately, this solid definition is immediately diluted with the requirement that many largely metaphorical concepts like thresholds, forces, regimes, basins, and adaptive cycles, be used to analyze resilience. The valued functions a system may provide are ignored. This is not a step toward operationalizing resilience.

Carpenter et al. (2001) propose that

In any study of resilience, we are concerned with the magnitude of disturbance that can be tolerated before a system moves into a different region of state space and a different set of controls (p. 766).

But they neither offer a way to characterize the magnitude of disturbance nor provide insight into the list of disturbances that might be implicated. Quantum resilience suggests that the *magnitude of a disturbance can only be characterized by observing the decrease in delivery of the valued function as quantified by the established quantum of resilience*. Note how little it has to do with the disturbance: no matter what strikes, the question is “how much less function is being delivered?” In fact, without some notion of the valued function, there is no way to define what it means to “tolerate” or move “into a different region of state space.”

The resilience literature is failing to converge because nobody has identified mechanisms to operationalize the concept. To be a useful system concept, resilience must

be *characterized* for each system. This is a specific word choice that defines resilience as a system “character”. That is, resilience is *not* an ability, capacity, function, or feature (that is, it is *not* something a system *does*). Further, resilience is *not* a property or attribute (that is, it is *not* something a system *has*). Instead, “resilient” is something a system *is*. Quantum resilience provides the mechanism for characterizing the extent to which a system is resilient.

With this in mind, I suggest there are two angles from which the idea of resilience must be defined. The first is oriented toward the fact that resilience is a “character” of systems: Systems that persist in their identifiable form (structure and function) are generally termed resilient, so **resilience is the idea that a system’s identity persists**. The second definition is operational and measurement-oriented: **Resilience is the extent to which a system delivers its valued function**. A longer form of the operationalized definition is: When properly characterized in accordance with its identity, resilience quantifies (or, “is a measure of”) the extent to which a system delivers its valued function.

As discussed at length above, there is a lot implied in the word “system” and those hidden parts of the definition are important. As a quick review, recall:

- *System identity* is defined in terms of structure and function. For a system to be resilient it must persist and be recognizable in some meaningful way—usually by continuing to provide its function as facilitated by its structure.
- In order to properly establish appropriate scale and to characterize the resilience of the system, *structure* is defined in terms of system-of-systems nearly

decomposable hierarchies (Simon, 1962) and their relationships (interfaces and dependencies between systems). This is discussed at length below.

- *Function* is defined as that valued behavior that a system performs or output that it generates (usually within some acceptable tolerances according to a service level agreement) and is generally observed at a system interface (e.g., an output or influence, etc.).
- *Perturbation* is specifically not involved in either the definition or characterization of resilience.

With the foregoing in mind, resilience can be characterized on a per-system basis. Such characterization should permit comparison of similar (homologous) systems or alternative system configurations, but confirms that resilience cannot be reduced to a universally absolute value. Neither would it make sense to do this. Consider, for example, what it would mean to say that a lake-agriculture socio-ecological system is more (or less) resilient than the Internet. Such comparisons are nonsensical. What this characterization does illustrate is that resilience can be bolstered in systems only in increments. That is, resilience grows as each quantum of a valued function is incrementally delivered by redundancies in the system. In this way, resilience is a quantum concept.

Characterizing Resilience

Characterizing the resilience of a system is done following a thorough system analysis. System analysis is an iterative process that coalesces as knowledge about the

system grows. The resilience characterization process has the following major components:

1. Enumeration of valued functions/services and identification of the quanta of resilience. As indicated, valued functions drive the analysis and must be elucidated first.
2. Enumeration of systems implicated in the function delivery. Here, the functions determine the scale of the system as the real system emerges.
3. Explication of the structure, the connectivity, and ultimately the identity of the system. This analysis phase completely characterizes the system to a level acceptable to participating domain experts.
4. Assignment of valued functions and quanta of resilience to implicated systems.
5. Calculation of apparent complexity and connectedness of the value delivery system. With the previously specified function quanta these values can be used in characterization of resilience.
6. Once characterized, resilience of alternative system configurations can be compared and contrasted in decision frameworks.

Systems engineers have long employed Simon's idea of nearly decomposable systems in their analysis work. In general, it is suggested that analysis *start* with that idea:

The fact, then, that many complex systems have a nearly decomposable, hierarchic structure is a major facilitating factor enabling us to understand, to describe, and even to "see" such systems and their parts. Or perhaps the proposition should be put the other way round. If there are important systems in the world that are complex without being hierarchic, they may to a considerable extent escape our observation and our understanding (Simon, 1962, p. 477).

Simon allows for complex systems to be incomprehensible, but he stands fairly firmly in the hierarchy camp:

I have already given some reasons for supposing that the former is at least half the truth—that evolving complexity would tend to be hierarchic—but it may not be the whole truth (Simon, 1962, p. 478).

Systems should be decomposed as necessary for the given context/domain of discourse. System decomposition can proceed along lines of system engineering best practices with consensus-building among domain-cognizant engineers. While consensus can be construed as a low bar for success, it does not free engineers to be less rigorous in the approach taken. Since definitional rigor is missing from the literature, quantum resilience demands “accurate” models, but this does not mean that a given complex system does not admit to multiple unique “correct” models. As with any boundary decision or modeling choice, consistency is required, and if experts disagree, transparency is enforced by tools so there are no murky system definitions. For example, when you step on the scale to check your weight, few will care what you are wearing or carrying since that is a matter of personal taste or convenience. No matter what your choice is, the scale will provide the actual weight of whatever is on it. If, however, you want to compare your body weight from day to day, the experts will demand consistency. They will come to some consensus that if you choose to carry your cat and computer onto the scale (as a matter of taste), that you must do so every day. This illustrates the need to understand that, though system definition is a matter of choice and consensus, the system being characterized is very real and must be treated with consistency.

To ensure proper identification of relationships, it is important to be fair and realistic in representation of where variables are available (i.e., for measurement) and where they are controlled. The resulting decomposition should garner general agreement among experts. If disagreement on system decomposition occurs, systems can be represented as black boxes as long as consensus can be gained on the decomposition depth at which the analysis is to proceed. Such system summarization will likely add a measure of uncertainty to the analysis (discussed below), but it will be completely transparent since it is captured in the tools. The degree of connectedness of the systems involved in valued function delivery will lead to some idea of the level of independence of the implicated systems.

There is much that goes into any measure of complexity (cf. Gell-Mann & Lloyd, 1996; McShea, 2001; Erdi, 2008, pp. 201ff; Farnsworth, Lyashevskaya & Fung, 2012; Tamaskar, Neema & DeLaurentis, 2014; see also Parrott, 2010 for a serviceable overview of measuring *ecological* complexity). Such measures generally include the number of system parts and how they are connected. Frequently, important system traits are ignored like the number *moving* parts, their cost, their provenance and pedigree (i.e., lineage or depth of supply chain), the *nature* of the connections (information, matter), etc. Generally such omissions can be shown to not greatly impact the value of the resulting metric, but each can be important in specific applications. The complexity measures used in characterization of resilience can be adjusted to garner approval by collaborators in the analysis, but it must be consistently applied. Herein, I recommend an approach that is both tractable and generalized. It not only recognizes hierarchical and relational complexity, but it enforces its use with transparency and consistency.

Kolmogorov (1963) complexity is generally defined as the minimum length of a program written in a description language required to produce a desired system (cf. Page, 2011, p. 29). For system analysis of the kind required for characterizing quantum resilience, any *formal* specification language would suffice. For example, this can be demonstrated to be a computable outcome of systems described according to the ITU-T SDL Z.100 specification (itu.int/en/publications), IDEF (NIST, 1993), or UML/SysML (uml.org), and many complexity metrics have been proposed and demonstrated. Gell-Mann and Lloyd (1996) coin the term “effective complexity” because

The meaning of the term ‘complexity’ that corresponds most closely to its use in ordinary conversation and in scientific discourse corresponds to *effective complexity*, the length of a concise description not of the entity but of its *identified regularities*.... Effective complexity measures knowledge, in the sense that it quantifies the extent to which an entity is taken to be regular, nonrandom, and hence predictable (p. 49, emphasis added).

Extending Simon, and following Gell-Mann and Lloyd in focusing on regularities, I define a system’s *apparent complexity* (C_a) as that complexity that we can actually manage and that impacts the function of our system. It is a function of system hierarchical structure, relational inter-connectedness, and dependence on other systems:

$$C_a = f(s, c, d)$$

For use in resilience characterization, apparent complexity consists of the total number of decomposed systems (s) implicated in delivery of valued functions, the total number of connections (internal and external) among those systems (c), and the in-degree of those systems. In-degree of a system (a.k.a. network node “fan-in”) is a well-known network science concept (cf. Borner et al., 2007) that is the total number of inputs to a particular system (“node”). It is used herein to represent the level of *dependence* systems

have on other systems. Note that the *size* of the state/phase space, i.e., the number of system states and supported transitions (which are obviously less *apparent*), may be added to this metric if it is determined to be an important driver of complexity. Generally, and as Simon would likely argue, deeper hierarchical complexity will subsume phase space complexity and permits black box treatment of systems. The formulation for apparent complexity is recursively calculated through the system hierarchy as follows:

$$C_a = s + \sum_{i=1}^s (c_i + d_i)$$

Note that the number of decomposed systems will be driven by the specific requirements of the analysis and may differ between experts (cf. Simon, 1962). What is required is that the decomposition is sufficient to adequately model the function and this will require consensus (thus enters the balance of *structure* and *function*). Note that connections are formed from known flows (both material and immaterial) between systems as well as shared measurable variables. For example, a local environment temperature that is employed in state definitions by two otherwise isolated systems contained within that environment constitutes a connection.

System *productivity* (P_s) vis-à-vis the valued functions is a function of the number of redundant systems (n) providing each unique valued function and the overall number of quanta of each valued function provided (q):

$$P_s = f(n, q)$$

Each system's contribution to delivery of the valued function is measured in the specified quanta and contributes to the quanta delivered by the overall system. For quantum resilience, overall system productivity is given by summing the product of the

total number of systems that provide each unique function and the total number of quanta of each function provided:

$$P = \sum_{i=1}^f \left(n_i \times \sum_{j=1}^n q_j \right)$$

Resilience (R) is characterized as a function of overall system productivity and apparent complexity: $R = f(P_s, C_a)$ as shown in the following equation:

$$R = \frac{P_s}{C_a} = \frac{\sum(n_i \times \sum q_j)}{\sum C_{a_i}}$$

Note that resulting resilience characterizations are not comparable across dissimilar systems—as described above, such comparisons are not meaningful—but can be used to compare similar systems or multiple deployment options for the same kind of system (one that delivers similar valued functions).

This characterization suggests that resilience is improved by increasing the redundancy of systems (n) that deliver the valued function (q), but it also supports the intuition that the complexity of the systems and their interdependence might diminish the resilience of the overall system.

As an example, Figure 3 depicts a hypothetical system of systems (S_0) that has been deduced from a hypothetical analysis process. This is known as the canonical example and can be used as a test case for future implementations in other tools. Note that S_0 is neither required to be a “real” system nor must it be a “complete” system. Instead, it is the boundary chosen by analysts to best assess the list of agreed-to valued functions. As a simplifying assumption for the example, one valued function (f) has been identified (P_s is simplified to $n \times q$). As shown, the valued function is redundantly

delivered by three systems (S_A , S_B , and S_C). For purposes of this example assume each provides a different quanta of that function: S_A provides 2 quanta of f , while S_B delivers 3 quanta and S_C offers 1. Note that while S_D does not directly provide the valued function, it is indirectly implicated in the delivery of the valued function because it provides inputs to S_A and S_B . Note as well that S_C relies on inputs from S_A and S_B .

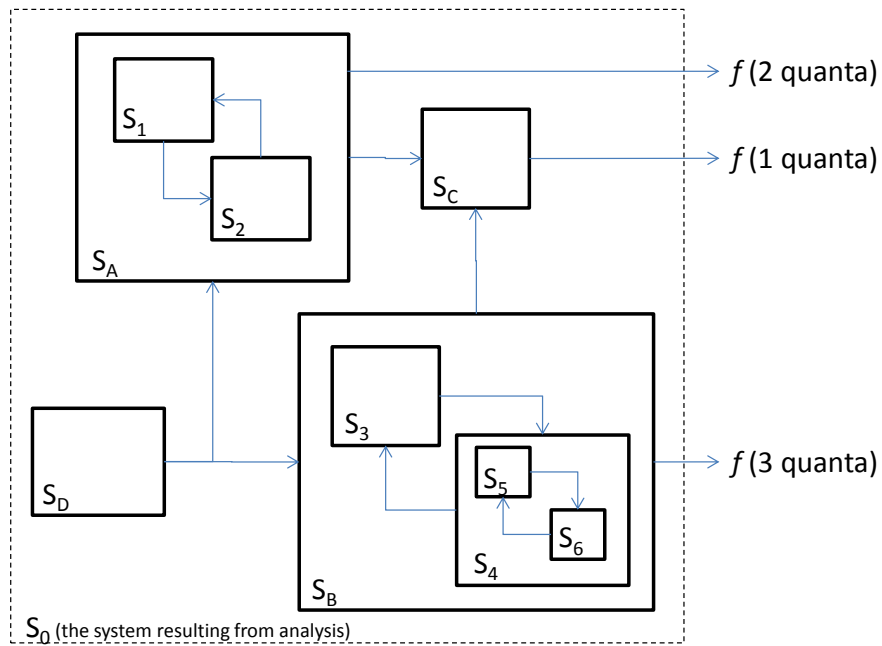


Figure 3. Canonical Example: System-of-Systems with One Valued Function

By hypothetical analysis, S_A was determined to be decomposed into systems S_1 and S_2 which work together to deliver 2 quanta of the valued function f . Hence (since the “parent” system is always counted in the decomposition), the number of decomposed systems s for S_A is 3. As shown, S_1 and S_2 are similar, each having a decomposition of one with one input connection and one output connection. In addition to its internal structure, S_A also has two output connections and one input connection for a total c of 13.

Hence, apparent complexity ($C_a = s + c + d$) for S_A is 13. The remaining systems are similarly characterized.

Table 7 shows pertinent values and system resilience as characterized using the method discussed above.

Table 7
R-characterization: Canonical S_0

	S_A	S_B	S_C	S_D	S_0
Delivered quanta (q)	2	3	1	0	$n=3, \sum q = 6$
In-degree (d)	3	5	2	0	$\sum d = 10$
Systems (s)	3	5	1	1	$\sum s = 10$
Connections (c)	7	11	3	2	$\sum c = 23$
Resilience (R)					$R = (3 \times 6) / (23 + 10 + 10) = 0.4186$

The example of Figure 3 shows that S_D provides inputs to both S_A and S_B . Hence, despite the redundant delivery of f by three separate systems, two of these (S_A and S_B) are directly dependent on input from S_D . Depending on the importance of its contribution in the overall delivery of f , this might suggest S_D is highly critical to system performance and function delivery. Intuitively, this should decrease resilience. Its criticality to the overall system can be demonstrated by removing it from the system and re-characterizing resilience as if it were not involved. Table 8 demonstrates the difference in resilience if S_D is *not* part of S_0 and hence not involved in delivering the valued function.

Table 8
R-characterization: S₀ without S_D

	S _A	S _B	S _C	S ₀
Delivered quanta (<i>q</i>)	2	3	1	$n=3, \sum q = 6$
In-degree (<i>d</i>)	2	4	2	$\sum d = 8$
Systems (<i>s</i>)	3	5	1	$\sum s = 9$
Connections (<i>c</i>)	6	10	3	$\sum c = 19$
Resilience (R)				$R = (3 \times 6) / (19 + 9 + 8) = 0.5$

Note the significant *increase* in resilience (+19%) when delivery of *f* is less highly dependent on one system (S_D).

Similarly, it also seems intuitive that high complexity in a specific subsystem might contribute to lower resilience values. Obviously, S_B is the most complex system in S₀ so it would be instructive to see the impact if that system were uninvolved. Note that removing S_B *also removes* 3 quanta of delivered valued function and one redundant provider of the function, reducing *q* to 3 and *n* to 2. Table 9 shows resilience characterization assuming S_B was not part of the overall system.

Table 9
R-characterization: S₀ without S_B

	S _A	S _C	S _D	S ₀
Delivered quanta (<i>q</i>)	2	1	0	$n=2, \sum q = 3$
In-degree (<i>d</i>)	3	1	0	$\sum d = 4$
Systems (<i>s</i>)	3	1	1	$\sum s = 5$
Connections (<i>c</i>)	7	2	1	$\sum c = 10$
Resilience (R)				$R = (2 \times 3) / (10 + 5 + 4) = 0.3158$

Here, comparing the result with that of Table 7, it is shown that the reduction in complexity *does not fully compensate* (~25% lower) for the reduction in functional redundancy and quanta in delivery of the valued function.

Alternatively, it would be interesting to perform a direct comparison by *decreasing the complexity* of S_B while leaving it in the system. Table 10 shows the outcome when S_B is hypothetically given the same decomposition and complexity as S_A .

Table 10
R-characterization: S_0 with Reduced Complexity S_B

	S_A	S_B	S_C	S_D	S_0
Delivered quanta (q)	2	3	1	0	$n=3, \sum q = 6$
In-degree (d)	3	3	2	0	$\sum d = 8$
Systems (s)	3	3 (as in S_A)	1	1	$\sum s = 8$
Connections (c)	7	7 (as in S_A)	3	2	$\sum c = 19$
Resilience (R)					$R = (3 \times 6) / (19 + 8 + 8) = 0.5143$

Comparing the resilience characterization of Table 10 (0.5143) with that of Table 7 (0.4186, where S_B was included with its full complexity) demonstrates a trade between complexity and resilience (lower complexity results in a resilience increase of approximately ~23%). Of course, this hypothetical example assumes S_B could still deliver 3 quanta of the valued function without its additional complexity and this is likely to be impossible in a real system. If the reduction in complexity of S_B resulted in a reduction of valued function delivery (for example, to 2 quanta, like S_A), the same comparison can be made with $q=5$, resulting in $R=0.4286$ which is very close (~2%) to the original resilience characterization shown in Table 7. This demonstrates the trade space of complexity, function, and resilience.

As earlier alluded, the decomposition of the system must be fair and accurate. Even if system modeling compromises are required in order to gain consensus among experts, this should not ultimately be deleterious to the characterization of resilience (that is, system models will not necessarily be decomposed to a level that is acceptable to all,

but should still render useful values). Table 11, Table 12, and Table 13 demonstrate resilience characterization given the assumption that *none* of the systems in Figure 3 had any decomposition (that is, imagine S_A and S_B are “simple” systems like S_C and S_D).

Table 11
R-characterization: Full S_0 (No Decomposition)

	S_A	S_B	S_C	S_D	S_0
Delivered quanta (q)	2	3	1	0	$n=3, \sum q = 6$
In-degree (d)	1	1	2	0	$\sum d = 4$
Systems (s)	1	1	1	1	$\sum s = 4$
Connections (c)	3	3	3	2	$\sum c = 11$
Resilience (R)					$R = (3 \times 6) / (11 + 4 + 4) = 0.9474$

If this result is compared to the result of Table 7, it can be seen that the relative simplicity of the systems increases the resilience of the overall system.

Table 12
R-characterization: S_0 without S_D (No Decomposition)

	S_A	S_B	S_C	S_0
Delivered quanta (q)	2	3	1	$n=3, \sum q = 6$
In-degree (d)	0	0	2	$\sum d = 2$
Systems (s)	1	1	1	$\sum s = 3$
Connections (c)	2	2	3	$\sum c = 7$
Resilience (R)				$R = (3 \times 6) / (7 + 3 + 2) = 1.5$

Table 13
R-characterization: S_0 without S_B (No Decomposition)

	S_A	S_C	S_D	S_0
Delivered quanta (q)	2	1	0	$n=2, \sum q = 3$
In-degree (d)	1	1	0	$\sum d = 2$
Systems (s)	1	1	1	$\sum s = 3$
Connections (c)	3	2	1	$\sum c = 6$
Resilience (R)				$R = (2 \times 3) / (6 + 3 + 2) = 0.5455$

Obviously, when all subsystems share similar complexity, any resilience gain expected from eliminating S_B (formerly of high complexity) is further attenuated because the impact from losing the quanta of valued function is far greater. This can be observed by comparing the difference in characterized resilience from Table 7 and Table 9 (where removing S_B reduces resilience by approximately 25%) with the difference in Table 11 and Table 13 (where removing S_B reduces resilience by approximately 42%).

Quantum Resilience, Reductionism, and Complexity

There do seem to be many “systems” that seem to fit the cliché phrase: “the whole is more than the sum of its parts.” For example, it is obvious that somehow “intelligence” emerges from the interaction of 100 billion (or so) neurons in a human brain. As a valued function, intelligence will prove notoriously difficult to quantify, but if quantifying such an idea becomes important, ultimately consensus (or at least grudging concession) will materialize. What is *not* debated is the connectionist structure of the underlying “intelligence delivery system” and its interfaces to the “real world” via sensory stimulus. What is also clear is that unless a particular “critical mass” of neuron count and connectedness is achieved, intelligence will not emerge. This is where the worlds of complexity and reductionism collide. That there is a physical system, with an actual structure, and with actual connectedness, is unquestioned. That intelligence is a valued function that emerges from it is unquestioned. How such a function emerges from such a system may continue to be an open question, but if we wish such a function to persist, the specific structure and organization that provides it must be codified in any characterization of that system’s resilience.

Quantum resilience could be considered reductionist since it starts with a hierarchical approach to quantifying system complexity. The “reductionist” epithet is sometimes intended as a stinging blow, but recall that reductionism is *not* a dirty word to engineers and they are to a great extent immune to the indictment. Engineers are required to make things work and they have an excellent track record of doing so within the paradigm of reductionism. Plus, while complexity theories are still brewing, reductionism is the *only* available paradigm. In this regard, if quantum resilience is reductionist, it is unapologetically so.

Nevertheless, any assertion that quantum resilience is *merely* reductionist is also easily deflected. Note first that quantum resilience specifically recognizes complexity in the denominator of the resilience characterization formulation. This is no accident. Quantum resilience amalgamates a number of approaches to calculating complexity in its formulation. Recognizing Kolmogorov, Gell-Mann, and Simon, as well as understanding systems as having a “fractal” dimension makes quantum resilience the most obviously complexity-oriented resilience theory available. In measuring hierarchical decomposition and intersystem relationships, quantum resilience has the best opportunity to correctly represent system complexity. Quantum resilience specifically acknowledges the “limits” of reductionism (though, note well, these limits rarely impact the ability to offer practical solutions to common problems) and specifically includes system relationships and dependencies in the formulation. Further, quantum resilience absolutely forces transparency and consistency in system modeling and permits it to be done with commercially available tools to ensure repeatable results. This ensures that engineers are not hiding behind complexity.

In addition to its firm grounding in general systems theory, development of quantum resilience purposefully considered complexity ideas. In fact, not only does quantum resilience take a large step to actually measuring complexity, it is architected to employ “true” complexity measures should any ever be invented (and assuming they are demonstrably better than the current approach). Further, if complexity is ever shown to impact the definition or quantification of valued functions it can be easily retrofitted in the numerator. While a future theory of complexity may contribute to “better” values for resilience characterization, is unlikely to further clarify the idea of resilience.

Again, note that most engineers happily and successfully solve problems using the reductionist paradigm, while a few have the luxury of indulging their interests and participating in the development of a complexity paradigm. Only a few researchers can afford to explore deeply esoteric notions that are only somewhat related to solving actual problems. Such exploration is vital and ensures continued progress as new ideas are brought to light, but for the rest of us, complexity is an aggregation of assorted and only somewhat related ideas that are not quite ready for application. To the extent they are, I have included them in the characterization formulated above.

Importantly, quantum resilience reminds us that we must not hide behind complexity. We must focus on the system discovered through analysis (where the scale is determined by the valued functions) as opposed to the system as we envision or romanticize it. Once we determine the scale at which we must operate, it becomes clear what can be accomplished whether or not the system is complex. The system is what it is and we need to focus on the parts that impact valued function. For resilience to be a meaningful concept, it must also be at some level pragmatic and operationalized.

Quantum resilience offers such a theory for now, and provides a way forward that allows future extension of its complexity measures.

Resilience and the Time Domain

As the backdrop against which all systems operate, interest in the concept of time is counterintuitively intensified by the fact that humans need *stability* from systems that are frequently very *dynamic*. In general, humans want systems that deliver their valued function(s) over long periods of time. For example, systems of government should be “stable.” A revolution every hundred years or so is about all humans can tolerate while expecting to remain productive (observation of current events is enough to demonstrate this). At a different, but still “long” time scale, telecommunications systems should be stable enough for frequent users to continue to exploit regular features even while new features are added. It is no secret that 4G LTE (fourth generation long-term evolution) is an attempt to provide such stability in the face of migration of significant underlying technologies. For ongoing maintenance of structures that evolve (like the law, or ethics, etc.) these “state changes” can be accommodated while stability is maintained (e.g., think of the impact of the anti-slavery or women’s suffrage movements which resulted in the 13th and 19th amendments to the US constitution). These systems can be viewed as stable but that doesn’t mean they are static.

Still, the need to consider time implies there might be important changes-over-time that must enter the analysis. Importantly, there are three separate ideas of such “movement” in a system:

1. System *dynamics* “happen” while a system is delivering its function. Such dynamics can be quite complicated depending on the function. For example, think of the dynamics that occur when you turn on a radio, or a printing press. As described below, the system analysis approach encapsulates (and sometimes hides) the dynamics in human-superimposed definitions of system states. For example, it is superficially easy to argue that *all* the dynamics of the system occur while it is delivering its function in a “working” or “operational” state. On the other hand, the dynamics would *not* be occurring when the system is in a “broken” or “off” state. Obviously, meaningful analysis involves deeper decomposition and greater granularity. As discussed below, required granularity of state definition can be determined on a system-by-system basis.
2. System *evolution* is part of the “adaptation” argument that I suggest should be beyond the scope of resilience theory and research. It is an important research area, but it must not be conflated with resilience. Note again that for quantum resilience it is the valued function that drives the analysis. The terminology is important here. Consider that if “values” surrounding a system change over time, these can lead to new functional demands on the system (as in the telecommunications example where what is valued now is far different than what was valued 20 years ago). This may compel an important and desirable system “evolution,” (whether internally or externally facilitated—with appropriate cautions about the teleology in the anticipated new system structure or function) but where resilience is concerned it is vital that system identity is considered. This is because if the valued functions change (in keeping with new “values”), the

resilience characterization must also change. In this way, resilience must be characterized for instances of system design (i.e., phenotypes). Evolved systems are generally new phenotypes and it is an important engineering challenge to determine whether or not they are homologous to the extent they can be compared to prior systems. Obviously, as artifacts, “designs” themselves can be considered as evolving systems, but this is a very different thing than the systems those designs represent.

3. System *decay* occurs for all systems over time. This sort of movement might be noticed as system changes result in diminishing functional delivery. Whether system decay is real or anticipated, quantum resilience is not ignorant of such movement. Though there are likely more efficient measures of such decay, if real trajectories are deemed desirable and useful, quantum resilience allows such decay to be tracked by changes in resilience characterization over time. That is, at periodic intervals the current system phenotype can be assessed for its current resilience characterization. If decay can be projected (e.g., in an anticipatory model), quantum resilience can serve in the same manner based on projected phenotypes. System decay is implicated in most discussions of thresholds found in the literature. In general, the literature takes an anticipatory position and warns of tipping points after which a system might fail. Sometimes the literature warns of cascading threshold effects. Because the importance of the system is understood, realistic approaches to such decay typically involve bolstering whatever part of the system is decaying. This places most discussion of thresholds and decay in the robustness space—protecting the system against anticipated

perturbations or projected future environments. Those scholars who anticipate unavoidable tipping points, or understand they can only be avoided at prohibitive cost, will fall back on planning for “transformation.” Of course, transformation results in a new system—which can be analyzed for its resilience as needed.

To fix the differences in the reader’s mind, consider that with, for example, the human slavery issue in the US, “the Law” (as a system; its identity; the system’s structure, relationships, etc.) did *not* change, but “laws” (operating parameters, possible configurations, etc.) changed. In this regard, one could argue there were significant “dynamics” but limited “evolution” in the formal sense. That “the law” was organized to support such revision was a matter of good design. Obviously casual use of phrases, e.g., “the law evolved,” is acceptable, but must be recognized as casual and colloquial. These are important distinctions when considering the time domain. The law did not decay and face a tipping point. Instead, society as the larger containing system forced changes. That is, societal values surrounding certain aspects of the system changed until eventually society directed system updates. After such updates, proper characterization of the resilience of the “law” as a system may or may not reflect improvements in resilience. Importantly, depending upon how *societal productivity* is quantified, a real increment in *societal resilience* might have occurred with these minor changes in the law. For example, if human capital was measured as part of the productivity numerator, it certainly can be shown to increase, or if “equity” can be somehow construed as a valued function, one could argue it was “increased” at a fairly minor cost in system complexity. Note well, however, that the changes in law (a “subsystem” of US society) contributed to

the complexity (even if minimally). Looking back on the transition, we can acknowledge that low complexity changes that result in high productivity outcomes and hence higher resilience are certainly adaptive, but it must be stressed that this is retrospective.

The “issue” of proper management of the time domain (or, more specifically, “system dynamics”) is dispelled by understanding that a system can be defined as having a (potentially large) collection of discrete configurations or states. Obviously, the granularity of state definitions will generally be an artifact of the fidelity chosen for a given system model. More will be said later but again, consider the extensive dynamics in play when a radio is turned on. While the dynamics are evident to engineers and can certainly be modeled as needed, most users would agree that a radio provides its valued function in the “on” state. The parameters and variables used to define states will usually have some tolerance (due to measurement accuracy and error), but these fluctuations can also generally be ignored once a given state is reached.

Assume a hypothetical system is in state #1 at some discrete time. As time elapses, “things change” and (depending on the granularity of the states and measured variables) eventually the system might enter state #2, and then, state #3 (or, it might revert back to state #1), and so on. Such analysis and modeling continues until the analyst is satisfied all the states are adequately codified for the level of fidelity agreed upon by the experts. This sequence of state transitions obviously shows a progression of time, but it is in discrete steps that reflect the discrete states of the system. In each of these discrete states, there may be a significant level of activity (system dynamics) by which the system will perform its valued function or functions (with whatever fidelity has been determined to be interesting or important). In many cases, however, such micro-dynamics are

unimportant to the overall character of the system (see example below). Resilience characterization requires that “system dynamics” be modeled only at the level it impacts delivery of valued function—only at the level the function *can* be valued.

It could be argued that this makes the chosen granularity an aspect of analyst “choice” or “art” and, hence, arbitrary and “not very scientific.” If “arbitrary” means a hard-won decision achieved through scientific and engineering consensus among experts, then this is true. But this is always so in any system analysis. For example, if in a particular system analysis we concern ourselves with dynamics at a period of one day, why did we not select one hour? And, had we chosen one hour, why not two minutes, or two milliseconds? In general, the level of granularity is selected at the lowest level humans can feasibly *care* about the system state, *not* at the lowest level they can assess and codify system state (this is somewhat akin to what control systems engineers consider the difference between *knowledge* and *control*). Some might argue that there is a “real” granularity, but it is easy to dispel that myth. If, for example, temperature is being measured and used in a control system, it is likely to be accomplished with a thermistor and an analog-to-digital converter that is sampled at some (humanly speaking) very fast rate. Can we track the temperature at 5 minute intervals, or must we model the system at the granularity of that very fast A/D conversion rate? If so, how does our decision change when the engineer reminds us that the A/D converter actually *oversamples* and *averages* its output before reporting it? And if we decide to consider *that* and alter our period to include some kind of averaging, what happens when the engineer reminds us that the transistors involved in the A/D conversion have their own dynamics which are faster still? Must we consider our system dynamics at *that* level? And what if the engineer

reminds us that the thermistor material itself is undergoing physical stresses that are impacting its own internal dynamics at the molecular and atomic levels? Or, that the silicon of the transistors has electron flow at yet a different rate? Must we consider that? The obvious answer is no. The idea of system dynamics has always been and will always be a matter of analyst choice. The choice made should certainly be defensible and agreed-to by the team of experts that are performing the analysis, but it will always be a choice.

The important point about the time domain is that we must realize we cannot glibly speak of resilience without any idea of why we care about a system's resilience. So it is not so much that resilience is static or dynamic, it is more that resilience is quantum. Analysts must be able to characterize system resilience no matter what state a system is in, or what activity it is performing. Resilience is increased or decreased based on how we tweak the subsystems that provide valued function. Note that this is logical given the emphasis on "structure" when resilience is characterized. Following Varela (see above), if structure is the "snapshot" of the system at a given time, then the resilience applies to the structure at that time (an *instance* of a design, or a *phenotype*). Systems can morph and change (with careful attention to identity maintenance), but resilience must mean something at a given time, so it must be defined over time scales for which it matters to those who value system functions.

Note well that discussions of system decay are actually discussions of perturbations and hence fall into the robustness space. But it is fair to consider this in light of a generalized theory of resilience. Here, "time" (or, if you will, "the whips and scorns of time") serves as the perturbation. In a "three little pigs" example, an analysis might show that a straw house is less resilient than a brick house if "huff and puff

resistance” is the only valued function. In this analysis perhaps it can be shown that the extra complexity of a brick house is easily overcome by its extraordinary wind resistance. Intuitively, it is also apparent that a straw house built 20 years ago and subjected to 20 years of wet and dry seasons might be far easier for the wolf to blow down than it was when originally erected. Recall that while as onlookers we might know that corrosion and decay has resulted in reduced function, to a resilience analysis, it is merely an instance of a new phenotype that delivers less function. Resilience might be measurably decreasing, but the only way to know that is to re-characterize the resilience in light of the decreased output of the valued function. In this case the complexity remains the same while the productivity decreases.

Also note that resilience is not about predicting precipitous failures at some future time. That is what the probabilistic risk analyses do. Monitoring trends is obviously important when we admit to human interest in maintaining (or improving) the status quo, but projections of tipping points will typically result in three human responses: (1) fix it, ensure enduring stability, and avoid the decay, (2) milk it for all its worth, and (3) cut and run. Given individual agency, purview, and hopefulness, these approaches are not mutually exclusive and all may be found on a given landscape at a variety of scales. No matter how a system phenotype is deployed (on purpose, by changes in state, or by partial or total failure of some previously productive system) resilience analysis must proceed by honestly assessing the productivity vis-à-vis valued function and by accurately documenting system complexity.

Management of Units in the Numerator

Quantum resilience is formulated to characterize resilience at any and all system scales. It is also formulated to permit multiple valued functions at any given scale. Because individual valued functions are usually expressed in different units, there is likely to be some concern over the management of units in the numerator of the quantum resilience characterization. That is, in calculating system productivity (the numerator) it is valid to sum numbers that have different units. This section demonstrates that, while this is true, it is not problematic, but if philosophical issues remain, it can be easily avoided by normalizing productivity values to percentages of total function, or rethinking the analysis to isolate functions at more appropriate levels.

Note first that resilience is a notion invented by humans. It is an idea we superimpose on systems. That is, you cannot locate resilience in a system and weigh it or measure its length. A system cannot be squeezed until its resilience drains into a graduated cylinder. This does not, however, make resilience a useless idea. Instead, resilience has become an important way humans talk about a specific character of a system. This is true for many calculated metrics that attempt to describe a character of a system. Since it is a human-defined concept, quantum resilience asserts that as a system character, *resilience can be represented with scale free and unit free values that are comparable only among consistently modeled homologous systems*. Since quantum resilience enforces transparency and demands consistency in its modeling approach, the characterization does not specifically require units. Units are used for their explanatory value to humans and to ensure consistency and consensus.

Quantum resilience characterization involves quantification of system *productivity* in the numerator. Note that system productivity is also a notion contrived to allow us to quantify the valued function of a system. Hence, even though quantum resilience requires units for quanta of valued functions, the valued functions are really just “productivity indicators” that are supplied for transparency so engineers can agree and come to consensus about how to measure the valued function of a system. It is perhaps obvious that units are used only when they are available so that consistency can be achieved both within and across models. Note well that sometimes units are unavailable (e.g., when indicating psycho-social elements of human resilience on a Likert scale). To be sure, use of quantities with definite units facilitates more rapid consensus and more accurate comparisons, but they are not specifically required.

It is best to think of the numerator as sums of quanta of functional productivity instead of sums of materials or other physical entities. In this way, a power plant that has two disparate valued functions (e.g., energy to the grid and employees in the job market) is quantized first by identifying the units (e.g., MW and head count) and second by identifying the quanta of each valued function provided by the power plant (e.g., 300 and 30). Both are important for transparency and consistency because, for example, it will prevent another engineer from assuming the units for energy are kW. The specific units, however, are unimportant once consistency is ensured because the results will only be comparable to a similarly modeled homologous system.

It is therefore somewhat important to understand that as long as consistency is employed, the units do not actually matter though, as mentioned, they provide important transparency which is vital in gaining consensus. It is not, therefore, illegitimate to

suggest that a power plant can offer 300 MW of energy to the grid and 30 employees to the job market and sum these for a “productivity” of 330. Engineers, however, are likely to be uncomfortable with such a sum for several reasons: (1) adding MW to employee count seems like bad engineering practice, and (2) it seems odd to glibly combine numbers of different orders of magnitude. The “skew” that results is sometimes difficult to swallow (i.e., does 30 employees really only constitute 9% of the overall valued function of 330?). Again, what must be stressed is that consistency matters. No numbers are lost in the calculations, and the resulting characterizations are only useful in comparison to other values that are similarly calculated.

If, however, engineering purity is required there are three approaches that can be taken. Two are simple changes to the algorithms and software, while one actually forces deeper understanding of the system. First, the software algorithms and tools can be configured to normalize all valued function contributions to percentages so that units disappear. Second, resilience characterization can be reported on a per-valued-function basis. That is, it is easy to report, for example, the resilience of the power plant in terms of *only* the energy provided to the grid, or *only* in terms of employees in the job market. Such an individual reporting makes sense in light of the (probable) fact that it is unlikely any team of engineers will ever fully exhaust a list of valued functions to the satisfaction of another team (that is, if we keep adding to our list, the numbers keep changing anyway). Both of these options are available in current versions of the software.

Third, sometimes such a dilemma is an indicator that the assignment of function to a specific system must be rethought. That is, it could be argued that employment is not a function of the power plant, but is instead a valued function of a larger employer that

operates multiple power plants. In this case, the engineers characterizing resilience might elect to remove employment from the list of valued functions offered by the plant, resolving the “units” dilemma. Note however that this decision must be consistently applied if comparisons are to be made between power plants.

Note that *weighting* certain functions as more important than others is certainly possible (e.g., using an analytical hierarchy process, cf. Saaty, 1980), but as long as consistency is used in comparisons, relative weights among valued functions do not matter. In fact, such *weighting is not recommended* because management of weights introduces yet another complication that must be tracked to ensure consistent treatment across system models.

Importantly, it should also be noted that the denominator makes some simplifying assumptions as well. For example, when the apparent complexity calculation adds number of subsystems in the decomposed hierarchy to number of dependencies between those systems, it can seem like adding apples and oranges. While specifically true, it must be observed that these are simple counts that are being added and no units are involved.

Resilient Design

Is there an archetypically “resilient” system—a system “shape” that regularly proves to be resilient? At this point it is certainly fair to propose that resilient systems are comprised of subsystems which follow three well-defined protocols: *belonging* (to the “parent”), *contribution* (toward the function of the overall system), and *connection* (with other sibling systems). Further, since connection between subsystems increases system complexity, each subsystem in an archetypically “resilient” system would be expected to exhibit *low per-member contribution (and loss)* to the valued function(s) of the system.

“Belonging” involves determining how and why a certain system fits into an overall system. In general, a “well-defined belonging protocol” simply “makes sense” and does not leave the observer wondering why a particular system should be a part of another. For example, it “makes sense” that a fuel supply chain like that employed by Phoenix has a West line and an East line, each of which provides about half of the fuel required by Phoenix (this example is discussed below). These are clearly redundant subsystems that comprise the larger system. It does not make sense, however, to suggest someone’s residence or the Phoenix Zoo be added to the Phoenix fuel supply chain. Any such inclusion would only lead to contrived arguments to defend such a system structure. Further, addition of a particular residence to the fuel supply chain can be met with the question “why not other (or all) residences?” or “is the contribution meaningful?” So the *belonging* question can be approached from both angles: why and why not. This is a vital step in system analysis as it permits scope and scale to be transparently selected.

“Contribution” usually makes it clear why a system *belongs* in another. A system’s membership in the overall system will be defined to some extent by what it

contributes to overall system function(s), but it might also be defined by the protections it brings, or the way it enables a system to survive in a given environment. Such contribution should be uncontrived, obvious, and measureable. For example, it is obvious that both the East and West lines of the Phoenix fuel supply chain make contributions to the overall system. This is obvious, uncontrived, and measurable. If a particular residence makes a contribution, then it can be included as well, if the engineers determine it is important to the overall system.

“Connection” between sibling systems is generally observed through direct interfaces, though in rare occasions can also occur indirectly or through feedback. Sometimes connections result from material exchanges and sometimes simply from sharing an environment variable (e.g., an environment temperature is an indirect connection). Other times coordinated parallelisms in sibling systems can be observed. For example, if the East line of the Phoenix fuel supply chain fails, operators can get on the phone (direct connection) to their counterparts in the West line and request increased delivery rates. Alternatively, a supervisory control system (indirect connection through another sibling system) can determine that the West line must increase delivery rates. Connections obviously contribute to the complexity of the system and are (in keeping with Simon’s nearly decomposable hierarchies) expected to diminish as the system is decomposed down to the level at which valued function is delivered.

Low per-member contribution (and loss) simply shows that the less a system contributes to the overall system function, the less it will be missed if it fails, and the more likely it is to be easily replaceable. With two “subsystems,” the percentage contribution of each line in the Phoenix fuel supply chain is fairly significant (e.g., 60-

40). It might seem intuitively obvious that adding more supply pipelines (hypothetically an additional East and West line or even a North and South line) would make the system *more* resilient (even if rather expensive). However, since this would also increase the complexity of the system in both number of subsystems and their interconnectivity, care must be taken before making a deployment decision based on intuition.

This important point can now be formalized. As described above, resilience is characterized as system productivity divided by system complexity:

$$R = \frac{P}{C}$$

As above, overall complexity is calculated as follows:

$$C = s + \sum_{i=1}^s (c_i + d_i)$$

and overall productivity is calculated as follows:

$$P = \sum_{i=1}^f \left(n_i \times \sum_{j=1}^n q_j \right)$$

where n (as before) is the number of systems providing a specific valued function.

Assuming system \mathbf{S} is comprised of function-contributing subsystems s such that the complexity of each function-contributing subsystem is the same (or nearly so):

$$C_{S_i} \cong C_{S_j}$$

And assuming that each of these subsystems contributes similar incremental quanta of valued function to the overall productivity measure such that:

$$q_{S_i} \cong q_{S_j}$$

Then, it is clear that each similar constituent subsystem contributes to resilience in equal ratios, that is, with similar numerators and denominators. Hence (noting that the second productivity sum below stops at $n-1$), for large n the following near equivalence holds true:

$$R = \frac{P}{C} = \frac{\sum_{i=1}^n (m_i \times \sum q_i)}{\sum_{i=1}^n C_{s_i}} \cong \frac{\sum_{i=1}^{n-1} (m_i \times \sum q_i)}{\sum_{i=1}^{n-1} C_{s_i}}$$

This demonstrates that when there are many systems which incrementally contribute similar quanta of the valued function, loss of a single system is not significantly deleterious to overall system resilience. This is illustrated in many of the examples that will follow.

This leads to the conclusion that the only practical design principle that comes from resilience analysis is that highest resilience results when small increments of valued function are provided by each of many redundant systems which are isolated in nearly decomposable hierarchies. If these findings are obvious and intuitive, it stems from two reasons. First, engineers learn from Nature which accomplishes its resilience in exactly that manner, and second, disapproving remarks about untoward focus on functional redundancy as the mechanism for resilience are unfounded.

INSTRUMENTING QUANTUM RESILIENCE

I invested a significant portion of my industry career developing, integrating, and using commercial model-based system engineering (MBSE) tools. For this reason, I thought it important to point out that commercial tools are already available and can easily be configured to perform the quantum resilience characterization I have proposed. This section demonstrates the ease with which two very complete and extremely capable commercially available system modeling tools were configured to perform quantum resilience characterization. Systems engineers who are familiar with tools will find these sections familiar and intuitive. Those who are not can easily skip these sections without missing anything that is particular to quantum resilience.

There are two overarching principles guiding this approach:

1. As much as possible, the default (as-shipped) data model (schema) of the commercial tool is used. This makes it easily understood by both new and experienced users of the tool.
2. As much as possible, the interface is kept simple and intuitive. There should be nothing contrived about the implementation. Any additions to data model elements should be “obvious” and use of default elements should not be forced.

Once minor schema updates were completed, all that remained was implementation of the characterization algorithm discussed above. This was done with under 500 lines of code in Java or C# (as needed) to exploit the vendor-published application programming interfaces (API) of the commercial tools. Both vendors expose a complete and powerful API that permitted the code to be nearly identical with only minor modifications.

Finally, this section demonstrates how my own FractalSys MBSE was easily adapted to support quantum resilience characterization. Importantly, it must be stressed that tools exist and can be easily adapted to implement quantum resilience. No custom tools are required.

Genesys™ for Quantum Resilience

The Genesys™ MBSE is offered by Vitech Corporation of Blacksburg, VA. This section outlines the approach taken for implementing quantum resilience in the Genesys tool and can serve as an appendix to most systems engineering specification practices guides.

Elements

Component

Components represent the physical “systems” being modeled. A component is “built from” other components (which are “built in” parent components). This is how the system hierarchy is modeled (see Figure 4). Genesys allows these hierarchies to be created easily in the user interface.

Special accommodations:

- Component has an optional “multiplicity” parameter added that allows for system boundary modeling (see Figure 5). A component with a multiplicity of zero is not included in the decomposition and complexity calculations. Instead, it is considered to be only an organizational entity. Components with multiplicity greater than one are counted as many times as indicated in the parameter value.
- Component also has its “resilience” parameter exposed.

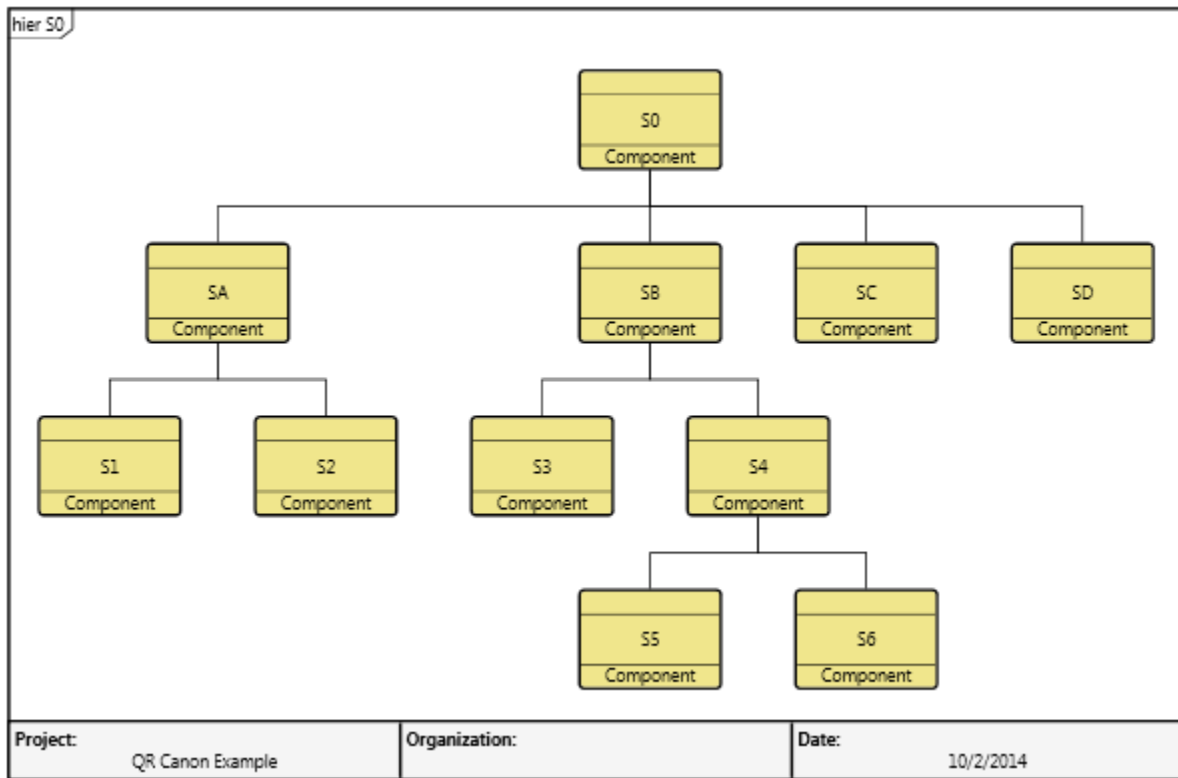


Figure 4. Genesys Representation of Canonical S_0 Decomposition

Properties - S0

Add / Remove Permissions

Multiplicity: 0 Units:

Resilience:

Figure 5. Parameters Added to "Component"

Function

Functions provide transformations and do the work of the components. Functions are “allocated to” components, and components “perform” functions. Also, a function “outputs” and “inputs” an *Item* and “services” a *Link* (see below).

Special accommodations:

- Function has been augmented in the schema to have a “quanta” parameter (visible to the user as “Quanta of Valued Function”, see Figure 6) which must be updated to contain the number of quanta of the valued function provided by the function. Though all functions are valued, not all functions are considered “valued functions” in the quantum resilience sense. Only those with values assigned to the parameters will be used in resilience characterization calculations. Additionally, a parameter called “valuedFunction” (visible to the user as “Name of Valued Function”) has been added where the “valued function” identifier (in the quantum resilience sense) is specified.

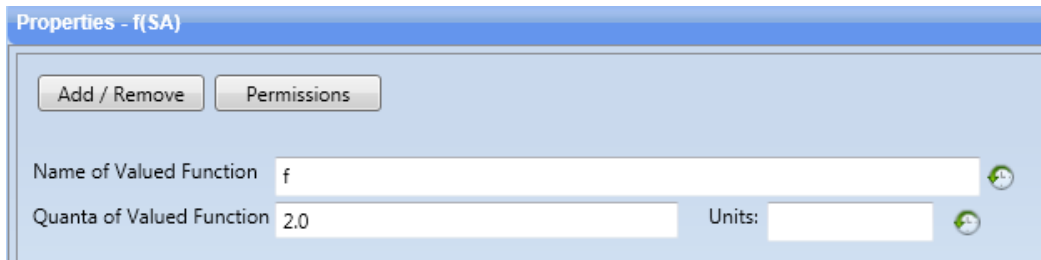


Figure 6. Parameters Added to "Function"

Item

Items represent data or material that traverses an interface between two systems. As long as good system engineering practices are used, it is not necessary to model specific interfaces or links (see Figure 7 and Figure 8) for full characterization of resilience (though it is perfectly fine to model them). An item is “input to” or “output from” a function. An Item’s input and output functions are adequate to serve as source and destination of interfaces.

It is possible that only “item” needs to be augmented by creating a new subclass called “Valued Function Output” (or similar). Then, any function can be used and

instances of that new class could have particular “valued function names” and “quanta” associated with them. For now, however, it seems simpler to merely augment the extant data model entities with parameters.

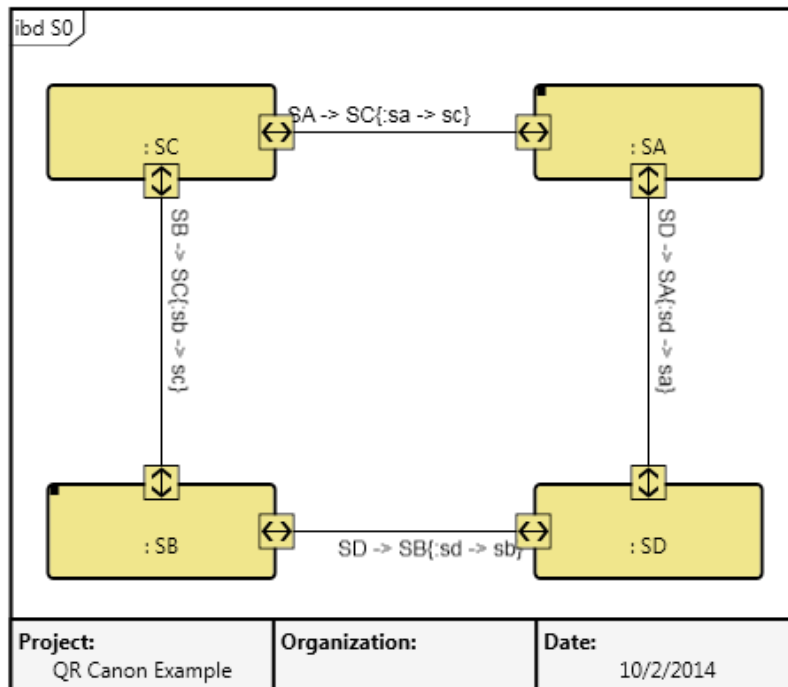


Figure 7. Block Diagram Showing Links between S0 Subsystems




Figure 8. Decomposition of SB (into S3 and S4) and of S4 (into S5 and S6)

Interface (optional)

Interfaces are logical connections between components. Note that an Interface is “comprised of” Links (see below) which are more typically the physical instantiation of an interface.

Link (optional)

Links are physical interfaces between two components. A link “connects to” a component and a link “transfers” an item. Complete system engineering approaches will require Links to be created (see Figure 7 and Figure 8).

Outputs

The calculations for quantum resilience characterization were implemented in C# for the GENESYS data model. The canonical hypothetical example discussed above results in the following outcomes:

```
-----  
R-characterization: <S0>  
-----  
<SB> in-degree: <SB rcv from SD> inputs <sd -> sb>  
<S3> in-degree: <S3 rcv from S4> inputs <s4 -> s3>  
<S4> in-degree: <S4 rcv from S3> inputs <s3 -> s4>  
<S5> in-degree: <S5 rcv from S6> inputs <s6 -> s5>  
<S6> in-degree: <S6 rcv from S5> inputs <s5 -> s6>  
<S3> connection: <S3 rcv from S4> inputs <s4 -> s3>  
<S3> connection: <S3 snd to S4> outputs <s3 -> s4>  
<S5> connection: <S5 snd to S6> outputs <s5 -> s6>  
<S5> connection: <S5 rcv from S6> inputs <s6 -> s5>  
<S6> connection: <S6 rcv from S5> inputs <s5 -> s6>  
<S6> connection: <S6 snd to S5> outputs <s6 -> s5>  
<S4> connection: <S4 snd to S3> outputs <s4 -> s3>  
<S4> connection: <S4 rcv from S3> inputs <s3 -> s4>  
<SB> connection (from valued function): <f(SB)> quanta: <3>  
<SB> connection: <SB rcv from SD> inputs <sd -> sb>  
<SB> connection: <SB snd to SC> outputs <sb -> sc>  
SB indegree: 5, decomp: 5, connec: 11, complexity: 21  
  
<SD> connection: <SD snd to SA> outputs <sd -> sa>  
<SD> connection: <SD snd to SB> outputs <sd -> sb>  
SD indegree: 0, decomp: 1, connec: 2, complexity: 3  
  
<SA> in-degree: <SA rcv from SD> inputs <sd -> sa>  
<S2> in-degree: <S2 rcv from S1> inputs <s1 -> s2>  
<S1> in-degree: <S1 rcv from S2> inputs <s2 -> s1>  
<S2> connection: <S2 snd to S1> outputs <s2 -> s1>  
<S2> connection: <S2 rcv from S1> inputs <s1 -> s2>  
<S1> connection: <S1 rcv from S2> inputs <s2 -> s1>  
<S1> connection: <S1 snd to S2> outputs <s1 -> s2>  
<SA> connection (from valued function): <f(SA)> quanta: <2>  
<SA> connection: <SA snd to SC> outputs <sa -> sc>  
<SA> connection: <SA rcv from SD> inputs <sd -> sa>  
SA indegree: 3, decomp: 3, connec: 7, complexity: 13  
  
<SC> in-degree: <SC rcv from SB> inputs <sb -> sc>  
<SC> in-degree: <SC rcv from SA> inputs <sa -> sc>  
<SC> connection (from valued function): <f(SC)> quanta: <1>  
<SC> connection: <SC rcv from SB> inputs <sb -> sc>  
<SC> connection: <SC rcv from SA> inputs <sa -> sc>  
SC indegree: 2, decomp: 1, connec: 3, complexity: 6  
  
System: <SB>, Function: <f>, quanta: 3  
System: <SA>, Function: <f>, quanta: 2  
System: <SC>, Function: <f>, quanta: 1  
  
S0 R-characterization: 0.4186, quanta: 6, numsys: 3, complexity: 43
```

```
System complexity contribution:
  SB: 48.8372%
  SD: 6.9767%
  SA: 30.2326%
  SC: 13.9535%
System productivity contribution:
  SB: 50.0000%
  SA: 33.3333%
  SC: 16.6667%
Function productivity contribution:
  f: 100.0000%
```

Though the canonical example is quite simple—the intent was to allow the illustrative calculations to be accomplished mentally—it is clear that it establishes a baseline for resilience characterization. Larger systems quickly become impossible to manage, thus requiring a modeling tool. Once modeled, system engineers interested in changing parts of the system, adding interfaces, decomposing systems, etc., would very quickly understand how such changes impact the overall resilience of their system.

Innoslate™ for Quantum Resilience

The Innoslate™ MBSE is offered by SPEC Innovations of Manassas, VA. This section outlines the approach taken for implementing quantum resilience in Innoslate and can serve as an appendix to most systems engineering specification practices guides.

Elements

Asset

Assets represent the physical “systems” being modeled. An asset is “decomposed by” other assets. This is how the system hierarchy is modeled (see Figure 9). Innoslate allows these hierarchies to be created easily in the user interface.

Special accommodations:

- Asset has an optional “multiplicity” property/attribute added that allows for system boundary modeling (see Figure 10). An asset with a multiplicity of zero is not included in the decomposition and complexity calculations. Instead, it is considered to be only an organizational entity. Assets with multiplicity greater than one are counted as many times as indicated in the parameter value.
- Asset also has a “resilience” attribute added to store the characterization value (see Figure 10).




Figure 9. Innoslate Representation of Canonical S_0 Decomposition

Attributes	
Name	<input type="text" value="S0"/>
Number	<input type="text"/>
Description	<div style="border: 1px solid #ccc; padding: 5px;"> <div style="font-size: 0.8em; margin-bottom: 5px;"> <> B <i>I</i> S ¶ ☰ ☷ ☹ ☲ ☱ ☳ ☴ ☵ ☶ ☷ </div> <div style="height: 100px;"></div> </div>
Multiplicity	<input type="text" value="Value"/> ▾ <input type="text" value="0"/>
Resilience	<input type="text" value="Value"/> ▾ <input type="text" value="0.4186046511627907"/>

Figure 10. Attributes Added to "Asset"

Action

Actions provide transformations and do the work of the assets. Actions are “performed by” assets, and assets “perform” actions. Also, an action “generates” and “receives” an *Input/Output* item.

Special accommodations:

- Action has been augmented in the schema to have a “quanta” attribute (visible to the user as “Quanta of Valued Function”, see Figure 11) which must be updated to contain the number of quanta of the valued function provided by the action.

Though all actions are valued, not all actions necessarily deliver “valued functions” in the quantum resilience sense. Only those with values assigned to the attributes will be used in resilience characterization calculations. Additionally, an attribute called “Valued Function Name” has been added where a string identifying the “valued function” (in the quantum resilience sense) is specified.

Attributes	
Name	<input type="text" value="f(SA)"/>
Number	<input type="text"/>
Description	<div style="border: 1px solid gray; padding: 5px;"> <> B I ↺ ↻ ☰ ☷ ☹ ☲ ☳ ☴ ☵ ☶ ☷ <input style="width: 100%; height: 60px;" type="text"/> </div>
Quanta of Valued Function	<input type="text" value="Value"/> ▾ <input type="text" value="2"/>
Valued Function Name	<input type="text" value="f"/>

Figure 11. Attributes Added to "Action"

Input/Output

I/O items represent data or material that traverses an interface between two systems. As long as good system engineering practices are used, it is not necessary to model specific interfaces or links (see below) for full characterization of resilience. An action “generates” and “receives” I/O items. Since actions are performed by assets, an input/output item’s source and destination serve to anchor both sides of the interface. Figure 12 shows an I^2 chart of the interfaces and the I/O items in the canonical example.

Logical Connection (optional)

Interfaces are logical connections between components. Note that a Logical connection “connects to” Assets.

Conduit Connection (optional)

Conduits are physical interfaces between two components. A conduit “connects to” an asset and “transfers” Input/Output items. Complete system engineering approaches will require Conduits to be created.

SA		SA-SC sa->sc	SD-SA	
	SB	SB-SC sb->sc	SD-SB	
SA-SC	SB-SC	SC		
SD-SA sd->sa	SD-SB sd->sb		SD	

Figure 12. I² Chart of Input/Output Items

Outputs

The calculations for quantum resilience characterization were implemented in Java for the Innoslate data model. The canonical hypothetical example results in the following outcomes:

```
<SC> in-degree: <SC rcv from SB> inputs <sb->sc>
<SC> in-degree: <SC rcv from SA> inputs <sa->sc>
<SC> connection (from valued function): <f(SC)> quanta: <1.0>
<SC> connection: <SC rcv from SB> inputs <sb->sc>
<SC> connection: <SC rcv from SA> inputs <sa->sc>
SC indegree: 2, decomp: 1, connec: 3, complexity: 6
```



```

<SD> connection: <SD snd to SA> outputs <sd->sa>
<SD> connection: <SD snd to SB> outputs <sd->sb>
SD indegree: 0, decomp: 1, connec: 2, complexity: 3

<SA> in-degree: <SA rcv from SD> inputs <sd->sa>
<S1> in-degree: <S1 rcv from S2> inputs <s2->s1>
<S2> in-degree: <S2 rcv from S1> inputs <s1->s2>
<S1> connection: <S1 snd to S2> outputs <s1->s2>
<S1> connection: <S1 rcv from S2> inputs <s2->s1>
<S2> connection: <S2 snd to S1> outputs <s2->s1>
<S2> connection: <S2 rcv from S1> inputs <s1->s2>
<SA> connection (from valued function): <f(SA)> quanta: <2.0>
<SA> connection: <SA rcv from SD> inputs <sd->sa>
<SA> connection: <SA snd to SC> outputs <sa->sc>
SA indegree: 3, decomp: 3, connec: 7, complexity: 13

<SB> in-degree: <SB rcv from SD> inputs <sd->sb>
<S3> in-degree: <S3 rcv from S4> inputs <s4->s3>
<S4> in-degree: <S4 rcv from S3> inputs <s3->s4>
<S5> in-degree: <S5 rcv from S6> inputs <s6->s5>
<S6> in-degree: <S6 rcv from S5> inputs <s5->s6>
<S3> connection: <S3 rcv from S4> inputs <s4->s3>
<S3> connection: <S3 snd to S4> outputs <s3->s4>
<S5> connection: <S5 rcv from S6> inputs <s6->s5>
<S5> connection: <S5 snd to S6> outputs <s5->s6>
<S6> connection: <S6 snd to S5> outputs <s6->s5>
<S6> connection: <S6 rcv from S5> inputs <s5->s6>
<S4> connection: <S4 snd to S3> outputs <s4->s3>
<S4> connection: <S4 rcv from S3> inputs <s3->s4>
<SB> connection (from valued function): <f(SB)> quanta: <3.0>
<SB> connection: <SB rcv from SD> inputs <sd->sb>
<SB> connection: <SB snd to SC> outputs <sb->sc>
SB indegree: 5, decomp: 5, connec: 11, complexity: 21

System: <SC>, Function: <f>, quanta: 1.0
System: <SA>, Function: <f>, quanta: 2.0
System: <SB>, Function: <f>, quanta: 3.0

R-characterization: 0.4186, quanta: 6.0, numsys: 3, complexity: 43
System complexity contribution:
    SC: 13.9535%
    SD: 6.9767%
    SA: 30.2326%
    SB: 48.8372%
System productivity contribution:
    SC: 16.6667%
    SA: 33.3333%
    SB: 50.0%
Function productivity contribution:
    f: 100.0%

```

FractalSys for Quantum Resilience

FractalSys is the author's model-based system engineering tool, a brief overview of which is contained in Appendix A. This section outlines the augmentation of

FractalSys to support quantum resilience characterization. It not only lists the changes made to instrument the calculation of system complexity (in-degree, connections, and decomposition) and tracking of system productivity vis-à-vis valued function and output quanta, but also demonstrates its use. With a few simple adjustments and methodological constraints, FractalSys can be used to characterize and compare resilience of modeled systems. Impact to FractalSys is discussed below as it pertains to *Systems*, *States*, *Transitions*, and *Variables*. Though not strictly required, *Functions* have been added to FractalSys to support quantum resilience more clearly. This is discussed below. The FractalSys overview (see Appendix A) should be consulted to clarify assumptions made herein.

It is important to remember that resilience is characterized as a function of system *complexity* and *productivity* as described by quantum resilience. System complexity is a function of system *decomposition* into subsystems, *in-degree* as a measure of dependence on other systems, and *connectivity* (in-degree *and* out-degree) with other systems. Measured this way, quantum resilience acknowledges both hierarchical and relational complexity. System productivity is calculated based on the *valued function* and the total output of the system as measured in *quanta* of the valued function. This document outlines how these key concepts (italicized) are managed in FractalSys.

Systems

The way systems are managed directly impacts the calculation of *decomposition* which contributes to system *complexity*. System hierarchies are adequately managed by FractalSys, but the overall containing system for which resilience characterization will be

done should be thought of as an “analysis boundary” (instead of—strictly speaking—a real “parent” system) though the distinction is somewhat moot. As shown in the canonical example above, the analysis boundary (S_0) is drawn with a dotted-line to indicate it is an analysis choice instead of a “real” system. S_0 effectively contains all systems the analyst deemed important to providing the valued function(s). Note that since quantum resilience expects “real” parent systems to be included in complexity measures, such a choice allows the artificial system boundary to not be implicated in the “decomposition” part of the complexity calculation. This can be accomplished by assigning the purely “organizational” systems a *multiplicity* of 0 in FractalSys.




Figure 13. FractalSys Representation of Canonical S_0 Decomposition

Figure 13 illustrates how the hypothetical system in the canonical example would be represented in FractalSys. Note that S_0 appears as a typical system in a typical hierarchy. By counting the number of systems in the hierarchy, its decomposition (self-included) is easily seen to be 11. Since S_0 is an artificial analysis boundary, however, it should not be counted in the analysis. FractalSys supports this as shown in Figure 14.

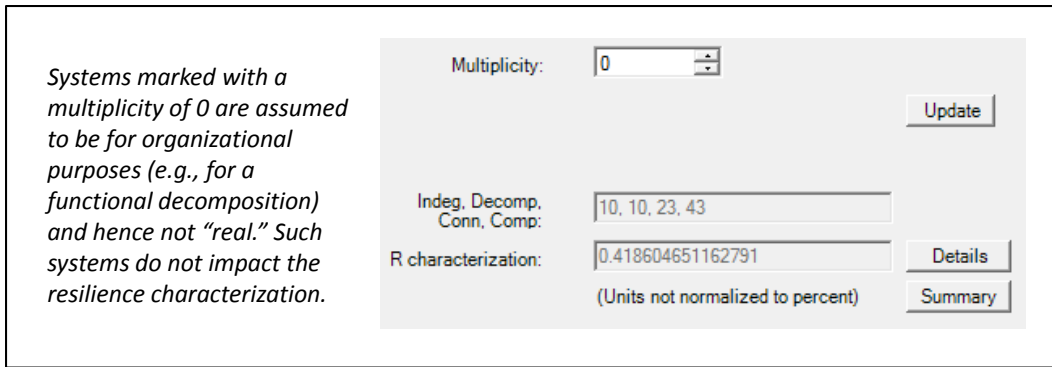


Figure 14. Clarification of “System Boundary” and “Parent System”

Note that FractalSys accurately reflects the decomposition of S_0 (in Figure 14) and only the decomposition of the systems implicated in the delivery of the valued function is counted in the complexity metric which is the sum of the in-degree, decomposition, and connectivity values ($10+10+23=43$).

States

Recognizing that system “state” can be an adequate proxy for “function” is important (see FractalSys overview). This is a fairly logical assumption considering that when a system is *in* a certain state (as indicated by the values of state variables selected and defined by human observers) it is likely to be undergoing certain dynamics in performing a certain function or functions. Since quantum resilience emphasizes *valued function*, FractalSys will allow identification of a particular system state as the one in which the system delivers its valued function. This does *not* mean a “state” can be resilient and, in fact, it is incorrect to refer to a “state” as resilient. This terminology has been casually employed in the literature, but it is vital to remember the “system” is what will be characterized for resilience, not a state.

Functions have been added to FractalSys to support the “valued function” concept (see Figure 15). Each system can support an arbitrary number of functions, but in practice, these are generally short lists. When a system state is defined, the user specifies the number of quanta provided for each function. These can remain zero as necessary. As shown in Figure 15, system S_B , provides 3 quanta of function f while operating in a state that has been arbitrarily named “providing f .”

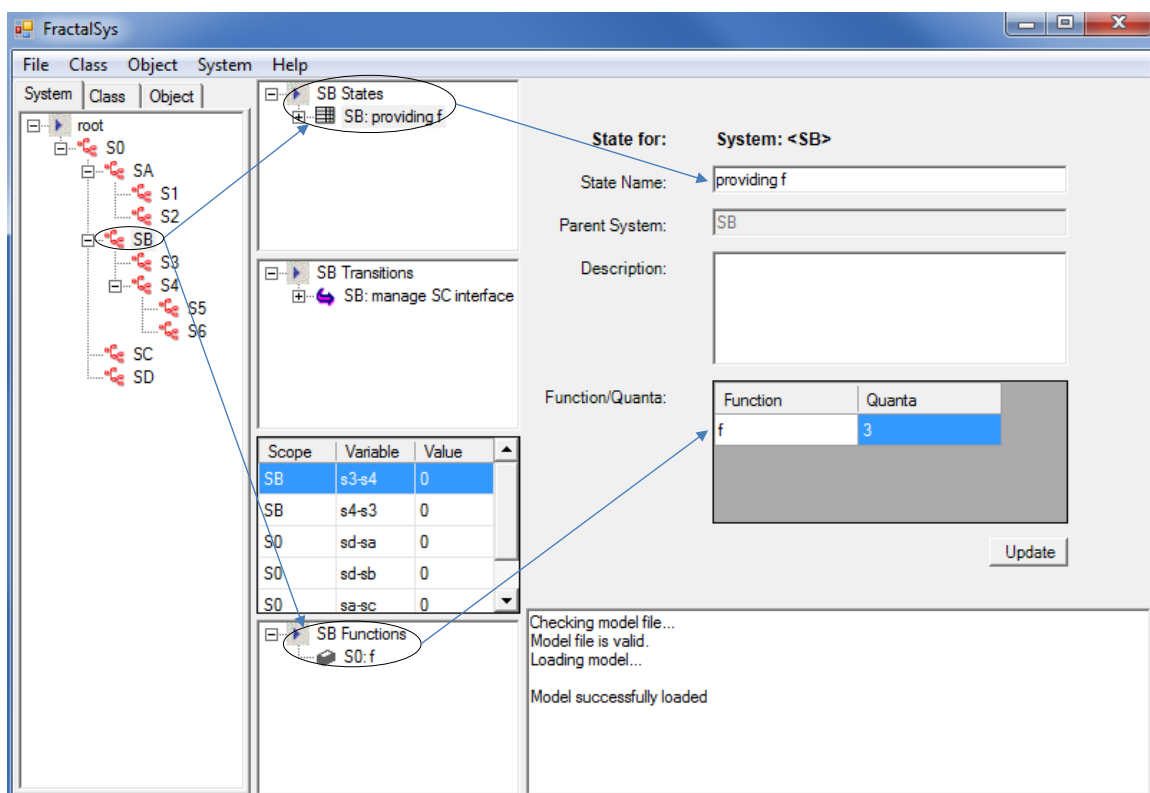


Figure 15. FractalSys State Panel Showing Provision of Valued Function

Note also that provision of a valued function is considered to be *one* external connection that contributes to the complexity measure. This is because, in general, the valued function will be consumed by some external system and does, in that way, introduce another connection to the system.

Variables

FractalSys assumes the extent to which variables are shared by systems defines the *connectedness* of those systems. Setting a variable that is used by another system increases *connections* for both systems. Using a variable which is set or controlled by another system increases *in-degree* for the using system. FractalSys also recognizes indirect connectivity. That is, if two otherwise independent systems use a variable set by a third system, they are indirectly connected through that variable. Such indirect connectivity is not currently exploited in resilience characterization though it is certainly within reach of FractalSys to measure this if it is deemed important. At this stage, since it creates no real connection or dependency, it is ignored.

FractalSys variables are used to ensure proper allocation of *in-degree* and *connections* in system complexity calculations. Obviously, any variable sharing between systems constitutes a connection (and recall that data sharing in FractalSys can represent material or non-material exchanges). Such *connections* will be counted based on *unique* variable names and multiple uses of the same variable between the same systems do not increase the overall count (though obviously this can be adjusted as necessary).

As discussed in the FractalSys overview, any *state* or *transition* that uses a variable defined in an *external* scope (that is, *inherited*, see FractalSys overview) indicates a system dependency. This serves to increment *in-degree* for that system. Similar to connection counting, variables incrementing in-degree will count only once per unique variable. The FractalSys overview also suggests that FractalSys does not care about directionality of interfaces since there was no driving reason to support such ideas

at its design. Since this is the case, an “incoming” data item would also serve to artificially inflate “out-degree” for the receiving system. This would result in situation in which in-degree and overall connections were always equivalent. To avoid this, a specific mechanism will be used to properly manage in-degree for resilience characterization. This is discussed in the next section.

Since resilience characterization requires the idea of “in-degree” to identify system dependencies, it may be easier on the user if they could specify interface directionality in FractalSys. This update can be considered as more experience is gained.

Transitions

To properly manage in-degree, a specific transition (and transition entry) must be created that “sets” the variable in the system which “owns” the variable within its scope. In a properly modeled system, such a mechanism will generally emerge as a matter of course (that is, there is nothing forced about it), but since it is specifically implicated in calculation of *in-degree*, it merits an example. The canonical example provides an adequate demonstration.

Figure 16 shows the S_0 decomposition and the specific implementation of the S_A - S_C interface. For ease of discussion, the data (or material) sent from S_A to S_C is referred to by the variable *sa-sc*. The hypothetical system proposed in the example shows this to be a unidirectional output from S_A that is presumably required by S_C in generating a valued function (otherwise it would not be modeled). The interface implied by the *sa-sc* exchange constitutes a connection for both S_A and S_C . As expected, this exchange will increment the total number of connections for these systems.

Since S_C uses a variable that comes from an external scope (see bottom center of Figure 16 where the S_A variables are listed and variable $sa-sc$ is shown to be within the S_0 scope), it also constitutes an input connection that increments in-degree for S_C . Hence, from this exchange (and as expected), S_C has a connection, and an additional in-degree. Note that since FractalSys interfaces are bidirectional, the same would be true for S_A . This would clearly go against the intent of the system engineer who (based on the directional arrow) specifically intended a unidirectional interface from S_A to S_C . To resolve this confusion, S_A can assert the direction of the interface by explicitly using the FractalSys “SET” keyword.

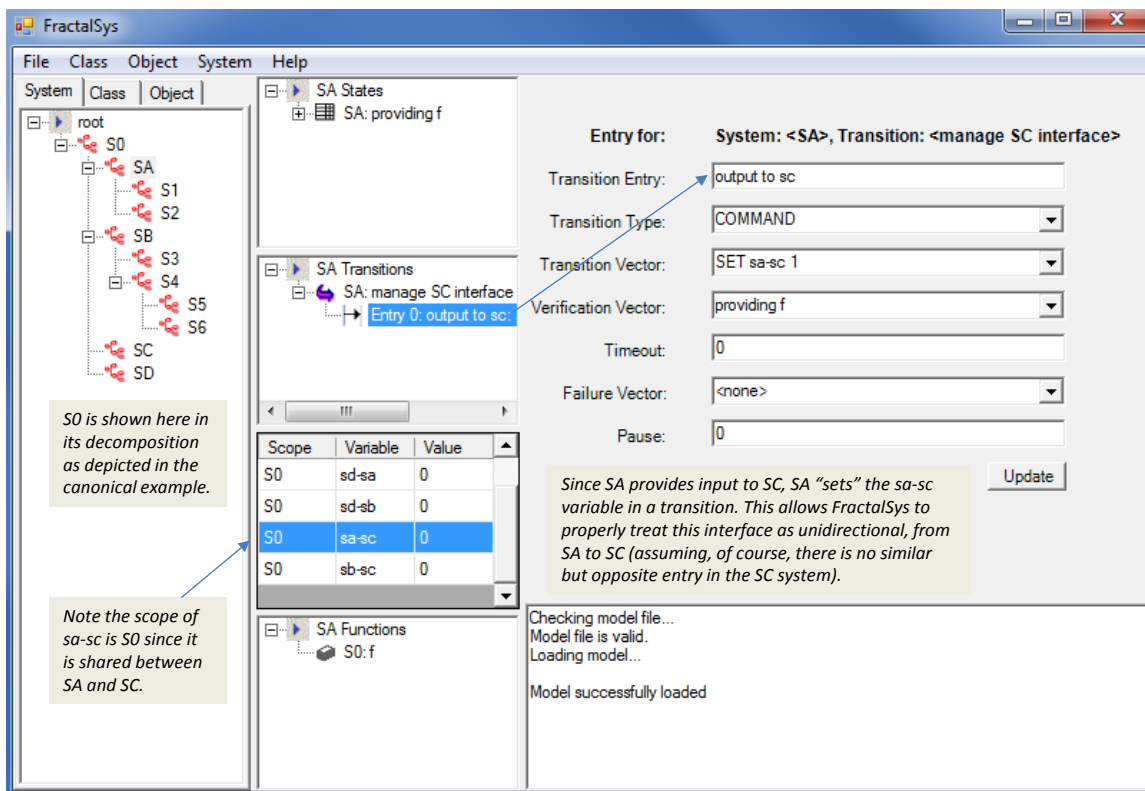


Figure 16. Example of Interface Modeling

“SET” is a reserved keyword in FractalSys that facilitates system simulation when FractalSys is not integrated with real systems that externally control variable value updates. To use SET to update a variable, it must be placed in a transition entry (and therefore, within a transition). As shown in the upper right panel of Figure 16, a transition and entry have been created for S_A in which the sa-sc variable is specifically “SET” to a value in the transition vector. For this example, though “1” is used, the actual value is inconsequential, but will likely make more sense for a real modeled system. Since S_A asserts control over this variable update with a SET keyword, FractalSys does not include this in its counting of in-degree for S_A . Hence, for S_C , both overall connection count and in-degree are incremented, while for S_A , only the overall connection count is incremented.

In this hypothetical example, such an approach may appear to be an onerous burden on the system modeler, but in practice, such interfaces are managed as a matter of course during a complete system analysis. Variables are set and used as needed and FractalSys manages the bookkeeping. Further, as indicated above, future modifications to FractalSys can be considered that specifically allow the specification of unidirectional interfaces.

Other Notes

System identity can be rigorously controlled (modeled and enforced) in FractalSys. Significant nuance can be shown since an arbitrarily large number of states and relationships are supported. Note that when multiple states are defined, a system is required to be *in* one of the states. If no defined state matches current variable values or

subsystem configurations, resolution of the unknown state is required. A model wherein a system can be in an unknown state is considered incomplete or underspecified and must be more fully defined. This also includes the enforcement of managing transitions between states based on changes in variables (i.e., “measurables” or telemetry). While transitions are not an adequate proxy for “function”, it can readily be seen that regular or periodic transitions between states can be a way to model (or even simulate) larger “functions.” The engineer is warned, however, that the need for such contrivances may actually be suggesting that a more “firm” state be specified for a parent system.

FractalSys *variables* can represent what Holling (and others) would suggest are the handful of fast- and slow-moving variables that drive most system states. Importantly, however, contrary to the literature, we need not be forced to limit them to a handful, or classify them as “fast” or “slow” *a priori*. Instead, if such a classification is deemed important, variables should be observed over timeframes of interest and then can be determined to be fast or slow-moving. It is important to moderate the tendency to too quickly assume what is fast and what is slow and to avoid making decisions based on such assumptions. Further, when tools like FractalSys are available to manage system models of significant size and complexity, it is unnecessary to be frugal in assignment or tracking of variables. Only when the system has been adequately codified can assessment of major driving variables and consolidation for simplicity be done.

Note that it is unlikely to be important that FractalSys is a tool for development of Anticipatory Systems (*sensu* Rosen, 1985). This might become interesting or useful in the future, but as currently envisioned, these features are not intended for use in quantum resilience characterization work.

An argument could be made that quantum resilience *only* characterizes resilience with respect to specifically identified valued functions/services and does not fully characterize a system. This is only true to the extent that the analysis (1) has not fully documented the system, and (2) has not fully identified all of the functions of value. If a system analysis is incomplete, the resilience characterization may be incomplete (but likely will still be very useful). This is not, however, an argument against quantum resilience. Instead it is an argument that suggests analyses must be completed and properly vetted. If a variety of valued functions are not identified, then yes, the characterization of the system will be more directly tied to those that are identified and it may discount those functions that are underrepresented. Again, an incomplete model can still be useful, but the applicability might be limited by the failure to acknowledge other valued functions or other contributing systems. Tools like FractalSys make it very easy to incrementally embellish models as needed and ensure transparency throughout the entire process. System modeling has a long history of balancing fidelity with cost/benefit trades. A similar “sweet-spot” is likely to be found with resilience characterization modeling.

ILLUSTRATING QUANTUM RESILIENCE

As a generalized theory, quantum resilience can be applied to all system regimes and all disciplines. Obviously, some systems are more physical than others, and gaining consensus among experts over the methods of modeling and measurement is more easily achieved. Some systems (especially psycho-social systems that are predominately about human resilience) will require more work to gain consensus. Importantly, quantum resilience ensures transparency and consistency in managing these systems. System regimes remain loosely defined, but involve the following rough delineations:

1. Mostly-human-engineered systems (e.g., the Internet, intelligence surveillance, infrastructure),
2. Mostly-natural systems (e.g., ecosystems like salt marshes, boreal forests),
3. Mixed socio-ecological systems (e.g., catchments and basins with economies based on mixed grazing and agriculture), and,
4. Mostly-human systems (e.g., organizations, governments).

The brief examples below serve to not only introduce a variety of system classes and illustrations on approach but also will instruct and sometimes correct intuition about resilience. In general, these examples are too simplistic to provide usable data, but they can facilitate discussion and provide fodder for future work. These examples are intended to build confidence in the method and tools and do not necessarily offer solutions.

Though they could easily be replicated in any of the aforementioned model-based systems engineering tools, all the examples documented herein have been modeled and characterized using the author's FractalSys tool.

Lake-Agriculture System

The history outlined by Carpenter et al. (2001) for a representative lake-agriculture system is a litany of protecting a natural lake from an onslaught of events it was not “designed” to withstand. In fact, from the analysis, the only conclusion that can be drawn is that the lake is *not* resilient and hence needs human-designed protection. Throughout, the degradation of the lake system away from its pristine condition is conflated with a degradation of resilience. This positions a former “state of nature” as the ideal high mark for resilience, which Carpenter et al. suggest was available in 1840. It is fairly typical for the eco-resilience literature to make this tenuous equation, so it must be made clear that they are speaking of ecological health and not resilience.

Instead of a notional analysis based on the adaptive cycle, the lake-agriculture system proposed by Carpenter et al. (2001) can be more completely assessed by quantum resilience. First, quantum resilience requires us to recognize that based on their discussion there are (at least) two valued functions: (1) incremental economic gain from agricultural production near the lake, and (2) some presumably measurable social utility of the lake based on (perhaps) recreational and aesthetic value (and perhaps a few vacation rentals). Note that for this particular example, management efforts were focused on restoring the clarity of the water by controlling the Phosphorous content. After redirecting sewage effluent away from the lake, the principle contributor to the nutrient increase in the lake was agricultural runoff. Note that while the lake must remain “environmentally healthy” to support aesthetic and recreational goals, its value is ultimately derived from functions that can be reduced to economic values as indicated in the proposed list of valued functions:

1. Agricultural production
2. Fishing licenses
3. Cabin rentals
4. Boat rentals
5. Boat launch fees

Obviously, the list can be argued and extended as necessary—including, if desired, specific valuation of low Phosphorous. Recall, however, that high P will be reflected in low economic return from recreation. Valuation of biophysical aspects is exemplified in the extended example of the Goulburn-Broken Catchment.

Note that “agricultural production” is forced here. If the real interest is to earn money from the real estate surrounding the lake, other options for land use are available (e.g., a housing development) and the valued function can be adjusted accordingly. Obviously, this might reveal other social dimensions that lead to other valued functions, but these can be formulated in a set of scenarios as needed. Once these valued functions have been elucidated to an extent satisfactory to the analysis team, quanta of resilience can be established and assigned (e.g., how much economic gain from agriculture over what periods and how much enjoyment of the clear water, etc.). Even with these first steps driven by quantum resilience analysis, it is clear that an actual system is being addressed and not simply a romanticized idea of a healthy lake. Certainly, subsystems can receive focused attention if desirable, but first the whole system must be properly acknowledged.

Second, observation of the scale of the system that delivers these functions can now occur. Just because, for example, economics is implicated in the valued function,

does not imply it is necessary to expand the system to include the entire global economy. If the quantum of resilience is something like net revenue per growing season from the land in the lake's vicinity, it can remain scoped fairly tightly. As suggested above, however, it must be admitted that this opens the doors to many alternative mechanisms for generating revenue, some of which do not even require that agriculture be the technology of choice. Still, if the valued function includes agriculture, alternative crops or alternative approaches to agriculture might be considered and might expand or contract the system scope (e.g., an orchard has different impact than an industrial wheat field). Further, since system structure is an important consideration, alternative protections for the lake could be envisioned that might expand the scale of the system at costs that must be acknowledged (e.g., if all runoff were captured by retaining walls and routed to a wastewater treatment plant this would change the scope to include the municipal water and sewer system). In that case, retaining walls and additional plumbing must be acknowledged to be protection-oriented and contributing to the *robustness* of the system in a way that may not be specifically measurable as valued function. Instead, it is a contributor to system complexity that is specifically required in order to gain the other valued functions that are closely related to the clarity of the water in the lake. Such protections do not necessarily increase resilience as much as they allow the system to do what it was designed to do.

Third, clear identification of systems and subsystems that are implicated in the valued function delivery can be accomplished. Once the systems that provide the function are identified it is easy to see how they can be required to maintain their identity. It is also easy to see how much they can change without impacting reliable delivery of the

function. Recall that it is disingenuous to suggest that the system can show its resilience by delivering different functions. A lake-agriculture system cannot reasonably be changed into a parking lot if recreational use of the water is one of the valued functions. For example, if agricultural production on surrounding land is non-negotiable, the aforementioned retaining walls could deflect or control dangerous runoff events, or a higher rate of lake water replacement could be established, but you cannot simply tell farmers to not use fertilizer. Obviously costs must be carefully considered and sometimes management costs will serve to moderate just how much value there is placed on specific valued functions. Note “maintaining a pristine natural environment” is a valid valued function, but it greatly limits what can be done if economic return remains a valued function. Balance must be sought through traditional cost-benefit analysis.

Fourth and finally, it can be determined how the valued functions are (or can be) redundantly supplied. Carpenter et al. (2001) recognize a few redundant approaches and this is a step in the right direction. For example, they recognize how bio-manipulation of the lake’s food web could improve water quality. Or, as previously stated, management may involve a more delicate balance of agriculture. After considering important warnings about focusing solely on technological fixes (cf. Allenby, 2012, p. 357), managers should certainly not *preclude* the possibility of employing elaborate and expensive technological remediation. Though costly, there is no reason technology could not be put in place to prevent all agriculture run-off from entering the lake. Sometimes SES managers are hesitant to include expensive technological fixes in their solution spaces, however, it is unfair to assume that their lamented socio-political “tumult of confusion” (Carpenter et al., 2001, p. 770) is without significant cost.

A simplified lake-agriculture system is depicted in Figure 17. It is understood that a real system is far more complex than this, but this will provide several talking points. For the simple example one farm is assumed to own and operate all surrounding land. They have installed an attractive stone wall to prevent erosion and protect the lake from runoff from a part of the farm, but runoff continues to happen in other areas. A local real estate manager rents two cabins on one corner of the lake. There is an access road, a dock with rental boats and a launching ramp. Fishing licenses are sold and catch fees collected. A small economy is generated in managing recreational use of the lake.

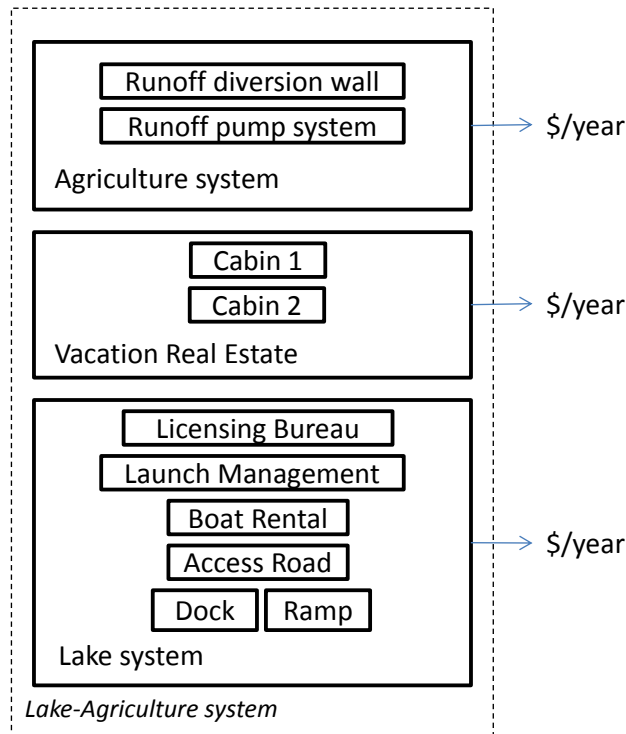


Figure 17. Simple Lake-Agriculture System

Based on the list of valued functions above, Table 14 summarizes the resilience of the system. Note that this is effectively an economic system and that its resilience is

based on the ability to generate economic value in a variety of ways. Obviously, if the lake becomes turbid, the economic value from recreation is diminished.

Table 14
R-characterization: Lake-Agriculture System

System (Function)	\$K/year
Agriculture (Economic Output)	200
Cabin 1 (Economic Output)	20
Cabin 2 (Economic Output)	20
Licensing Bureau (Economic Output)	1.5
Launch Management (Economic Output)	2.5
Boat Rental (Economic Output)	5
Productivity	1494.00
Complexity	19
R-characterization	78.6316

This provides a baseline resilience characterization for the simple system. Now assume that for some reason, the non-agricultural economic production ceased. Note well, this can be caused by any number of disturbances including eutrophication after the farmers carelessly dumped 100 tons of fertilizer into the lake, ten consecutive years of uncontrolled flooding and unrestrained runoff, a labor union strike in the recreation sector, or a bomb scare at the dock. *The particular disturbance does not matter*, but assume that eventually, all non-agriculture economic production disappears. This constitutes an alternative system deployment and the resilience can be reassessed with the results depicted in Table 15.

Table 15
R-characterization: Lake-Agriculture System (Ag Only)

System (Function)	\$K/year
Agriculture (Economic Output)	200
Productivity	200.00
Complexity	14
R-characterization	14.2857

Despite the marginal contribution (~20%) of the non-agricultural economic production in the original assessment, the resilience is dramatically reduced by the loss of the recreation sector. Though in this case, it is clearly an artifact of the way the system is modeled, it is a good example of how incremental delivery of valued function contributes to resilience. In a more thoroughly modeled system, the negative impact of the non-agricultural productivity loss might not have been so dramatic, or it might have been worse. This reiterates the importance of gaining consensus on the system modeling but also demonstrates the power of the approach in instructing engineers about resilience and intuition.

Importantly, this example also demonstrates how quantum resilience reverses the equation on disturbances. As discussed at length above, *the resilience of a system cannot be based on disturbances*, but once the resilience is characterized, it can be used to “measure” the size of a disturbance. With this simple lake-agriculture system, no matter how the new system configuration came about, the difference in the pre- and post-resilience characterizations can be used to measure the impact in terms of resilience.

Hence, the magnitude of a “lake eutrophication event” or a “labor union strike” can be measured in diminished resilience. As always, caution must be used to ensure that homologous systems are being compared.

Simple Desk

A trivial “desk system” with a top, four drawers, one file drawer and four legs might offer valued functions of work surface (sq. ft.), locked storage in drawers (cu. ft.), and filing space (feet of standard file width). Table 16 depicts the resilience characterization for the simple desk.

Table 16
R-characterization: Desk

System (Function)	Productivity	Multiplicity
Top (Level Work Surface)	20	1
Drawer (Storage)	1	4
File Drawer (Filing)	2	1
Productivity	39.00	
Complexity	16	
R-characterization	2.4375	

As alluded above, a desk is manufactured to operate in a given environment (e.g., the level floor of a climate-controlled office building), so it would be ludicrous to expect it to function properly when positioned on the stairs in a stairwell. It therefore makes no sense to suggest it is not resilient to that particular disturbance. Recall that resilience characterization is unconcerned with specific perturbations, so it cannot matter why the desk is deployed in an unlevelled environment. It matters only that its valued function is impacted. Once the desk’s resilience has been properly characterized, positioning the desk precariously on the stairs can be visualized as an alternative system deployment and then resilience can be re-characterized. In this case, the contents of the drawers might be a bit jumbled, but only the 20 sq. ft. of “Level Work Surface” would be lost, resulting in a new resilience characterization of 1.267. The same kind of re-characterization can be

done if, say, somebody attacked the desk with a chain saw and removed a corner, destroying two sq. ft. of work surface and one of the drawers. In this case, the new resilience characterization is: 2.143.

The important thing to understand about this is that quantum resilience changes the equation on perturbations. Instead of attempting to define resilience by corralling an infinite number of perturbations, specific perturbation impacts can be “measured” by changes in resilience characterization. This properly focuses resilience analysis on the system instead of its environment.

Imbricated Theater

When the show must go on, it might be important that a theater is resilient. Imagine a simple theater with some physical offices, dressing rooms, entry foyer for concessions sales, and a stage. This theater will obviously have some management, concessions staff, production staff, and a troupe of actors. Hierarchically, the theater system might be arranged as depicted in Figure 18. The concessions staff and the production staff report to the manager but form separate organizations within the theater.

Note that for this example the concessions staff contains three normal “sellers” and one “special” seller. This is because, in a pinch, this sales attendant can also act. The theater demonstrates redundancy in its degenerate form because even though the sales staff are differently trained and have different “structure” than the actors, one can apparently substitute and provide the acting function if the need arises. This also demonstrates the way degeneracy can lead to imbricated redundancy since the sales attendant reports to a different organization than the actors. In this illustration, the larger

“containing” system (the entire theater organization) is getting involved in providing the function when the specific sub-organization (the acting troupe) cannot deliver it. Such functional redundancy may seem inconsequential until the resilience of the system is characterized.

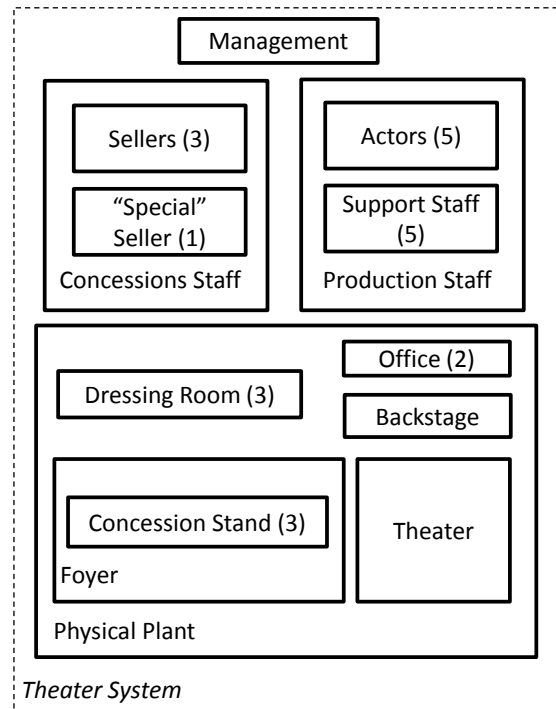


Figure 18. Simple Theater System

For this simple system, assume each actor produces one unit of performing, while each concessions attendant produces one unit of selling. Additionally, the special sales person also contributes her capacity to act. While we may not understand the “magic” behind a person who can sell concessions *and* perform, Table 17 shows that it confers added resilience on the system.

Table 17
R-characterization: Simple Theater

System (Function)	Special Salesperson <i>can Act</i> Productivity (Multiplicity)	Salesperson <i>cannot Act</i> Productivity (Multiplicity)
Generic Sales Persons (Sell)	1 (3)	1 (3)
Special Sales Persons (Perform)	1 (1)	
Special Sales Persons (Sell)	1 (1)	1 (1)
Support Staff (Support)	1 (3)	1 (3)
Actors (Perform)	1 (5)	1 (5)
Productivity	61.00	50
Complexity	40	39
R-characterization	1.5250	1.2821

This example also presents a way to think of human and social capital. Here, if human capital is measured by training (perhaps she is attending acting school), exposure (has seen the play several times), experience (has been called on to fill in before), and daring (no fears about being on stage), the special sales person clearly demonstrates higher human capital than a typical seller who has no interest in expanding her horizons and stepping on stage. It is clearly the human capital of the theater that contributes to its higher resilience in this case.

Electric Power Plant/Grid

After I heard it suggested that the Fukushima Daiichi nuclear power plant “should have been resilient” in the face of the 2012 earthquake and tsunami I realized that resilience is apparently a tricky concept. Though quantum resilience takes a different tack, some scholars argue that resilience might imply the ability of a system to bounce back to typical function after it is perturbed in some manner. Since there is a huge literature supporting that idea, this is most likely what was intended in that casual remark,

though it is anybody's guess what they might have envisioned. Other scholars imply resilience instills the ability to do even more than bounce back, perhaps even allowing for adaptation to be an important aspect of the response (cf. Holling, 1973). Typically, desirable adaptations are not elucidated and, in the context of a submerged and earthquake-toppled nuclear power plant, could scarcely be envisioned. Recently, however, scholars seem to be standardizing on the concept that adaptation is either requisite for resilience or a close cousin (Folke et al., 2010). Quantum resilience proposes a different and more practical approach and the specific industrial example of Fukushima will help to make this clear.

The earthquake and tsunami that destroyed the Fukushima nuclear power plant and forced Japan out of the nuclear power industry provides an excellent example to which the tenets of quantum resilience can be applied. Many would argue that the power plant is the system that should be resilient. But what might “resilience” have looked like in the Fukushima nuclear power plant? Recall that back when Fukushima was designed and built, “fault tolerant” and “ultra-dependable” were the precursors to “resilience” that implied there were designed-in problem management features that included something more sophisticated than simple redundancy. Engineering teams spent significant time in failure modes effects analysis and the FMEAs done for Fukushima (though obviously done over 40 years ago), were very likely as sound as anything going at the time. This resulted in a highly robust system that was augmented with some redundancy and some assorted system performance degradation options. The plant was built to withstand magnitude 9 earthquakes and be failsafe, which it managed to do (Braun, 2011). It was designed to lose power and shutdown gracefully, which it was in the process of doing,

when the over 7 meter tsunami struck (it was designed to withstand 6.5 meter tsunamis). Does this qualify as resilient? It would certainly qualify given the current literature about projecting and planning for disturbances. If so, why did Fukushima fail to survive the tsunami? And if Fukushima's many defenses did not qualify it as resilient, could it have been done better? This provides a hint that resilience is being misinterpreted.

There are two reasons to consider this in the context of quantum resilience. The first is to point out that what is typically considered resilient is really more about robustness (see above). Obviously, making a system robust can only go so far in ensuring its survivability. But in this case, not only did Fukushima operate well in its designed-for environment it also resisted an extreme environment for a time. Second, we must ask if we are thinking about "resilience" at the right scale. Can a standalone power plant really be expected to be resilient? What would the valued functions be in that case? Consider that *Japan's energy infrastructure was very resilient* to the disaster (in fact, so resilient that they quickly phased-out nuclear power generation entirely and are consequently *less* resilient now because of the removal of this functional redundancy). Consider that it was only the areas that were hit by the disaster at Fukushima that were impacted by power outages—and even in that region portable generators were available fairly quickly to provide necessary power. Even if a few weeks or months were required to restore power where infrastructure was damaged, that is an impressive "bounce-back" capability considering the many years required to do the initial configuration. So perhaps resilience must be considered at a larger system level. For Japan, their *energy infrastructure resilience* was instrumented by significant levels of redundancy and broad distribution of generating capability, so it exhibited significant resilience.

While a given power plant cannot be effectively “resilient against” a tsunami (recalling that quantum resilience suggests that phraseology is meaningless), each plant contributes to the resilience of the larger energy production system because of the redundant incremental contributions each makes. In this particular example (and in the energy industry at large), the valued function is electrical power generation, and an ability to provide incremental power to the grid is the quantum of resilience. For this reason, while we may be interested (because of cost) in hardening each plant against a long list of disturbances, it makes no sense to design them for resilience beyond their own local redundancy. Can more be done to make the plants robust? Sure. For example, backup power systems (for pumps and supervisory control systems) can be wired from distant generating plants, ensuring that only physical severance of power cables could dissociate the power from the need—but such solutions are costly and subject to their own failure modes. Further, increasing robustness in this manner greatly increases the complexity of the grid while only marginally contributing to its resilience.

Assume an electric power grid is supported by several power plants offering two valued functions: energy to the grid (MW) and employees to the job market. Table 18 provides a summary resilience characterization of a grid with six power plants with varying outputs and employment. Note that these numbers are notional and resilience characterization is high because none of the real complexity is modeled. Obviously, if one plant is removed (e.g., for maintenance or because of a tsunami, see 5-Plant column), resilience characterization will be lower because of less overall productivity. Table 18 shows a ~15% decrease in resilience for this alternative system deployment. This might be expected given that Plant 1-1 was one of 6 and was one of the smaller ones.

Table 18
R-characterization: Multi-operator Grid

System (Function)	6-Plant Productivity	5-Plant Productivity
Plant 1-1 (Energy)	200	-
Plant 1-1 (Employment)	20	-
Plant 1-2 (Energy)	150	150
Plant 1-2 (Employment)	15	15
Plant 1-3 (Energy)	350	350
Plant 1-3 (Employment)	30	30
Plant 1-4 (Energy)	250	250
Plant 1-4 (Employment)	20	20
Plant 2-1 (Energy)	400	400
Plant 2-1 (Employment)	30	30
Plant 2-2 (Energy)	400	400
Plant 2-2 (Employment)	25	25
Productivity	11340.00	8350.00
Complexity	20	17
R-characterization	567.0000	491.1765

Clearly, removing one plant from the grid shows obvious impact but the redundant plants allow it to not be crippling. Table 19 shows that if the system were further (if trivially) decomposed to have multiple turbine/generators providing the energy for each plant, the same impact (loss of plant 1-1) would only impact resilience by 12%. Such simple examples reinforce the intuition that incremental delivery of valued function from redundant systems contributes to resilience.

Table 19
R-characterization: Multi-operator/Multi-turbine Grid

System (Function)	6-Plant Productivity	5-Plant Productivity
Turbine/Generator 1-1-1 (Energy)	100	-
Turbine/Generator 1-1-2 (Energy)	100	-
Plant 1-1 (Employment)	20	-
Turbine/Generator 1-2-1 (Energy)	50	50
Turbine/Generator 1-2-2 (Energy)	50	50
Turbine/Generator 1-2-3 (Energy)	50	50

System (Function)	6-Plant Productivity	5-Plant Productivity
Plant 1-2 (Employment)	15	15
Turbine/Generator 1-3-1 (Energy)	100	100
Turbine/Generator 1-3-2 (Energy)	100	100
Turbine/Generator 1-3-3 (Energy)	150	150
Plant 1-3 (Employment)	30	30
Turbine/Generator 1-4-1 (Energy)	100	100
Turbine/Generator 1-4-2 (Energy)	150	150
Plant 1-4 (Employment)	20	20
Turbine/Generator 2-1-1 (Energy)	100	100
Turbine/Generator 2-1-2 (Energy)	100	100
Turbine/Generator 2-1-3 (Energy)	100	100
Turbine/Generator 2-1-4 (Energy)	100	100
Plant 2-1 (Employment)	30	30
Turbine/Generator 2-2-1 (Energy)	100	100
Turbine/Generator 2-2-2 (Energy)	100	100
Turbine/Generator 2-2-3 (Energy)	100	100
Turbine/Generator 2-2-4 (Energy)	100	100
Plant 2-2 (Employment)	25	25
Productivity	32340.00	25400.00
Complexity	50	44
R-characterization	646.8000	577.2727

This illustration also demonstrates that a specific power plant’s resilience is attributable to its redundant energy supply capacity (i.e., more turbines and generators), and that it is difficult to otherwise increment the resilience of a specific power plant. In general, it makes sense to increase the *robustness* of an individual power plant (as was done with Fukushima) but to understand energy *resilience* at a broader system level.

Pipe System

Intuition can be tested by considering a simplistic pipe system where the valued function of the system is delivery of some fluid. In this simple case assume one pipe delivers 100 units of liquid from one place to another (pressure, pipe size, and flow rate are all ignored). Table 20 depicts the resilience characterization results (Single Pipe

column). It seems intuitive that adding a redundant pipe will increase the resilience of the delivery system. Note, however, that unless the *volume* of fluid delivered is also increased, adding another pipe simply increases the system complexity (and cost).

Table 20
R-characterization: Simple Pipe System

System (Function)	Single Pipe	Dual Pipe	Shared Dual Pipe
Pipe 1 (Fluid Flow)	100	100	50
Pipe 2 (Fluid Flow)	-	0	50
Productivity	100.00	100.00	200.00
Complexity	2	3	4
R-characterization	50.0	33.33	50.0

Consider two alternative configurations with a redundant pipe, and note that *neither increases* resilience. First, if the backup pipe is added and left *empty* (see Dual Pipe column), it is merely a backup, and provides no function—only complexity and cost. This actually *reduces* resilience. Second, if the “backup” pipe is used to provide half the valued function (i.e., 50 units flowing through each pipe, see Shared Dual Pipe column), the redundancy can be thought of as a “wet” backup. Still, this only matches the resilience of the original configuration because now two systems (with their attendant complexity) are used to perform the same overall function (delivering 100 units) that one system previously provided.

In the two latter cases, the redundant pipe might only be added if it is considered critical infrastructure where limited downtime is permitted. In the “dry” backup case, there may be a delay while the empty pipe takes over for the first (depending on pressure and distance), but it might be deemed a reasonable delay when repair times are considered. Importantly, however, it must not be glibly assumed that resilience from

redundancy will always be cost effective. If the mean time to repair (MTTR) for the failed pipe warrants the additional cost for a hot backup, clearly the third configuration provides the best choice in spite of the fact that its resilience is identical to the single pipe configuration.

Note, however, that multiple valued functions could be considered. If function is distinguished and modeled separately as “fluid flow” and “backup”, the resilience characterization is the same as the original characterization because now there are two valued functions (see Table 21). Alternatively, engineers may decide that modeling flow *capacity* (the capacity to have flow) makes more sense. This allows for both pipes to contribute the same valued function even if not being used (results shown in Table 22).

Table 21
R-characterization: Simple Pipe System (Differentiated Functions)

System (Function)	Dual Pipe
Pipe 1 (Fluid Flow)	100
Pipe 2 (Backup)	100
Productivity	200
Complexity	4
R-characterization	50

Table 22
R-characterization: Simple Pipe System (Combined Function)

System (Function)	Dual Pipe
Pipe 1 (Flow Capacity)	100
Pipe 2 (Flow Capacity)	100
Productivity	400
Complexity	4
R-characterization	100

Importantly, quantum resilience does not force a specific approach, and leaves the domain-specific modeling decisions to the experts involved. Quantum resilience does, however, enforce transparency and calls attention to the consensus required while at the same time enforcing consistency for any models that might be compared among homologous systems. Any of these approaches could be appropriate and each reinforces the intuition of the engineer.

Referent Organizations

Trist (1983) suggests referent organizations provide vital connective tissue that allows organizations in specific domains (e.g., energy or water delivery and regulation) to better perform their functions as the problems they face become increasingly complex. Trist refers to these as “meta-problems” while Rittel and Webber (1973) would call them “wicked.” Referent organizations contribute not by providing delivery of domain-related functions (e.g., they do not produce energy or pump water), but by filling the inter-organizational space by clarifying values and rules, facilitating communication and information flow, tracking emerging trends, and identifying alternative futures.

In terms of resilience, the “value” of referent organizations is buried in the complexity they add to the structure of the system. This makes the trade between complexity and productivity obvious while pointing out that it is sometimes difficult to observe, putting significant pressure on the analysis team to thoroughly understand and document the systems for which they are characterizing resilience. More specifically, in terms of quantum resilience, the referent organization makes the system more complex without providing a specific increment to productivity by exporting a valued function.

Instead, referent organizations are a piece in the background that increases the productivity of other organizations that are actually producing the valued function. While their overall effect is multiplicative, this can only be through productivity that is increased in *another* organization.

Figure 19a imagines two isolated organizations generating their valued functions (f) while Figure 19b imagines the insertion of a referent organization which connects the two organizations in such a way that they can both be more productive ($f + \epsilon$). As shown, it is questionable whether or not the overall industry represented in Figure 19a will be more or less resilient than the one indicated in Figure 19b. This depends on the incremental value of the productivity represented by ϵ .




Figure 19. Referent Organization as Productivity Enhancer

For this simple example, assume two identical organizations each with a supervisor, an assistant, and two departments with 10 workers. Both organizations have the same output in terms of valued function (e.g., 100 units of something). Their baseline complexity allows resilience for the overall “industry” (consisting of two organizations) to be characterized as shown in Table 23 (“Two Orgs” column). Adding a referent organization with five workers that connects the two organizations together (similar to what is shown in Figure 19b) *without* increasing the functional output of the industry results in a lower resilience due to the increased complexity of the industry (“Two Orgs with Ref” column). As discussed however, referent organizations are supposed to *increase* output of the organizations with whom they interface. In the last column of Table 23, productivity has been increased to demonstrate the restoration of the original resilience of the simple system. Note that for this particular referent organization to be an effective addition to the landscape, it must make each of the other organizations more than 50% more productive in delivering their valued function.

Table 23
R-characterization: Two Organizations with Referent Organization

System (Function)	Two Orgs	Two Orgs with Ref	Two Orgs with Ref and Increased Productivity
Organization 1 (Industry Output)	100	100	155
Organization 2 (Industry Output)	100	100	155
Productivity	400.00	400.00	620.00
Complexity	52	79	79
R-characterization	7.6923	5.0633	7.8481

Optionally, a referent organization can be modeled as providing unique valued functions which would contribute to the overall productivity of the system. This would

serve to counteract their complexity in the overall resilience characterization. While not wrong, such an approach would not be recommended since it *assumes* the value of the referent organization instead of forcing that it be proven through actual productivity increases in other organizations. Once again, this stresses the need for consensus to be established among the experts modeling the systems.

Importantly, though it is not done in this simple example, there is significant human capital that can be modeled for organizations (such valuation is demonstrated in the extended example of the Goulburn-Broken catchment below). In this simple example, if human capital were modeled as a valued function, this would allow referent organizations to contribute to overall system productivity without contributing specifically to the “Industry Output” valued function. The point of this specific example is not to be exhaustive but to isolate the important idea that sometimes hidden complexity indirectly contributes to system productivity.

Urban Foraging

Though scholars differ on the degree to which they are willing to acknowledge cultural differences between foraging and agro-urban societies, there is at least one consistent observation: agrarian societies store things; foraging societies (because of their nomadic and transitory existence) do not. Ultimately, it must be acknowledged that agro-urban societies “stock” their local environments and make them target rich for “urban foraging” while foraging societies depend on nature to stock the environment and sometimes must migrate in order to best exploit the much slower processes of nature. Clearly, hunter-gatherers follow the food, but less obviously, urban foragers also follow

the food. For them, however, food is regularly replenished in their cupboards and on their market shelves.

Though the systems are not perfectly homologous, in general, urban food systems are more resilient than hunter-gatherer approaches because of their adoption of storing. Storing operationalizes resilience through redundancy. This is especially true if food accessibility, availability, and variety are considered among the valued functions of a food system. From a natural landscape you can harvest what is in season, when storage is brought to bear, seasonal supply is less controlling. The availability of multiple stocks of goods has allowed significant food resilience in urban populations. Cities provide means for storage of large and redundant quantities of food and goods. Foraging in a supermarket is easier than in the forest due to the density and variety of the goods in a small space (not to mention the ease). Further, because of the redundancy, there is enough food in a supermarket for many people, and should one market be exhausted, others are readily available until resupply happens.

Kelly (1995) makes it clear that foraging as a lifeway is a spectrum and that there can be no generalized hunter-gatherer society which eventually gave way to a generalized agrarian society. He also speaks of the way tribal intermarriage provided a mechanism for access to a larger quantity of foraging and hunting grounds (Kelly, 1995, p. 274). Intermarriage permitted bi-locality to emerge as an effective response to fluctuating environments since safe and uncontested migration becomes a simple solution when there are inadequate food resources in the current locale. Kelly's summary is that such cultural evolutions "decrease long-term variance in returns and reduce risk." In fact, bi-locality

effectively increases the redundancy of available food sources and enhances the resilience of the foraging food system.

It may have started on the savanna with intermarriage and exploitative family relationships, but exploitation of redundant food sources was mastered in the city. Urban humans still participate in a foraging lifeway. Cities are a depot of goods that provide a target-rich environment for foraging. In the developed world, it has been a very long time since humankind has been predominately involved in food production. Even those who work in agro-business seldom “live off the land.” Instead, humans forage in markets that have been stocked by other humans who focus on growing or husbanding. The trade and supply networks that have evolved over time have contributed greatly to the resilience of the human food system.

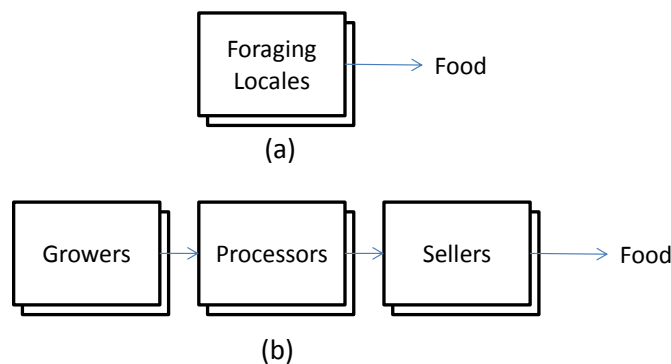


Figure 20. Simplistic Foraging Supply Chains

Figure 20 depicts highly simplified foraging supply networks. Figure 20a represents hunter-gatherer foraging from one or several natural locales while Figure 20b shows a more modern supply chain for food provisioning. In each case a valued function of food provisioning is proposed (as suggested above, a more complete list might include accessibility, availability, variety, nutritional value, reliability, etc.). In each case assume

that 100 units of food are provided. Note that for this simple example only supply is being addressed since demand (or, need) is more properly thought of as a perturbation on the system. Table 24 shows a comparison of some simple functional allocations based on Figure 20a. These include a minimalist environment of one locale (Single Locale), a situation where two locales are accessible, and a situation where two locales are available, each providing 50 units of food.

Table 24
R-characterization: Urban Foraging

System (Function)	Single Locale	Split-Multi-Locale	Multi-Locale
Foraging Locale 1 (Food)	100	50	100
Foraging Locale 2 (Food)	-	50	100
Productivity	100.0	200.0	400.0
Complexity	2	4	4
R-characterization	50.0	50.0	100.0

As shown above with the trivial pipe system example, it can be observed that redundant access to a food supply increases resilience.

Figure 20b depicts a simplified food supply network. Since each of the nodes can be multiplied, care must be taken in reviewing the following tables. Note as well that I am not specifically arguing that modern food supply networks are comparable and homologous with natural foraging locales, but the numbers provided are intended to demonstrate the higher resilience of the urban foraging approach since these simple examples are about intuition.

Table 25 shows a simple progression of complexity increases in urban food supply chains. Column I shows resilience with one grower, one processor, and one seller. Column II assumes a single grower and processor but expands to 10 sellers. This

demonstrates how, with the same food volume and some supply chain complexity increases, the resilience is still greater because of the redundancy in sellers who provide access to the food consumers. Column III expands the network to 10 individual supply chains (10 of each; growers, processors, sellers) but maintains the original volume of food. Note how in this case the complexity “catches up” with the constant productivity values and effectively cancels the gains in resilience. Finally, column IV recognizes that it is not realistic to assume the volume of food will remain the same if the network is expanded like that, so restoring each seller’s volume to 100 (and assuming the growers and processors can support that volume), the dramatic increase in resilience can be seen.

Supporting intuition, this simple example demonstrates increasing resilience with increasing supply redundancy.

Table 25
R-Characterization: Urban Foraging Simple Progression

System (Function)	I Function (Multiplicity)	II Function (Multiplicity)	III Function (Multiplicity)	IV Function (Multiplicity)
Seller (Food)	100 (1)	10 (10)	10 (10)	100 (10)
Productivity	100.00	1000.00	1000.00	10000.00
Complexity	10	46	100	100
R-characterization	10.0000	21.7391	10.0000	100.0000

RESILIENCE IN THE GOULBURN-BROKEN CATCHMENT

Introduction

It has recently become fashionable to be concerned about the resilience of cities and large socio-ecological systems (SESs). Such concern seems warranted since cities and SESs provide livelihoods for most people, and we are interested in understanding how enduring they might be. Further, current approaches to resilience tend to lead us to believe that such assessments are doable, meaningful, and actionable. While quantum resilience enables such assessments, it is unclear that resilience characterization at such scales is actually meaningful—especially if performed from the top-down. This is, in part, due to the broad (and sometimes conflicting) list of perceived valued functions associated with such scales. It is also in part due to the tendency to want to compare the outcomes of specific resilience characterizations and suggest why, say, one city is more resilient than another. Quantum resilience asserts that such comparisons are meaningless since resilience characterizations are only comparable between homologous systems. The section on urban resilience above provides an alternative view and explains why analysis at such scales may not be as meaningful as hoped.

The Goulburn-Broken Catchment (GBC) in Victoria, Australia is a region that provides a suitable example of the power of quantum resilience to characterize the resilience of a socio-ecological system (SES), though it makes no statements as to the value of such assessments beyond their ability to compare alternative deployments of the same system. The GBC has significant complexity including important economic productivity through agriculture production at the expense of ecological deterioration. This, of course, makes it difficult to address with anything less than an interdisciplinary

staff of scientists, but since this venue precludes that approach, the analysis will be representative and remain high-level where required. Still, it will provide a very complete demonstration of quantum resilience and illustrate its dramatically different approach to resilience characterization.

This example is facilitated by a regional assessment performed by the Resilience Alliance. Walker et al. (2009) provides a package that is among the more hopeful “resilience” reports, but not because it is about resilience, *per se*. To be fair, the paper mentions resilience, but it is more of a regional status report that recommends a particular way forward based on particular normative values. Specifically, they have developed an “approach for assessing sustainability” (p. 1) and after prominently featuring the triple bottom line throughout, the authors conclude the region is unsustainable and argue for transformation. The value of the report is found in the small step they take toward a “systems” approach that provides a solid segue for the introduction of quantum resilience. Importantly, the paper provides enough hints to serve as a springboard into a quantitative resilience analysis and it will provide the basis of this example. Where the data is incomplete or notional, it will be estimated and left for future, more rigorous quantification. In all cases, the simplifications made for this illustration remain completely transparent since quantum resilience enforces that in the system modeling approach.

The Resilience Alliance defines resilience very broadly as “a measure of a system’s capacity to cope with shocks and undergo change while retaining essentially the same structure and function” (Walker et al., 2009, p. 1). Obviously, in arguing for transformation, the authors have concluded the system is neither sustainable nor resilient.

As discussed above, “coping with shocks” is a matter for robustness analysis and would introduce infinities into a resilience analysis. Throughout the paper, resilience is conflated with robustness and it is also suggested at one point that “insurance” is “one aspect of resilience” (p. 13). I will return to that topic in the discussion below.

“Undergoing change,” which is presumably a euphemism for adaptation, is simply too soft to consider in any definition which purports to “measure” something. Further, while “retaining structure and function” is appropriate, it is unclear what “essentially” might mean, so that word is similarly unhelpful when measures are intended. Such a broad and malleable definition for resilience leaves them with the ability to recommend “transformation” (which in this case is an externally induced change resulting in a new system) as a possible next step, but as described above, this clearly steps outside the idea of resilience. While their open-ended idea of resilience cannot provide a true characterization of resilience in the GBC, there are worthwhile contributions that can be used in a system analysis that leads to a real resilience characterization.

To highlight differences between quantum resilience and their approach, it is important to observe their plan which they outline as follows:

First we characterize the region as a system by defining the key subsystems, identifying the main issues, drivers, and potential shocks (including changes in drivers). We then assess the capacity of the system to deal with these shocks based on the major benefits currently generated by the region and the biophysical, economic, and social sub-systems that underpin their continued supply. Next, we assess the resilience of the region... (Walker et al., 2009, p. 2).

Their plan includes the following details:

1. characterize the system,
2. identify issues, drivers, and shocks,

3. assess capacity to manage shocks,
4. acknowledge current benefits provided by the region, and
5. assess the resilience.

As mentioned, their emphasis on “shocks” can become an important part of a robustness analysis and can be revisited once the resilience analysis is complete. Only after the system has been characterized can it be determined if robustness features that protect against shocks contribute to, or detract from, resilience. The promising part of the plan is their interest in characterizing the system as it currently is, and acknowledging its benefits (steps 1 and 4). This is an important step toward finally achieving a real resilience analysis. This is where quantum resilience proposes resilience analysis should start: first by identifying the valued functions (“benefits”), and then by codifying the system that delivers them. This will properly scope the analysis and ensure transparency of scope and scale.

Their plan, however, is only a thin veneer on their philosophical position which becomes clear nearly immediately as they list the “issues” that negatively impact the region (cf. Resilience Alliance, 2010, p. 10 for their definition of “issue”). In fact, every “issue” is directly related to the economic productivity of the region which comes at the expense of the pristine environment (Walker et al., 2009, p. 3). Land was cleared and fertilizers were applied so it could be productively farmed. Water is stored to manage seasonal fluctuations in supply and energy is generated and consumed accordingly. While it is not wrong to argue for the negative environmental impacts of each of these, it is not helpful to neglect the fact that the “issues” actually serve to define the very productivity of the current system. Fixing any or all of these “issues” would radically change the

system and have deleterious impacts on the economy. If a true triple bottom line approach is desired, it cannot start with overemphasis on the need to restore nature to what they might estimate to be closer to a pristine reference condition.

Further, in speaking of the GBC (and implicating agriculture development worldwide), Walker et al. (2009) suggest “its resilience has declined since colonization” (p. 6). Unfortunately, they offer no “before” and “after” characterizations of resilience to support their claim, so the claim remains notional. This assessment is clearly based on conservationist values since it is declared that resilience started to decline at the start of colonization in 1830. Their intent is to demonstrate the loss of pristine nature, but in establishing pristine nature as the ideal high mark for resilience, they have already assumed their conclusion. Further, they have skipped over important science that is, ironically, wrapped up in bullets 1 and 4 of their plan (as summarized above). Hence, their idea of resilience is nearly equal to environmental quality. Their equation is simple but not proven and reflects strong normative bias. Further, it makes resilience about what we dream it might have been, or even about what it could be once again (if nature were left to itself), but it completely misses the point that *resilience must be about what is*. When properly characterized, resilience represents the extent to which a system delivers its valued functions. If a triple bottom line is assumed, all three aspects must be included in the resilience characterization. Casual assertions that resilience has declined must be supported. Assuming only the ecological aspects, however, makes it obvious that their only conclusion can be “transformation” that shuts down all productive industry and allows the landscape to revert to pristine nature.

Importantly, while these scholars want to assess the resilience of the entire catchment (to help with future policy initiatives, etc.), their analysis never actually happens at that scale. Instead, it is stated that “human capital” has “high” resilience, that the irrigation sector is “leaky” and has “low resilience,” that the rivers have “increased resilience”, etc. This is an important hint that while scholars suggest they want resilience characterized at the scale of entire basins or catchments they still can only manage to comprehend specifics at smaller scales. This hint should drive us to better analysis at appropriate scales. Again, just because quantum resilience can characterize resilience at any scale does not mean it should.

System Analysis

Figure 21 shows a notional workup of an analysis of the GBC system that arises from reading Walker et al. (2009). In general, the authors segregate the system in accordance with the triple bottom line into biophysical, social, and economic sectors. Note that though properly contained within the “economic” system, the agricultural system straddles the biophysical and the economic systems and is the specific “problem” that forces the analysis. That is, if the catchment had not been developed for agriculture, the authors would agree it would have remained (closer to) pristine and (to them) demonstrated high resilience. Unsurprisingly, the problem devolves into a balance of economic production and environmental protection. Further, by reminding us of the values that allowed for private property rights to evolve and which ultimately led to the current dilemma, Walker et al. illustrate the near identity of the social dimensions with the economic dimensions of the sustainability discourse. That is, at least in the Western

mind, we might agree that “money isn’t everything” and that it “can’t buy happiness”, but we are also willing to agree that having a productive economy is better than not having one. A strong economy has significant social value.

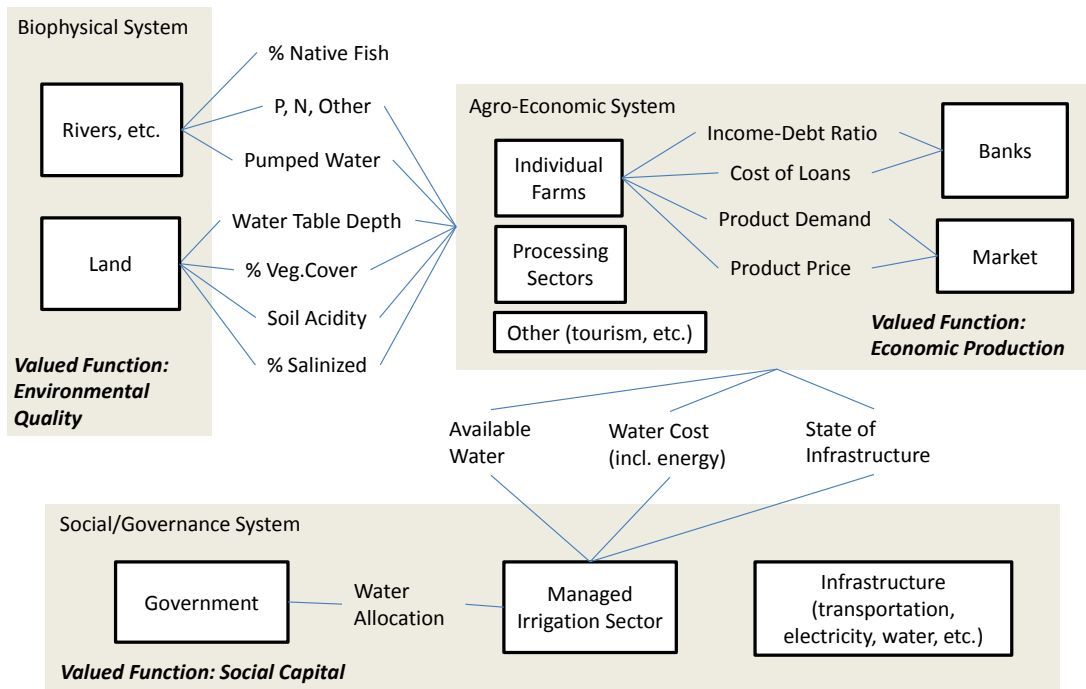


Figure 21. Notional Overview of Actors and Forces in the GBC

The primary concern of Walker et al. (2009) is the significant ecological impact of the agricultural sector on the biophysical system. A complex combination of land clearing, irrigation, and fertilization has resulted in a rising water table that is increasingly saline. Since trees and crops do not grow well in oversaturated soil, groundwater pumping serves to lower the water table, but also results in nutrient rich water discharge into the river systems. Water chemistry alterations have resulted in algal blooms and fish death. As the land deteriorates, it becomes a vicious cycle of fertilizing, irrigating, and ground water pumping in order to keep it productive. Further, water procurement and management has non-negligible cost and requires expensive

infrastructure that must also be maintained. This causes an expected spillover into the economic sector. As crop yields drop due to deteriorating land conditions, farms must borrow money in order to maintain their viability. Of course, financing becomes more difficult to procure with the erosion of projected future crop values. Further, since farms are less financially secure, market fluctuations due to recession can initiate a downward spiral that drives farm owners into receivership.

In the social sector, it is easy to understand the ensuing values conflict between conservationists and agriculturalists. Walker et al. (2009) applies what they call “resilience thinking” to the problem in order to determine and elucidate a way forward. After rehearsing the dilemma in which the catchment finds itself, and providing some notional assessments of the resilience of the systems involved, they determine the only way forward is transformation, effectively stating that the only winner can be the conservationists. As the quantum resilience analysis demonstrates, such a conclusion might be premature.

This current analysis remains high-level but is formalized, hoping to lead to further fidelity improvements over time. It accepts Walker et al.’s general decomposition of the system into agro-economic, biophysical, and social sectors, though such segregation is not strictly required and analysis might better proceed on a more physical basis that starts with a bottom-up approach. Their proposed decomposition can be managed in the model and it allows for reporting on what seems to be the accepted triple bottom line axes. For this example, there is very little system decomposition done. Figure 22 shows the decomposition of the agro-economic sector into three sub-sectors that permit establishment of specific inter-sector relationships. For example, dairy farms

(Figure 22a) produce milk which is processed by dairy processing plants (Figure 22b). The infrastructure systems (Figure 22c) also maintain interfaces with many other systems as dictated by their role in the production system.

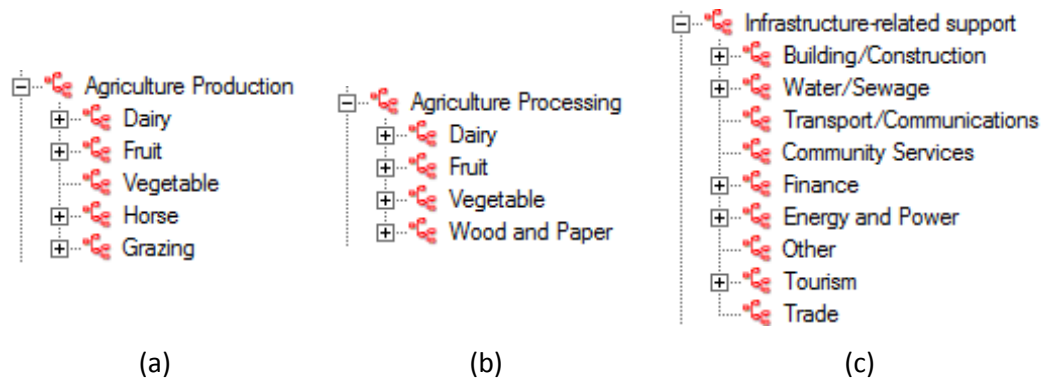


Figure 22. Simple Decomposition of Agro-Economic Sector

Where systems in Figure 22 are not expanded (+ signs), this is intended to hide the identical duplicate systems that represent farms, orchards, ranches, processing plants, businesses, and other organizations that are used in the model. For this analysis, no specific physical decomposition of the infrastructure sectors (energy, water, irrigation, etc.) was performed. Ideally, the full system should be codified down to the conduits, pipes, and pumps to ensure an accurate complexity measure, but for this analysis simple connections were made between the providers and the consumers. This is tantamount to suggesting that, for example, the power provider has a direct connection to each farm. Obviously, this is a simplifying assumption that future fidelity expansions can resolve.

Figure 23 shows the decomposition of the biophysical and social systems. For this simple analysis, the biophysical sector contains only the two primary river catchments and the surrounding dryland acreage (which seemed particularly important to Walker et al.). A more complete analysis might decompose the region into specific farms in order to

emphasize the peculiarities of sub-catchment land uses. This would certainly serve to focus the entry of farmland runoff into the waterways and to emphasize the specific areas that are damaged by acidity and salinization. Since the MDBA assessments of the river valleys is used in the quantification of valued function (more below), this high-level decomposition will suffice for this example. Decomposition of social systems is similarly abridged and representative. Several classes of worker are employed to simulate the differences in human and social capital and only a few organizations are represented.

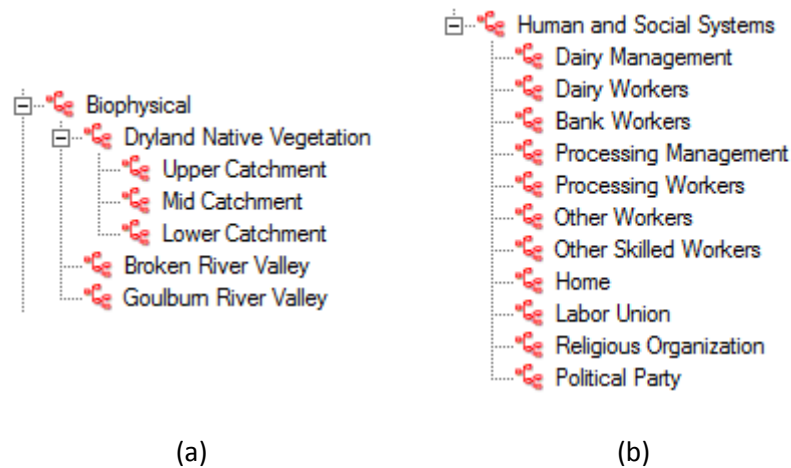


Figure 23. Simple Decomposition of Biophysical and Social Sectors

Though the model is representative only, it can be easily expanded as needed. Even with the simplifications, the model contains hundreds of thousands of relationships. Figure 24 shows a notional idea of the interfaces added to the model. Note the introduction of the “biophysical” to represent impacts and externalities to rivers, etc. For clarity not all the individual interfaces and organizations are shown.

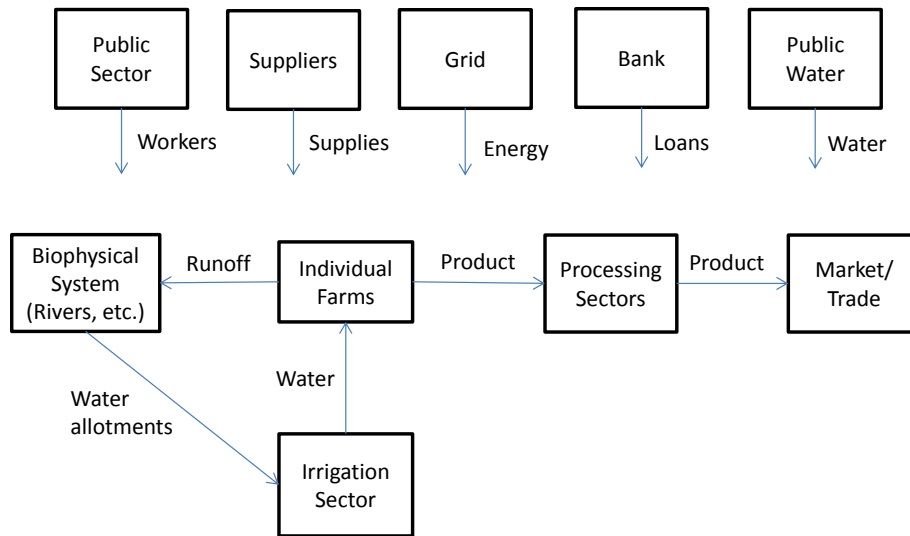


Figure 24. Typical Farm Relationships

As the analyses below will show, the additional interfaces and relationships will dramatically increase the complexity measures and will tend to decrease the resilience characterization since they are implicated in the denominator of the resilience characterization. However, the incredible increase in social and human capital (represented by the many workers) will tend to counter this effect.

Valued Function

As indicated above, Walker et al. approach the GBC value through a fairly common sustainability sieve and focus on the triple bottom line. They specifically refer to the ecological (biophysical), the economic, and the social aspects. From these three “subsystems” (p. 20) they derive GBC value, though as is fairly typical in the literature only the economic aspects are quantified, the rest are indicated only in the sense of possible thresholds that increase risk. A resilience analysis must quantify all valued function and quantum resilience enforces this and makes the outcomes transparent.

Economic value, as always, is easy to measure and is featured prominently in Walker et al., so limited discussion of these measures is required. The economic production figures used in the simple model presented here are merely transcribed as estimated from Walker et al. (p. 7, figure 3). Higher fidelity or granularity can certainly be added to the model, but for an example analysis these will suffice. Walker et al. express interest in farm debt, but this does not seem to be a serious issue in that Victorian debt compares well to other regions. According to the Australian Bureau of Agricultural and Resource Economics and Sciences (ABARES) (Dharma & Dahl, 2013), Victorian dairy farms have average annual income of ~\$150,000 whereas average farm debt is ~\$700,000. Approximately half of farm debt is for land purchase, and, in general, land values increased between 2000 and 2010. Another 10% of debt is working capital to manage cash flow as is expected for low margin concerns. With average equity ratios of greater than 70%, farm debt seemed well in line with expectations. If such dairy-specific figures can be safely extrapolated to general farm debt, it seems reasonable to not consider debt in this analysis, but it can easily be added as necessary.

Walker et al. frequently conflate “resilience” with “environmental quality” (e.g., p. 6, 19-20) without any specific calculation of value. In general, pristine nature (c. 1830) seems to be valued as the reference condition, and this becomes clear with their recommendations for “transformation” away from agricultural production and toward conservation and (perhaps) ecotourism as a way to limit the damage done but still exploit nature for financial gain. They acknowledge that people “will lose from the transformation” (p. 21) and that “strategic new investments in social and human capital” (p. 21) will be required, but they offer no plan as to how such investments will be

financed. Instead, they “propose there is a possible tipping point between market values vs. preferences for non-market, intrinsic, and option values” (p. 20). The assumption seems to be that people will eventually be willing to give up their way of life and standard of living once they realize the environment is in danger. This is difficult to accept and is certainly not something that is actionable from a policy standpoint (i.e., you cannot simply recommend that people close their businesses and relocate).

Despite “clear evidence that transformation is needed” (p. 21), it remains unlikely that the economic and social goals can be met if transformation is achieved. Since the ecological values are in direct conflict with the economic values (as well as *most* currently held social values despite a possible “greening” of the Australian mind), three alternatives can be envisioned: (1) maintain a “current” (exploitative) regime that may result in a future crash, (2) implement a “multiple values” (balanced) regime where a crash might be avoided or put off for an extended period, or (3) revert to an “intact” (pristine) regime where the crash is effectively self-inflicted but with foreknowledge and hopefully some mitigation planning. Walker et al. (p. 9) outline these three regimes specifically for the river channels, but the ideas can easily be extrapolated and applied to the entire catchment. While there are no specifics offered, a “multiple values” regime seems to assume some moderate economic return can be gained while some specific conservationist goals are achieved. Again, the authors do not specifically state this (favoring instead the radical transformation recommendation), but it appears that an interim approach *could* be managed between now and the radical transformation. For example, perhaps economic output could be reduced by 10% per year over the next 20 years while relocating people to other regions that have alternative means of production

and that are either less environmentally fragile or where the environment is already irrevocably damaged. No matter what the plans are (radical transformation or adaptive management in hopes of restoring pristine nature while removing the humans), an honest assessment of the social impact would be required.

In an attempt to take a rational step forward, the catchment's ecological valuation could be modeled in each of the three regimes and its resilience characterized accordingly. Such a valuation would employ the variables of high concern including water table depth, salinity, soil acidity, vegetative coverage, river condition (Walker et al., 2009, figure 6) and project their values accordingly. In an attempt to make the measures as directly usable as possible, the inverse of certain measures may be used, but such decisions would usually be determined by consensus of the analysis team. In any case, the quantum resilience model makes the decisions clear and transparent.

Importantly, the river system's value is obviously extremely high to the agriculture sector. But ecological values are measured differently (conservationist instead of exploitative). Because it is being exploited for agriculture development, the valuation of the catchment's river systems is reflected in the economic production figures. When it comes to the natural values (i.e., those that would exist prior to agricultural production) these are codified in the biophysical sector of the model. They are separated for clarity and to enable them to be demonstrated in isolation. Table 26 lists values recently assessed by the Murray-Darling Basin Authority (MDBA) in the second Sustainable Rivers Assessment (SRA-2) (MDBA, 2012). The numbers shown are scores provided by a team of experts out of a "pristine" reference condition of 100. Assessments for the Broken and Goulburn river valleys are shown with the surrounding catchments. Since these items are

deemed important by Walker et al., valuation for fish, vegetation, and overall hydrology will be used in the biophysical portion of the sample GBC model. Use of the MDBA assessment values allows the simplifying assumption that other experts have already considered the important variables mentioned above (water table, salinity, acidity). Ideally, these values would be assigned to smaller parcels, but this will suffice for an example.

Table 26
Partial List of SRA-2 GBC River Valley Assessments

MDB River Valley	SRA-2 overall rating	Fish	Vegetation	Hydrology
Broken	very poor	7	21	97
Goulburn	very poor	15	46	43
Campaspe (West of GBC)	very poor	20	18	64
Murray Central (North of GBC)	poor	20	100	56
Ovens (East of GBC)	poor	40	48	99

Walker et al. speak of human and social capital (and once, political capital) without clarifying or distinguishing the terms. This allows the terms to be defined for use herein following widely accepted scholarship. Human capital is individual knowledge, skills, and experience combined with a willingness to employ such attributes in contributing value to society (cf. Baron, 2011). A more complete and formal analysis might use a numerical method to measure *aggregate* human capital (e.g., as proposed by Mulligan & Sala-I-Martin, 2000). For this example, human capital will be modeled based on education level and skill sets of employees in given sectors. For example, assume human capital is one point each for being employed, being in management, graduating high school, graduating college, being an entrepreneur or business owner, etc. Obviously,

the list can be expanded as necessary to achieve consensus. Some representative values are used in the sample model.

Social capital is distinguished from human capital since it is not individualistic, but relational. Putnam (2000) suggests it emerges from repeated interactions among individuals engaged in purposeful activity. Walker et al. specifically refer to “equity” so it seems this should figure prominently, though it is difficult to understand what equity might mean in this context. Very likely it is undifferentiated and included as required by the “social” aspects of the sustainability discourse. In a liberal democracy like Australia, equity might mean equal access to the rule of law, police protection, healthcare, etc.—all of which are “standard” in Australia. Though there are likely exceptions and failings (e.g., displacement of aboriginals), social equity must be assumed and hence would be difficult to model. Walker et al. also mention “participation” (presumably in governance) which is easily measured by statistics like membership in political parties and voter turnout. There are likely far more nuanced means if higher fidelity is required. For use herein, and following Fukuyama (2001), social capital is derived from a network of relationships among those with shared values, and the consequent positive participation it engenders in society. Specific items like safety, trust, and voting can be measured and tracked but, as Fukuyama points out, are derivative of the actual relationships which are the important thing. For this example analysis, assume social capital is comprised of one point each for membership in various groups: family, school, political party, business, religious organization, etc.

Note that one could persuasively argue there are many more valued functions of a society. This simple model summarizes all that value in the social capital function. Future

expansion with a team of experts is encouraged if higher fidelity resilience characterization is required. Table 27 lists a summary of the valued functions in each interest area.

Table 27
Valued Functions of the GBC

Ecological (Biophysical)	<ul style="list-style-type: none"> • Water table depth by hectare (meters: 0 m is bad for agriculture whereas 2 m is acceptable) • Dry land vegetative cover by hectare (percent: 20% is bad, 50% is acceptable, 80% is best) • Vegetation (as assessed in MDBA 2012) • Native fish (as assessed in MDBA 2012) • River basin health (as assessed in MDBA 2012) • Wetland health (percent of 113 total deemed healthy per Walker et al.)
Economic	<ul style="list-style-type: none"> • Economic productivity (dollars)
Social	<ul style="list-style-type: none"> • Social capital (points for relationships, participation, etc.) • Human capital (points for education, employment level, etc.)

For this sample analysis no claim about completeness or correctness is made. In the absence of collaborators (and data) I can only be representative and model in broad trends. My only claim is that such modeling is absolutely necessary. Scholars can argue, persuade, and come to consensus on what is permissible and what best reflects real value, but they cannot ignore this important step.

Resilience

This analysis will demonstrate the quantum resilience approach with discussion of a several GBC resilience models. First, a simple illustration of only the economic value of the agricultural output of the region is demonstrated (*Simple Economy*). This is done with a fairly flat decomposition of the economic sectors. No interfaces have been created between sectors, so no relational complexity is involved. This demonstrates the manner in

which incremental delivery of valued function by many systems enhances resilience. Second, representative sector interfaces are created (as in Figure 24) to add relational complexity to the system (*Complex Economy*). Obviously, considerable interaction among sectors is required to generate the economic productivity of the GBC. While this remains representative and high level its impacts on resilience are obvious. The complex economy also includes estimated valuation of the social and biophysical aspects discussed above. This analysis results in what can be considered a baseline resilience characterization of the GBC. Third, in an attempt to demonstrate use of quantum resilience in comparing alternative system deployments, the analysis moderates the economic output and shows restoration of the ecological aspects in accordance with the aforementioned “shared values” regime (*Moderated Economy*). Since this necessarily results in a smaller economy and lower population, it clarifies how valuation of the ecological aspects contributes to the resilience characterization. Finally, and with limited discussion, the *pristine reference condition* model is shown to illustrate the contribution of the economic and social sectors to overall resilience. After these models are presented some concluding remarks summarize the outcomes and suggest whether or not it is appropriate to consider resilience at such scales.

Simple Economy

A simple illustration of only the economic value of the agricultural output of the region serves as a baseline for comparison. Note that this is done with only the (nearly flat) hierarchical decomposition of Figure 22 (i.e., hierarchical complexity is represented, but no interfaces have been created between sectors, so no relational complexity is

involved). This example demonstrates the manner in which incremental delivery of valued function by many systems enhances resilience.

Walker et al. (2009) suggest the top 16 economic sectors shown in Table 28 contribute nearly \$8B to the economy. For the quantum resilience analysis, each pair of columns in Table 21 illustrates a progressively larger number of subsystems delivering the same quantity of valued function (economic production in dollars). For example, Walker et al. suggest dairy farms provide \$0.45B to the GBC economy. In the first pair of columns the sector is not decomposed, so one system (e.g., one dairy farm) delivers all the economic value. In the second pair of columns the sector is decomposed into 45 systems or representative dairy farms with each providing \$0.01B. Finally the third pair of columns suggests there are 450 dairy farms each contributing \$0.001B. In all cases the same total economic value is provided. The same is done for all 16 sectors. Note that none of these decompositions is likely to approach the actual decomposition of the region. The Australian Bureau of Statistics (ABS, 2013) suggests there are over 30,000 businesses involved in agriculture in Victoria, but it has proven difficult to determine the exact number in the GBC. Further, vetting Walker et al.'s estimates has proven difficult. ABS (2013) suggests Victoria had a total agricultural production of over \$11B in 2007-8, and Walker et al. suggest the number for the GBC alone at over \$3B in 2003 (counting farms and processing), though other sources put that as low as \$1.2B), so it is difficult to ascertain the actual values. Still, the numbers can be considered representative and certainly have utility in demonstrating quantum resilience. If an actual resilience analysis were to be done, fully vetted numbers would be used.

Resilience is characterized for each configuration and shown at the bottom of Table 28. Notice particularly how the resilience increases as the more distributed functional deployment is implemented. This satisfies the intuition that more systems that deliver the valued function are better for the resilience of the overall system. This also provides at least a preliminary answer to Walker et al.’s assertion that they must “cope with shocks”. Assuming shocks are individualized to the decomposed subsystems (e.g., specific farms, or specific land areas) the important take away is that since there are many farms, if the “shocks” are regionally focused, this will only marginally impact overall productivity performance of the system.

Note well that this resilience characterization is incomplete since it has no relational complexity. It is intended only as a first step in demonstrating importance of distribution of valued function. It is clear that with the sectors partially decomposed into separate farms, processing plants, etc., redundancy of function contributes to higher resilience. With the addition of the interfaces and relational complexity, we should anticipate lower (and more accurate) resilience characterizations.

Table 28
R-characterization: GBC Agricultural Economic Production

Sector	Sectors not decomposed		Sectors partially decomposed		Sectors further decomposed	
	\$ B	#	\$ B	#	\$ B	#
Dairy Farm	0.45	1	0.01	45	0.001	450
Orchard	0.19	1	0.01	19	0.001	190
Horse Farm	0.13	1	0.01	13	0.001	130
Ranch	0.18	1	0.01	18	0.001	180
Dairy Plant	1.6	1	0.1	16	0.01	160
Fruit Plant	0.5	1	0.1	5	0.01	50
Vegetable Plant	0.22	1	0.01	22	0.001	220
Wood and Paper	0.17	1	0.01	17	0.001	170

	Sectors not decomposed		Sectors partially decomposed		Sectors further decomposed	
Plant						
Companies	0.35	1	0.01	35	0.001	350
Water/Sewage	0.25	1	0.25	1	0.25	1
Transport/ Communications	0.9	1	0.9	1	0.9	1
Community Services	0.95	1	0.95	1	0.95	1
Bank	0.7	1	0.1	7	0.01	70
Other	0.2	1	0.2	1	0.2	1
Tourist Organizations	0.13	1	0.01	13	0.001	130
Trade	0.9	1	0.9	1	0.9	1
Overall Productivity	125.12		1681.3		16461.1	
Complexity	37		434		4214	
R-characterization	3.3816		3.874		3.9063	

Complex Economy

Second, a more complete model of the system is provided. This includes creation of representative economic sector interfaces to add relational complexity to the system. Obviously, considerable interaction among sectors is required to generate the economic productivity of the GBC. While this remains representative and high level its impacts on resilience are obvious. Further, the biophysical system valuation is added as described above. Biophysical interactions are also created but probably well under represented. Finally, social capital and human capital are added with their representative relationships. Like computer networks, human social networks tend to be scale-free (i.e., high degree of interaction with relatively few and little interaction with many). While analysts of the SES should not trivialize the actual connections, for the purposes of this notional model, it is fair to model only a few of the most impactful relationships. The complex economy analysis represents the entire catchment and it is based on the “partially decomposed” system discussed above (i.e., 45 representative farms, etc.).

Since economic productivity is based almost entirely on the ability to irrigate and manage land in the farms, it is vital to tie that somehow to the surface water systems. The Australian Bureau of Infrastructure, Transport and Regional Economics (BITRE) reports Victoria's surface waters provide approximately 2 trillion liters of water per year. Of this, approximately 1.65 trillion liters are used for irrigation (BITRE, 2013, pp. 274, 283). In making our simplifying assumptions, assume that water is distributed equally to the 45 farms we are using in the simplified study. For this we envision a valued function (water supply to farm connections) assigned to the irrigation sector and measured in billions of liters. Obviously, the infrastructure required for this is massive (as are the pumping energy costs and ripple effects on the power generation industry). It is outside the scope of this brief example to decompose the infrastructure to any great extent, but any complete resilience analysis should do so. Since the agriculture sector has been similarly simplified, this analysis will assume that the energy and water sectors each maintain 45 connections, one to each of the farms.

Table 29
R-characterization: GBC Complex Economy

System (Function)	Productivity	Multiplicity
Dairy Farm (Economic output)	0.01	45
Orchard (Economic output)	0.01	19
Horse Farm (Economic output)	0.01	13
Ranch (Economic output)	0.01	18
Dairy Plant (Economic output)	0.1	16
Fruit Plant (Economic output)	0.1	5
Vegetable Plant (Economic output)	0.01	22
Wood and Paper Plant (Economic output)	0.01	17
Companies (Economic output)	0.01	35
Water/Sewage (Economic output)	0.25	1
Transport/Communications (Economic output)	0.9	1
Community Services (Economic output)	0.95	1

System (Function)	Productivity	Multiplicity
Bank (Economic output)	0.1	7
Other (Economic output)	0.2	1
Organizations (Economic output)	0.01	13
Trade (Economic output)	0.9	1
Water Connection (Water Supply)	35	45
Power Connection (Energy Supply)	170	45
Upper Catchment (Vegetative Cover)	0.5	900
Mid Catchment (Vegetative Cover)	0.2	1000
Lower Catchment (Vegetative Cover)	0.02	500
Broken River Valley (Native Fish)	7	1
Broken River Valley (River health)	97	1
Broken River Valley (Vegetation)	21	1
Goulburn River Valley (Native Fish)	15	1
Goulburn River Valley (River health)	43	1
Goulburn River Valley (Vegetation)	46	1
Dairy Management (Human Capital)	2	135
Dairy Management (Social Capital)	2	135
Dairy Workers (Human Capital)	1	4500
Dairy Workers (Social Capital)	1	4500
Bank Workers (Human Capital)	2	200
Bank Workers (Social Capital)	2	200
Processing Management (Human Capital)	2	250
Processing Management (Social Capital)	2	250
Processing Workers (Human Capital)	1	20000
Processing Workers (Social Capital)	1	20000
Other Workers (Human Capital)	1	20000
Other Workers (Social Capital)	1	20000
Other Skilled Workers (Human Capital)	2	5000
Other Skilled Workers (Social Capital)	2	5000
Productivity	5578465164.30	
Complexity	721787	
R-characterization	7728.6861	

Table 29 shows the dramatic increase in productivity (due to the additional valued function) and complexity (due to the additional system elements and relationships). It is clear that modeling the human and social capital of over 50,000 workers greatly increases the valued function output (despite the small numbers assigned). Though the complexity

contributed through their relationships (on average, each human modeled as three relationships) tends to decrease the overall resilience characterization it is clear that measuring human capital in the manner shown has a dramatic impact on the cumulative resilience characterization. This is most clearly evidenced when the valued functions are listed according to their individual contribution to the overall resilience as shown in Table 30. Shown like this, it is clear that human and social capital contribute over 99% of the resilience.

Based on the simplified “partially decomposed” approach, a resilience baseline for the full GBC system has been established.

Table 30
GBC Resilience by Valued Function

Function	R-Contribution
Economic output	0.0023
Water Supply	0.0982
Energy Supply	0.4769
Vegetative Cover	2.1946
Native Fish	0.0001
River health	0.0004
Vegetation	0.0002
Human Capital	3862.96
Social Capital	3862.96

Moderated Economy

Moderating the economy of the GBC necessarily entails significantly decreasing the economic output. For this simple assessment, economic productivity will be roughly halved by removing half the farms and processing infrastructure (all fractions were rounded up). Though tourism was not changed (assuming that industry may even be *helped* by the changes), the total economy is reduced to \$4.45B (57%) of the original

\$7.82B. Obviously, with the decrease in the economy, human workers will be impacted and this will result in the loss of nearly half the human workers—reducing human and social capital. Given the limited production, it is assumed there will be lower impact on the environment allowing some recovery toward the pristine reference condition. The model increases the biophysical assessment by setting the assessments to half-way toward 80% of the pristine reference condition (it is unclear if this is realistic, but rapid recovery of the environment is assumed). For example, if “fish” was assessed at 15 by the SRA-2 (MDBA, 2012), then fish will be set to $15 + (80-15)/2 = 48$ (again, rounding up as necessary). Table 31 lists the results of the moderated economy configuration.

Table 31
R-characterization: GBC Moderated Economy

System (Function)	Productivity	Multiplicity
Dairy Farm (Economic output)	0.01	23
Orchard (Economic output)	0.01	10
Horse Farm (Economic output)	0.01	7
Ranch (Economic output)	0.01	9
Dairy Plant (Economic output)	0.1	8
Fruit Plant (Economic output)	0.1	3
Vegetable Plant (Economic output)	0.01	11
Wood and Paper Plant (Economic output)	0.01	9
Companies (Economic output)	0.01	18
Water Connection (Water Supply)	35	23
Water/Sewage (Economic output)	0.25	1
Transport/Communications (Economic output)	0.5	1
Community Services (Economic output)	0.5	1
Bank (Economic output)	0.1	4
Power Connection (Energy Supply)	170	23
Other (Economic output)	0.2	1
Organizations (Economic output)	0.01	13
Trade (Economic output)	0.5	1
Upper Catchment (Vegetative Cover)	0.65	900
Mid Catchment (Vegetative Cover)	0.5	1000
Lower Catchment (Vegetative Cover)	0.42	500
Broken River Valley (Native Fish)	43	1

System (Function)	Productivity	Multiplicity
Broken River Valley (River health)	97	1
Broken River Valley (Vegetation)	50	1
Goulburn River Valley (Native Fish)	48	1
Goulburn River Valley (River health)	60	1
Goulburn River Valley (Vegetation)	63	1
Dairy Management (Human Capital)	2	70
Dairy Management (Social Capital)	2	70
Dairy Workers (Human Capital)	1	2300
Dairy Workers (Social Capital)	1	2300
Bank Workers (Human Capital)	2	100
Bank Workers (Social Capital)	2	100
Processing Management (Human Capital)	2	125
Processing Management (Social Capital)	2	125
Processing Workers (Human Capital)	1	10000
Processing Workers (Social Capital)	1	10000
Other Workers (Human Capital)	1	10000
Other Workers (Social Capital)	1	10000
Other Skilled Workers (Human Capital)	2	2500
Other Skilled Workers (Social Capital)	2	2500
Productivity	1403016801.00	
Complexity	364056	
R-characterization	3853.8489	

As expected, the result is dramatically lower than the baseline resilience of the complex economy. What is most interesting about the result is that though economic production remained at over half (57%) of the complex economy, and though the biophysical productivity was nearly doubled, *the overall resilience assessment is less than half that of the complex economy* (3854 compared to 7729). Whether or not the numbers are “real”, the formulation is consistently applied in both cases demonstrating two things: (1) the importance of the redundant and incremental delivery of valued function, and (2) the importance of the human capital contributions and the need to gain consensus among experts on how they are quantified.

Pristine Reference Condition

As indicated in Table 27, the vegetative cover is calculated on a per-thousand-hectares basis (hence, the upper catchment with 900,000 hectares, etc.). This may be thought to greatly skew the outcome, but recall that consistency is the only requirement in managing the values. In the pristine reference condition, it is assumed that 80% vegetative cover is regained in the catchment. Further, it is assumed that the MDBA assessment of the river valleys would achieve 100% of reference condition. Table 32 demonstrates the resilience of the pristine catchment. It is clearly far lower than the catchment while under production.

Once again, whether or not the valuations are approved by a consensus among experts, consistency and transparency are maintained in the characterization process. With collaboration among experts, the numbers may be refined and granularity added as needed.

Table 32
R-characterization: GBC Pristine

System (Function)	Productivity	Multiplicity
Upper Catchment (Vegetative Cover)	0.8	900
Mid Catchment (Vegetative Cover)	0.8	1000
Lower Catchment (Vegetative Cover)	0.8	500
Broken River Valley (Native Fish)	100	1
Broken River Valley (River health)	100	1
Broken River Valley (Vegetation)	100	1
Goulburn River Valley (Native Fish)	100	1
Goulburn River Valley (River health)	100	1
Goulburn River Valley (Vegetation)	100	1
Productivity	4609200.00	
Complexity	4808	
R-characterization	958.6522	

Discussion

As expressed in the introduction, it is unlikely that assessing the resilience of region so large has practical utility when done from the top down as illustrated here. Since the outcomes devolve into a conservationist v. agriculturalist argument, and any recommended changes will be only marginally helpful, it is probably more important to focus on the resilience of specific infrastructure (e.g., just the water distribution network or the electrical grid), or on the resilience of the individual farms themselves (e.g., exploring alternative deployments of crop portfolios, etc.). This obviously leads to a much more complete analysis several hierarchical levels down from the catchment itself. Tighter-scoped assessments could lead to actionable results since exploring “what-if” scenarios with alternative system deployments remains tractable at those levels.

As with most systems engineering efforts, smaller models can be merged to create larger (catchment-level) models if desired. If analysts have vetted the smaller models and ensure consistency in the approach, the tools and methods of quantum resilience are clearly able to provide the insights required at all levels. Even in the simplified form presented here, assessments like this simply cannot proceed without the instrumentation offered by model-based system engineering tools. In order to properly execute such an analysis, an entire team would need to be assembled, not because the modeling is hard, but because almost every decision is debatable. As consensus is achieved, the models, the methods of valuation, and the consequent resilience characterizations can be trusted.

Importantly, it is clear that valuation of human and social capital will dominate any resilience characterization in which it is included. Especially at the size of a catchment, the large populations and cumulative impact of the human and social capital

very quickly results in extremely large numbers. Due to this, it is tempting to think this is artificially inflating the human side of the characterization at the expense of the other sectors. Obviously, since consistency is enforced by the quantum resilience approach, whether they are aggregated (with economic and ecological aspects) or segregated to be published separately, the numbers are useful when it is remembered they are all relative. There is no theoretical maximum for resilience and the resilience characterizations are not absolute in any sense. That is, they can only be compared with similarly modeled homologous systems. What this means is that if experts agree on the approach and the valuation methods, it matters very little if the social dimension contributes 99% of the total outcome (as it does here). None of the numbers are lost, they all remain transparent, and resilience is properly characterized. Still, if communicating specific differences or nuances is important, reporting the social and human capital numbers separately might lead to better understanding of results.

Similarly, it is tempting to think resilience of ecosystems would better proceed on an individual basis and not be included with human and economic systems. It is dangerous, however, to separate the outcomes since it might tempt some to argue that pristine nature is indeed the ideal high mark for resilience. This is only so in normative senses and can lead to discounting the economic productivity gained from an ecosystem simply because such gain “uses up” (or, damages) the ecosystem. Still, it might be important to consider each regime in its own right. For example in the analyses above, three different biophysical regimes were modeled, but until the pristine nature regime is shown alone at the end, it is easy to miss the differences in the wash of numbers. Since

the tools of quantum resilience allow the characterizations to be isolated, it is easy to present them together as shown in Table 33.

Table 33
Biophysical Resilience in all three regimes

System (Function)	Multiplicity	Complex Economy Productivity	Moderated Economy Productivity	Pristine Productivity
Upper Catchment (Vegetative Cover)	900	0.5	0.65	0.8
Mid Catchment (Vegetative Cover)	1000	0.2	0.5	0.8
Lower Catchment (Vegetative Cover)	500	0.02	0.42	0.8
Broken River Valley (Native Fish)	1	7	43	100
Broken River Valley (River Health)	1	97	97	100
Broken River Valley (Vegetation)	1	21	50	100
Goulburn River Valley (Native Fish)	1	15	48	100
Goulburn River Valley (River Health)	1	43	60	100
Goulburn River Valley (Vegetation)	1	46	63	100
Productivity		1584458	3108722	4609200
Complexity		4808	4808	4808
R-characterization		329.55	646.57	958.65

Given the presentation in Table 33, it is easy to see how some would assume pristine nature represents the idealized goal of maximum resilience. The numbers are clearly growing as nature returns to its pristine reference condition. Obviously this is due to the normative approach taken in its valuation (which introduces the systemic trend). As mentioned above, however, it is vital to see that *overall resilience of the catchment is much higher when the agricultural production is added*. This is because all the valued

functions are taken into consideration in a proper characterization of resilience. Though it could be provocatively suggested that based on these outcomes (and *contra* Walker et al.) resilience has *improved* since pre-colonial times, it must be remembered that such comparisons are unscientific since they constitute entirely different systems. Only when consensus is achieved about how to balance the normative valuations with the actual measured values can a true resilience characterization be accomplished. Quantum resilience provides this way forward since it ensures that transparency and consistency is maintained in all assessments.

Finally, it was noted earlier that Walker et al. suggest insurance is an aspect of resilience. This is far too casual a remark and demands further comment. If insurance is an aspect of resilience, it is because it can be construed as a “backup plan,” i.e., a form of redundancy. Since redundancy drives resilience this is a tempting equation, but care must be taken in any effort to equate insurance with resilience. First, on the positive side, since insurance is generally a financial arrangement, it provides a uniform medium of conversion between “everything” and money. If insured infrastructure fails, money is paid. If someone dies, money is paid. If a crop fails, money is paid. Such object-to-cash conversions (and the institutions which perform them) may be important in determinations of valued function quantification. As has been exemplified, economic valuation is by far the easiest way to gain consensus on value. So if the tools of the insurance industry can assist, this is a potentially positive approach. Second, since Walker et al. define resilience in terms of disturbances, and since insurance is generally paid after a particular disturbance removes the ability of a system to provide its valued function, insurance payouts suggest the valued function might ultimately be restored as

the money is reinvested. This assumption requires important caveats: (a) the process of insurance payout and rebuilding temporarily leave a system with less function and only restore the provision of function at a future time, and (b) it allows for the creation of new (different) systems which provide different functions. Obviously in the latter case, this is not resilience. In both cases, loss of function implies a change in the system so the system's resilience must be re-characterized. Once the rebuilding (or whatever) is accomplished, a similar re-characterization must occur. Third, insurance is generally sourced from *outside* the system scope (the same is true for bank loans, but, in general, loans are paid back, whereas insurance payouts are *always* losses for the insurance company). What this means is that either the system scale has accidentally been broadened leaving the resilience assessment in question, or analysts are expressly allowing for "resilience" to be bolstered by external or containing systems. This is important. When resilience is defined with "recovery" or "restoration" in mind, the latter is generally the case, and this calls into question the validity of the assessment. When this happens the only appropriate thing to do is admit to the system expansion and admit that the analysis is really not about the specific system but includes external contributors.

In summary, when performing resilience analysis, especially in a normative sustainability framework as attempted by Walker et al., it is vital to consider all aspects of the triple bottom line together, but it might be more effective to do it in smaller chunks. If an overall catchment assessment is required, a bottom up approach can be used to assemble larger analyses once they have been vetted. In all cases, teams of experts must achieve consensus on valuation and decomposition of the systems and their relationships.

Future Experiments

While Walker et al. have suggested that the region is a likely candidate for “transformation”, they leave the plan undeveloped. The intent seems to be to allow the region to revert to Nature and become a tourist mecca. This is, of course, completely in accordance with their conservationist hopes, but does not deal honestly with the significant economic and human impact that would result. Sure, values are changing and even those living there are willing to admit to the damage done, but certainly they would rather continue their livelihood than be forced to relocate and learn a new trade (euphemistically referred to as “development of human capital”). Further, expecting that the region can be turned into a tourist haven is legitimate, but to expect the economic value of that particular sector to equal the agriculture production is a false hope.

Interestingly, there may be an opportunity to experiment with the idea of opportunistic exploitation of redundancy. It has been reported that Eastern Australia will have a surplus of electrical power generation capacity in the next decade to 2025 (AEMO, 2014). Most are responding to this situation by applauding their ability to lower carbon emissions by shutting down power plants. However, since the impact of that is so small (consider the Chinese directly to the North installing over 30 coal-burning power plants per year), it might be better to explore alternative uses of the power. Instead of idling the power plants and the operations staff, perhaps this would be a good time for Australia to consider subsidizing some level of power generation that could be used for pumping the increasingly saline groundwater from the Murray-Darling basin into the Pacific Ocean (or perhaps something a bit less radical but equally productive). This

would serve the dual purpose of preserving jobs in the power industry as well as preserving the economy of an agriculture producing region that is failing due to rising water tables. Southeastern Australia might benefit from a national contribution to its power needs during a period when they need to do additional groundwater pumping. Making the scale of the system a national one makes it possible to sustain that \$20B annual economy in grazing, fruit production and processing that the basin provides.

The implication of the report is that AEMO future demand projections are being attenuated by rooftop PV uptake. Unfortunately, this is not well documented, and a 24% growth in that industry may not say much if the numbers are small. It also seems their demand model would be much more sophisticated than what is shown in their data files (given the intermittencies of solar power, seasonal change in demand, etc.). That they allude to future developments in storage technology is good, but they freely admit the technology is not yet available. If, however, the reduction in demand comes from closing of a high demand industry (to which they allude) we are free to assume it has something to do with agriculture processing in that area. If closing a high demand industry is a leading indicator of a highly depressed economy, it is important to understand. If it is caused by crop failures due to water table and salinity issues it may be something that can be rectified.

Such a change in approach demonstrates the human ability to plan and alter futures, but still, is not an example of adaptation until after the proposal is implemented and proven to be a positive step. Otherwise, from a resilience perspective, all we can say is there remains some extra generation capacity that can be redundantly applied as needed.

RESILIENCE IN THE PHOENIX FUEL SUPPLY CHAIN

Introduction

Resilience analysis of water distribution networks (WDN) has not been distracted by the metaphors and analogies proposed by thought leaders in the resilience literature. Instead, engineers have focused on developing specific resilience indices that can be applied to real networks (Todini, 2000; Prasad & Park, 2004; Jayaram & Srinivasan, 2008). While this is a hopeful step, unfortunately resilience is conflated with efficiency as these so-called resilience metrics are generally measured as surplus internal power in the network divided by the maximum power that could be dissipated internally after satisfying the delivery constraints. While understandable and perhaps an effective metric for this discipline, it is not specifically resilience.

Fortunately, and for several reasons, it is still a largely productive path. First, they are focusing on the performance of the network (i.e., as measured in pressure head) which is akin to finding the valued function and calculating productivity for the system. Second, they are (inadvertently) introducing structural complexity ideas into the calculations by considering pipe dimensions, number of connections, performance of pumps, etc. This is important because it shows that these engineers recognize that the structure of the system is an important contributor to resilience characterization. Third, though this tends to conflate resilience with performance during failure, extensive Monte Carlo or Genetic Algorithm approaches introduce failure modes, and then with each failure mode, the index is recalculated for the system. This is effectively an exploration of alternative system deployments (e.g., with this or that pipe missing). Usually, cost estimates are developed for non-inferior solutions and multi-criteria analysis is done to

provide recommendations that target a “sweet spot” which balances cost and “resilience” for proposed WDNs. Again, while this might be exactly what is needed for their discipline, it is not specifically resilience that is being assessed.

In a review of three available WDN “resilience” indices, Baños et al. (2011) suggest:

It can be concluded that, as none of the resilience indexes consider where over-demand is applied, but rather the global excess of pressure in the network, they do not accurately determine the capability of the network to provide adequate supply under demand uncertainty. Therefore, it is suggested that resilience indexes consider the topology of the network in order to determine its critical points, where over-demand could make solutions unfeasible, i.e. where the head pressure is lower than the required pressure (p. 2365).

Importantly, they recommend consideration of the topology of the network and a focus on critical points. This suggests several important amplifications are required. First, as quantum resilience asserts, these scholars are beginning to understand that modeling the structure and connectedness (topology) of the system is vital for any resilience characterization. Second, they allude to an important evolution of thought that suggests pressure head might not be the only valued function they must track. The suggestion that “over-demand” may impact the “capability of the network to provide adequate supply” suggests that an additional service level agreement might be emerging (i.e., volume or flow at terminals, rather than just pressure in pipes). This observation leads to, third, the realization that for WDNs, resilience may not be so important a concept. Instead, while resilience can certainly be characterized, for water distribution it could well be that *robustness* is a far more vital concept. If “critical points” and “over-demand” can make solutions “unfeasible”, the implication is that there is more work to do on ensuring

appropriate supply. In essence, it may well be that WDNs should be considered more like appliances where mean time between failures (MTBF) and mean time to repair (MTTR) drive deployment decisions far more than resilience. If MTBF is long, and MTTR is short, it is likely more cost effective to have a repair staff on hand than it is to design for resilience. While these are not yet factored into the current WDN “resilience” metrics, they should be.

System Analysis

This is significant for this example because a large part of the Phoenix petroleum supply chain is a special case of a water (fluid) distribution network. Interestingly, though it is very long, from the standpoint of a fluid distribution engineer, the system is very simple because it has effectively no looping and very low complexity. Since the ultimate goal is fuel availability to consumers, however, the system involves more than pipes. Fuel availability can be modeled in the final delivery of fuel to the filling stations and consequent availability to consumers. This involves long-haul pipelines from Los Angeles and El Paso as shown in Figure 25 and Figure 26 (KinderMorgan, 2015). Once the fuel has arrived in the Phoenix terminal it is shipped via truck to nearly 1000 filling stations in the Phoenix metropolitan area.

Note that capricious consumer demand constitutes a perturbation on the system, not part of the supply system. This demand can be projected with stochastic demand models if necessary (the telecommunications industry calls these “traffic models”), but they will not play a role in this resilience analysis. In fact, such demand is quite stable when considered over the period of time for which consumers value the function of the Phoenix fuel supply chain. Importantly (and as presaged by the Baños discussion above),

we have only passing interest in the pressure head of the fluid supply system. While this is a vitally important engineering design aspect, at the system level, where the valued function can be appreciated, it can be ignored. Focus here will be on available volume.



Figure 25. Phoenix Fuel Supply - West Line



Figure 26. Phoenix Fuel Supply - East Line

Selection of the “headwaters” of the Phoenix fuel supply network sets the scope of the system, but is essentially an arbitrary decision. Selecting the source as back under the sands of Saudi Arabia (and Canada, and Mexico, etc.) would implicate a significant number of infrastructure systems, dramatically increasing analysis time and model complexity. Selecting the crude deliveries at refineries in California and Texas as the source implicates a smaller number of systems, but is still quite involved. Note that crude harvested in Saudi Arabia can be sent anywhere and crude arriving in Los Angeles or El

Paso can be refined and sent (nearly) anywhere as well, so these are not specifically appropriate choices for this system. For this analysis, the pumping stations in Los Angeles and El Paso have been selected as the source since these stations are ultimately what start the fuel on the path to Phoenix. This selection has the benefit of focusing the problem on the specific fuel and transport infrastructure that impacts the final goal of delivery to filling stations and availability to consumers. Refined petroleum sent from Los Angeles to Phoenix (via Yuma) over a specific 20-inch pipe can *only* arrive in Phoenix, so this is the best definition of the Phoenix fuel supply network source. Similar statements can be made about the El Paso source.

Note, however, that it is not wrong to select a wider scope. The Watson pumping station's resilience is likely very important to someone, and it is certainly enhanced by a variety of supply lines that could easily be modeled. Similarly, the Colton terminal can have its resilience characterized and this would likely be very interesting to the operators of the California-Nevada petroleum network. These could certainly be separately analyzed and then included in a broader system model once vetted and understood. As will be shown here, however, the Phoenix fuel supply network is itself comprised of several other important subsystems which can be seen to contribute to the resilience of the overall system, but should also be considered in isolation. These details will be highlighted in this simple example.

Figure 27 provides a simple overview of the system as modeled for this analysis. Five terminals (Watson, Colton, El Paso, Tucson, and Phoenix) are modeled with their storage (though the storage is probably not entirely for Phoenix). For this analysis, no further decomposition has been done though it can certainly be added if details become

available. Each segment of pipe is also modeled with the specifics of fuel flow as allocated to Phoenix. Note that ethanol is received from regional as well as Midwest suppliers for use in the clean-burning gasoline blending, but this is not modeled for this simple example. There are many considerations that would result in higher fidelity.

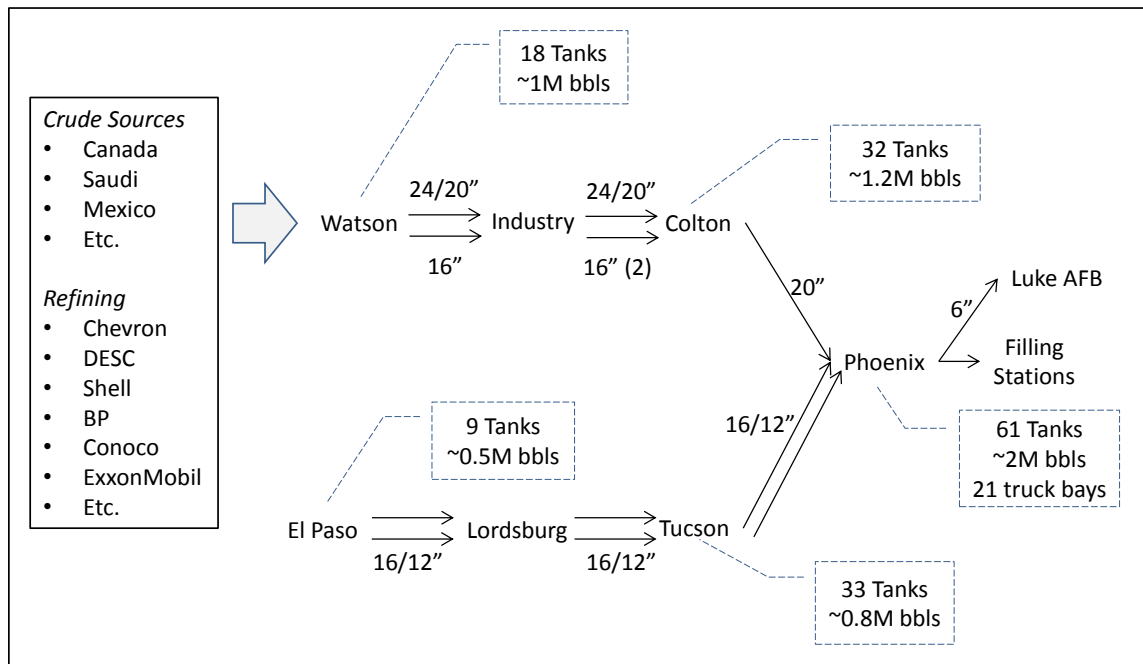


Figure 27. Phoenix Fuel Supply System Overview

Valued Function

For this model three valued functions were selected: pipeline transport and commercial distribution (both in thousands of gallons per day), and tank storage (in thousands of gallons). Applying what was learned in the trivial pipe example discussed above, pipeline transport “capacity” is modeled. The single 20” West line has a capacity of 100,000 barrels (4.2M gallons) per day. After the August 2003 failure that hobbled the Phoenix fuel supply for two weeks, the 60,000 barrels-per-day East line was expanded to support a capacity of 200,000 barrels per day in two redundant pipes, commencing

operations in 2007. Each of the pipes in the redundant East line has 100,000 barrel per day capacity for a total of 8.4M gpd. Note that it is apparent that not all the capacity is used. Though likely out of date, *azenergy.gov* (2004) suggests that approximately 2.8M gpd (60%) arrive from the West line, while 1.8M gpd (40%) arrive from the East line.

Commercial distribution is modeled with a minimally sized fleet (100) of 8,700 gallon trucks that are filled at the Phoenix terminal loading rack (21 truck bays) and deliver fuel to the approximately 1000 filling stations in the Phoenix area. Further, this fuel is modeled as gallons per day available to Phoenix fuel users whose consumption is approximately 4.6M gpd. Given this consumption, 100 trucks can supply the filling stations with 5 trips on a given work day. Because of some of the distances involved, this is very likely an oversimplification that should be revisited for fidelity enhancements.

Since there must remain a mass-flow balance, storage tanks are used to manage fluctuations in demand that might occur over periods over which scheduled inbound fuel flow from the long haul pipelines cannot be responsive. Since it takes approximately a week for fuel to arrive from Watson and El Paso, demand fluctuations on the scale of days can be managed by the 3-5 day storage capacity provided at the Phoenix terminal. This makes storage another important valued function.

Resilience

Based on the system analysis discussed above (with the caveat that complexity has not been fully modeled) and the quantification of valued function, the resilience of the Phoenix fuel supply network is depicted in Table 34.

Table 34
R-characterization: Phoenix Fuel Supply

System (Function)	Productivity	Multiplicity
Long-haul Transport		
Watson Tank (fuel storage)	2333	18
Watson to Colton (fuel transport)	4200	1
Colton Tank (fuel storage)	1575	32
Colton to Phoenix (fuel transport)	4200	1
El Paso Tank (fuel storage)	2333	9
EP-Tucson Line 1 (fuel transport)	4200	1
EP-Tucson Line 2 (fuel transport)	4200	1
Tucson Tank (fuel storage)	1000	33
Tucson-Phx Line 1 (fuel transport)	4200	1
Tucson-Phx Line 2 (fuel transport)	4200	1
Phx Tank (fuel storage)	1350	61
Local Distribution		
Filling Station Tank (fuel storage)	10	1000
Filling Station (fuel distribution)	5	1000
Truck (fuel distribution)	43.5	100
Productivity	285704573.00	
Complexity	4005018	
R-characterization	71.3367	

Though this provides a resilience baseline for the overall system, because of the number of subsystems (multiplicity) involved, it should be immediately obvious that there are really *two* systems here. That is, even at this simplified modeling level, long-haul transport requires roughly two orders of magnitude *fewer* systems than local distribution. While higher fidelity modeling would alter this somewhat, it is unlikely to result in dramatically different ratios in outcomes. There are simply many more systems involved in the local distribution network due to the many trucks and filling stations. For this reason, in order to glean the most information from the resilience characterization, it is probably more appropriate to model these as separate and non-comparable systems, than it is to include them together.

This is illustrated in

Table 35 and Table 36 which highlight the dramatic differences in complexity even when no real complexity has been emphasized (for example, no attempt has been made to consider any of the human dimensions or the variety of organizations and contracts that are involved in Phoenix fuel distribution).

Table 35
R-characterization: Phoenix Local Distribution

System (Function)	Productivity	Multiplicity
Filling Station Tank (fuel storage)	10	1000
Filling Station (fuel distribution)	5	1000
Truck (fuel distribution)	43.5	100
Productivity	20285000.00	
Complexity	4004403	
R-characterization	5.0657	

Table 36
R-characterization: Phoenix Long-haul Transport

System (Function)	Productivity	Multiplicity
Watson Tank (fuel storage)	2333	18
Watson to Colton (fuel transport)	4200	1
Colton Tank (fuel storage)	1575	32
Colton to Phoenix (fuel transport)	4200	1
El Paso Tank (fuel storage)	2333	9
EP-Tucson Line 1 (fuel transport)	4200	1
EP-Tucson Line 2 (fuel transport)	4200	1
Tucson Tank (fuel storage)	1000	33
Tucson-Phx Line 1 (fuel transport)	4200	1
Tucson-Phx Line 2 (fuel transport)	4200	1
Productivity	13619172.00	
Complexity	339	
R-characterization	40174.5487	

Discussion

Obviously, if a more complete model of the long haul pipeline were developed the resilience characterization would have higher fidelity and become more accurate. The same can be said about the model of the local distribution network. Even simplistic models, however, allow some important observations to be made. As shown above, the dramatic differences in the decomposed system count points to a need to segregate these models and suggests that the Phoenix fuel supply network is probably better thought of as (at least) two separate systems. The first system—worthy of analysis but not particularly interesting from a resilience standpoint—would follow fuel delivery from Los Angeles and El Paso to the Phoenix terminal where it would end with fuel storage. The second system, local distribution of fuel within Phoenix, should be understood to constitute its own system and is far more interesting from a resilience standpoint.

For the long haul pipelines, it is obvious that additional redundancy could contribute to resilience, but only at significant cost. This was adequately demonstrated in the 2007 upgrades to the East line which were precipitated by the 2003 pipeline rupture that impacted fuel availability in Phoenix for several weeks (Evans, 2003). Though redundancy was added, the important upgrades consisted of new pipe and weld technologies which did more to increase the robustness of the system than add additional capacity. Further redundancy is unlikely to be cost effective and focus is better placed on maintenance and responsiveness in the event of failures.

Note well that the local distribution system (from Phoenix terminal to filling stations) is already highly redundant (and hence, resilient). In fact, it is this redundancy

which supplied adaptive capacity allowing trucks to be routed on roundtrips to Tucson in order to fill-in the gaps left by the failed East line during the 2003 crisis.

It is tempting to consider what additional redundant capacity in the West line might contribute to the Phoenix fuel supply. Obviously this could easily be modeled and resilience characterization performed, but it is perhaps more interesting to consider such redundancy as adaptive capacity. That is, an additional (redundant) pipeline between Phoenix and West Coast could certainly provide petroleum to Phoenix, but since it is not specifically required at this time, it could be exploited for other purposes. For example, Intel's Phoenix area manufacturing plants generate two million gallons per day of wastewater that is treated in a local reverse osmosis desalination plant (Hackley, 2013). While 60% of the treated water is returned to the underground aquifer, there remains a large quantity of brine that is pumped to evaporation ponds in South Chandler where there is significant fly infestation. Though a complete analysis would be required, there is no specific reason the brine could not be pumped from Phoenix to Los Angeles using the redundant pipeline. With some additional infrastructure, the brine could then be distributed into the Pacific Ocean. This approach (or some other) allows the redundant pipeline to be exploited for alternative uses and opens the door to further experimentation with adaptive capacity.

CONCLUSION

Quantum resilience pragmatically focuses on operationalizing resilience. To this end, a rigorous and enforceable definition has been provided which clarifies and positions resilience in the literature. Further, a system analysis approach has been demonstrated to permit characterization of resilience in accordance with valued function delivery and the complexity of the delivery system. This frees resilience analysis from ambiguous analogies and metaphors and permits even normative values to be acknowledged and quantified while enforcing transparency and consistency. Importantly, this demonstrates that resilience is not an end in itself, but simply one contributor to a more complete understanding of a system.

Quantum resilience analysis both reinforces and challenges human intuition. It highlights the specific mechanisms whereby resilience can be incrementally increased in important systems and makes it clear that protectionist features frequently do not increase resilience but instead contribute to robustness. Formal characterization permits designers and managers to compare alternative or extant system deployments and develop trade spaces that contribute to decision-making processes. Since it works at all scales, quantum resilience is helpful in determining exactly where resilience can be bolstered and where it makes no sense to consider it. This permits understanding of the trade between system productivity and complexity, and reminds designers that increasing system robustness is frequently vital, but will not necessarily augment system resilience.

Since quantum resilience frees analysts from the crippling infinities of perturbation-oriented definitions of resilience, it removes the mystery from the concept and clarifies its utility in forward engineering and management initiatives. Quantification

of function forces analysts to define exactly what they value about their systems instead of permitting a generalized sense that reflects their hope that the system persist. Focus on actual system structures that deliver valued function forces rigor in system definition and transparency in definition of actual system scope. Further, quantum resilience *reverses* the equation on perturbations by first removing them from the definition and then by providing a mechanism by which the magnitude of specific perturbations can be measured. That is, once system resilience is properly characterized, the impacts of alternative system deployments can be quantified whether such alternatives are designed, evolved, or occur as the outcome of a disturbance.

In addition to making it clear that resilience is not measured on an absolute scale and cannot be quantified as an absolute number, quantum resilience demonstrates that even the relative resilience characterizations can only be compared between homologous systems. It is perhaps obvious that it makes no sense to compare the resilience of a bridge to the resilience of the Internet, but it is now also clear that comparable systems will offer the same valued functions and be modeled at fidelities for which it makes sense to compare them to other similar systems.

As a novel and complete theory, quantum resilience provides many advantages beyond clarifying the discourse and providing a way to characterize resilience for all systems. It can make predictions about how to increase resilience in given systems and guide research that leads to discovery of alternative deployment options for systems. It can also project approaches by which adaptive capacity can be incorporated through redundancy. Further, by specifically addressing system complexity it demonstrates how a more complete theory of complexity can be incorporated if one is ever created. Defining

apparent complexity in the denominator leaves the door open to characterizations that might include complexity that is *not* apparent.

Quantum resilience not only provides the required characterizations that permit resilience to be operationalized, but also stages problems in a way that forces acknowledgement of what is actually valued, how it must be measured, what its true scale and extent are, and what the actual contributors to the valued functions are. This serves to clarify the scales at which resilience can be successfully pursued and demonstrates that frequently the analysis scope must be tightened to provide meaningful results.

Finally, quantum resilience is deployed openly and transparently so it can be easily integrated into commercially available model-based system engineering tools. This is a reminder that resilience is not a mysterious idea that must be addressed in vague terminology with custom tools and partial solutions. Instead, resilience characterization flows out of solid systems engineering practice and is a concept that is available and useful to all.

REFERENCES

- AEMO (2014) Electricity Statement of Opportunities for the National Electricity Market. Australian Energy Market Operator Ltd. Available online: www.aemo.com.au.
- Ahl, V. and Allen, T.F.H. (1996). *Hierarchy Theory: A vision, vocabulary, and epistemology*. New York: Columbia University Press.
- Allenby, B. (2012). *The Theory and Practice of Sustainable Engineering*. Upper Saddle River, NJ: Pearson/Prentice Hall.
- Anderies, J. M., Folke, C., Walker, B. and Ostrom, E. (2013). Aligning key concepts for global change policy: robustness, resilience, and sustainability. *Ecology and Society* 18(2): 8.
- Ashby, W.R. (1956). *An introduction to cybernetics*. New York: Wiley.
- Australian Bureau of Statistics (ABS) (2013). Year Book Australia, 2009–10. Retrieved February, 2015 from <http://www.abs.gov.au/AUSSTATS/abs@.nsf/Lookup/756B638E3405C430CA25773700169CB8?opendocument>
- azenergy.gov (2004). Arizona Motor Fuel Supply & Distribution. Online: azenergy.gov/doclib/energy/az_motor_fuel_and_supply_distribution.pdf
- Baños, R., Reza, J., Martínez, J., Gil, C., and Márquez, A. (2011). Resilience Indexes for Water Distribution Network Design: A Performance Analysis Under Demand Uncertainty. *Water Resources Management*, 25(10) 2351-2366.
- Baron, A. (2011) Measuring human capital. *Strategic HR Review*, 10(2) 30-35
- Baskerville, G. (1995) The Forestry Problem. In Gunderson, L., Holling, C.S. and Light, S.S. (Eds.) (1995). *Barriers and bridges to the renewal of ecosystems and institutions*. New York: Columbia University Press.
- Biggs, R., Schluter, M., Biggs, D., Bohensky, E.L., et al. (2012) Toward Principles for Enhancing the Resilience of Ecosystem Services. *Annual Review of Environment and Resources*, 37:421-448
- Borner, K., Sanyal, S., and Vespignani, A. (2007) Network Science. In Cronin, B. (Ed). *Annual Review of Information Science & Technology*, 41. Medford, NJ: Information Today, Inc./American Society for Information Science and Technology, chapter 12, pp. 537-607.

- Braun, M. (2011). *The Fukushima Daiichi Incident*. Areva. Available on-line at: www.ecjones.org/JAPAN_Nuclear_Incident.pdf or similar content at: www.physics.harvard.edu/~wilson/AREVA_Fukushima.ppt
- Bureau of Infrastructure, Transport and Regional Economics (BITRE). (2013). *Yearbook 2013: Australian infrastructure statistics, Statistical Report*, Canberra ACT. Available at: http://www.bitre.gov.au/publications/2013/files/infra1886_r_bitre_infrastructure_yearbook_0813_web.pdf
- Campbell, M.K. and Farrell, S.O. (2003). *Biochemistry* (4th Edition). New York: Thomson Learning, Inc.
- Carpenter, S., Walker, B., Anderies, J., & Abel, N. (2001). From metaphor to measurement: Resilience of what to what? *Ecosystems*, 4(8):765–781.
- Cascio, J. (2009). Resilience. *Foreign Policy*, 172: 92.
- Champe, P.C. and Harvey, R.A. (1994). *Lippincott's Illustrated Reviews: Biochemistry* (2nd Edition). Philadelphia, PA: Lippincott Williams & Wilkins.
- Dharma, S. and Dahl, A. (2013). *Australian dairy: financial performance of dairy producing farms, 2010–11 to 2012–13*. ABARES research report 13.9, Canberra. Online: daff.gov.au/abares/publications.
- Edelman, G.M. and Gally, J.A. (2001). Degeneracy and complexity in biological systems. *Proceedings of the National Academy of Sciences USA*, 98(24) 13763-13768.
- Erdi, P. (2008). *Complexity Explained*. Berlin, Heidelberg: Springer-Verlag.
- Evans, B. (2003). Kinder Morgan Says Arizona Products Pipeline May Restart at Weekend. *Platt's Oilgram News*, 81(160) p. 1.
- Farnsworth, K.D., Lyashevskaya, O. and Fung, T. (2012). Functional complexity: The source of value in biodiversity. *Ecological Complexity*, 11: 46-52.
- Fiering, M.B. and Holling, C.S., (1974). Management and standards for perturbed ecosystems. *Agro-Ecosystems*, 1: 301-321.
- Fiksel, J. (2003). Designing resilient, sustainable systems. *Environmental Science and Technology*, 37(23):5330–5339.
- Fiksel, J. (2006). Sustainability and resilience: toward a systems approach. *Sustainability: Science, Practice, & Policy*, 2(2): 14-21.

- Flood, R.L. (1988). Situational complexity, systems modeling and methodology. *Transactions of the Institute of Measurement & Control*, 10(3) 122-129.
- Folke, C., Carpenter, S. R., Walker, B., Scheffer, M., Chapin, T. and Rockström, J. (2010). Resilience thinking: integrating resilience, adaptability and transformability. *Ecology and Society*, 15(4): 20. Online: www.ecologyandsociety.org/vol15/iss4/art20/
- Folke, C., Carpenter, S., Elmquist, T., Gunderson, L., Holling, C. and Walker, B. (2002). Resilience and sustainable development: building adaptive capacity in a world of transformations. *Ambio*, 31(5):437–440.
- Fukuyama, F. (2001). Social Capital, Civil Society and Development. *Third World Quarterly*, 22(1) 7-20.
- Gell-Mann, M. and Lloyd, S. (1996). Information measures, effective complexity, and total information. *Complexity*, 2(1) 44-52.
- Gould, S.J. (1997). The exaptive excellence of spandrels as a term and prototype. *Proc. Natl. Acad. Sci. USA*, 94: 10750-10755.
- Gould, S.J. and Lewontin, R.C. (1979). The Spandrels of San Marco and the Panglossian Paradigm: A Critique of the Adaptationist Programme. *Proceedings of the Royal Society of London. Series B, Biological Sciences*, 205(1161) 581-598.
- Gould, S.J. and Vrba, E.S. (1982). Exaptation—A Missing Term in the Science of Form. *Paleobiology*, 8(1) 4-15.
- Governor's Office of Energy Policy. 2013. Arizona Motor Fuel and Supply Distribution. Online: www.azenergy.gov/doclib/energy/az_motor_fuel_and_supply_distribution.pdf
- Grossberg, S. (1980). How does a brain build a cognitive code? *Psychological Review*, 87, 1-51.
- Gunderson, L. and Holling, C.S. (Eds.) (2002) *Panarchy: Understanding transformations in human and natural systems*. Washington, D.C.: Island Press.
- Gunderson, L. and Pritchard, L. (Eds.) (2002) *Resilience and the Behavior of Large-Scale Systems*. Washington, D.C.: Island Press.
- Gunderson, L. H. (2000). Ecological Resilience—In Theory and Application. *Annual Review of Ecology and Systematics*, 31(1) 425-439.

- Gunderson, L., Holling, C.S. and Light, S.S. (Eds.) (1995) *Barriers and bridges to the renewal of ecosystems and institutions*. New York: Columbia University Press.
- Gunderson, L., Holling, C.S., Pritchard, L., and Peterson, G. (2002). Resilience of Large-Scale Resource Systems. In Gunderson, L. and Pritchard, L. (Eds.) (2002) *Resilience and the Behavior of Large-Scale Systems*. Washington, D.C.: Island Press.
- Gunderson, L., Pritchard, L., Holling, C.S., Folke, C., and Peterson, G. (2002). A summary and synthesis of Resilience in Large-Scale Systems. In Gunderson, L. and Pritchard, L. (Eds.) (2002) *Resilience and the Behavior of Large-Scale Systems*. Washington, D.C.: Island Press.
- Gunderson, L.H., Allen, C.R., and Holling, C.S. (2010) *Foundations of ecological resilience*. Washington, D.C.: Island Press.
- Hackley, R. (2013). Intel Plant Replenishes Aquifer for Thirsty Arizona City. Retrieved from www.bloomberg.com/news/2013-09-15/intel-plant-replenishes-aquifer-for-thirsty-arizona-city.html
- Holling, C. S. and Meffe, G. K. (1996). Command and control and the pathology of natural resource management. *Conservation Biology*, 10: 328–327.
- Holling, C.S. (1973). Resilience and stability of ecological systems. *Annual Review of Ecology and Systematics* 4:1-23.
- Holling, C.S. (1986). Resilience of ecosystems: local surprise and global change. In W.C. Clark and R.E. Munn (Eds.) *Sustainable development and the biosphere*. Cambridge University Press, Cambridge, UK. Pp. 292-317.
- Holling, C.S. (1987). Simplifying the complex: The paradigms of ecological function and structure. *European Journal of Operational Research*, 30: 139-146
- Holling, C.S. (1996). Engineering resilience versus ecological resilience. In P. Schulze (Ed.) *Engineering within ecological constraints*. National Academy Press, Washington, D.C., USA. Pp. 31- 44.
- Holling, C.S., Gunderson, L. and Peterson, G. (2002). Sustainability and Panarchies. In Gunderson, L. and Holling, C.S. (Eds.) (2002) *Panarchy: Understanding transformations in human and natural systems*. Washington, D.C.: Island Press. Pp. 63-102.
- Hollnagel, E., Paries, J., Woods, D. and Wreathall, J. (Eds.) (2011). *Resilience Engineering in Practice: A Guidebook*. Surrey, GBR: Ashgate Publishing Group. Retrieved 9 October 2012, from <http://lib.myilibrary.com?ID=290724>

- Hollnagel, E., Woods, D.D., Leveson, N. (Eds.) (2006). *Resilience Engineering: Concepts and Precepts*. Ashgate, Aldershot, UK. Online: <http://site.ebrary.com/lib/asulib/Doc?id=10211227&ppg=34>
- Jayaram, N. and Srinivasan, K. (2008). Performance-based optimal design and rehabilitation of water distribution networks using life cycle costing. *Water Resources Research*, 44(1) W01417ff.
- Jen, E. (2005). *Robust Design: A Repertoire of Biological, Ecological, and Engineering Case Studies*. New York: Oxford University Press.
- Kelly, R.L. (1995). *The Foraging Spectrum: Diversity in Hunter-Gatherer Lifeways*. Washington, D.C.: Smithsonian Institution Press.
- KinderMorgan (2015). Pacific Operations Specification Manual. Retrieved March, 26, 2015 from http://www.kindermorgan.com/pages/business/products_pipelines/pacific_specs.aspx
- Kolmogorov, A.N. (1963). On Tables of Random Numbers. *Sankhyā: The Indian Journal of Statistics, Series A*, 25(4) 369-376.
- Low, B., Ostrom, E., Simon, C., and Wilson, J. (2002). Redundancy and diversity: do they influence optimal management? In Berkes, F., Colding, J., and Folke, C. (Eds.). *Navigating Social-Ecological Systems: Building Resilience for Complexity and Change*. West Nyack, NY: Cambridge University Press. 83-114.
- Maynard Smith, J. and Szathmary, E. (1999). *The Origins of Life*. New York: Oxford University Press.
- McShea, D.W. (2001). The Hierarchical Structure of Organisms: A Scale and Documentation of a Trend in the Maximum. *Paleobiology*, 27(2) 405-423.
- Mulligan, C. and Sala-I-Martin, X. (2000). Measuring aggregate human capital. *Journal of Economic Growth*, 5(3) 215-252.
- Murray–Darling Basin Authority (MDBA) (2012). *Sustainable Rivers Audit 2: The ecological health of rivers in the Murray–Darling Basin at the end of the Millennium Drought (2008–2010). Summary*. MDBA: Canberra, AUS. Online: mdba.gov.au/what-we-do/mon-eval-reporting/sustainable-rivers-audit
- NIST (1993). Federal Information Processing Standards Publication 183: Standard for Integration Definition for Function Modeling (IDEF0). Available online at: idef.com/pdf/idef0.pdf.

- Page, S. (2011). *Diversity and Complexity*. Princeton, NJ: Princeton University Press.
- Parrott, L. (2010) Measuring Ecological Complexity. *Ecological Indicators*, 10(6) 1069-1076.
- Peterson, G., Allen, C., and Holling, C.S. (2010) Ecological Resilience, Biodiversity, and Scale. In Gunderson, L.H., Allen, C.R., and Holling, C.S. (Eds.) *Foundations of Ecological Resilience*. Washington, D.C.: Island Press. (Originally published in 1998).
- Prasad, T., and Park, N. S. (2004). Multiobjective genetic algorithms for design of water distribution networks. *Journal of Water Resource Planning and Management*, 130(1): 73–82.
- Putnam, R.D. (2000). *Bowling Alone*. New York: Simon & Schuster.
- Rapoport, A. (1986). *General Systems Theory*. Cambridge, MA: Abacus Press.
- Resilience Alliance (2010) *Assessing resilience in social-ecological systems: Workbook for practitioners*. Version 2.0. Online: <http://www.resalliance.org/3871.php>.
- ReVelle, C., Whitlatch, E.E. & Wright, J.R. (2004). *Civil and Environmental Systems Engineering*. Upper Saddle River, NJ: Pearson Prentice Hall.
- Rittel, H. and Webber, M. (1973). Dilemmas in a General Theory of Planning. *Policy Sciences*, 4(2) 155-169.
- Rosen, R. (1985) *Anticipatory Systems: Philosophical, Mathematical & Methodological Foundations*. Oxford: Pergamon Press. (second edition published 2012 by New York: Springer)
- Saaty, T. (1980). *The Analytic Hierarchy Process: Planning, Priority Setting, Resource Allocation*. New York: McGraw-Hill.
- Simon, H.A. (1962). The Architecture of Complexity. *Proceedings of the American Philosophical Society*, 106(6) 467-482.
- Tamaskar, S., Neema, K., and DeLaurentis, D. (2014). Framework for measuring complexity of aerospace systems. *Research in Engineering Design*, 25(2) 125-137.
- Todini, E. (2000). Looped water distribution networks design using a resilience index based heuristic approach. *Urban Water*, 2(2): 115–122.

- Tononi, G., Sporns, O. and Edelman, G.M. (1999). Measures of degeneracy and redundancy in biological networks. *Proc. Natl. Acad. Sci. USA*, 96(6) 3257-3262.
- Vale, L.J. and Campanella, T.J. (2005). *Resilient City: How Modern Cities Recover from Disaster*. Cary, NC: Oxford University Press.
- Varela, F.J. (1979). *Principles of Biological Anatomy*. New York: North-Holland/Elsevier.
- von Bertalanffy, L. (1968). *General Systems Theory*. New York: George Braziller.
- Walker, B. and Salt, D. (2006). *Resilience Thinking*. Washington, D.C.: Island Press.
- Walker, B.H., Abel, N., Anderies, J.M., and Ryan P. (2009). Resilience, Adaptability, and Transformability in the Goulburn-Broken Catchment, Australia. *Ecology and Society*, 14(1): 12. [online] www.ecologyandsociety.org/vol14/iss1/art12/.
- Walker, B.H., Holling, C.S., Carpenter, S.R. and Kinzig, A. (2004). Resilience, adaptability and transformability in social–ecological systems. *Ecology and Society* 9(2): 5. [online] www.ecologyandsociety.org/vol9/iss2/art5.
- Wigner, E. (1960). The Unreasonable Effectiveness of Mathematics in the Natural Sciences. In *Communications in Pure and Applied Mathematics*, 13(1). New York: John Wiley & Sons, Inc.
- Williamson, J. (1951). *Seetee Ship*. New York: Gnome. Based on short stories published in *Astounding Science Fiction* magazine, 1942-3.
- Williamson, J. (2001). *Terraforming Earth*. New York: Tor.
- Woods, D.D. (2009). Escaping failures of foresight. *Safety Science*, 47: 498–501.
- Zolli, A. (2012). *Resilience: Why things bounce back*. New York: Free Press.

APPENDIX A
FRACTALSYS OVERVIEW

FractalSys is a system modeling tool (and model-based reasoning engine) that represents any system as a collection of an arbitrary number of other systems. FractalSys can be used in off-line system modeling applications for analysis, design, or trade space characterization, or it can be directly embedded in operational control systems. The model-based approach allows engineers to develop their systems in an incremental manner which better supports the common merging of top-down and bottom-up design approaches observed in a given project life cycle.

The “big picture” is quickly and easily captured in the model and details and fidelity can be added as they become important or necessary. Since the model-based approach to design and development is incremental, there is no problem defining operations and progressively augmenting them to increase their fidelity. Details can be added whenever they are required and can be simulated (or simply not referenced) until it is appropriate to augment the model for higher fidelity. System demonstrations do not require all the subsystems to be functioning and interim progress is easily assessed by taking an overall view of the model. Internal consistency is enforced and transparency is facilitated.

Systems can be treated as black boxes, placed in a library and used as necessary in other efforts. They can be included in larger systems or isolated for individual test and simulation. When verified and tested, they can be used as necessary without concern for what is on the inside. They manage messages as designed and affect no external data. Their low level details can be invisible to the casual user when necessary or exposed to the critical eye if required. The built-in simulation capability is useful to verify all the required parts are defined. Function of one system is easily borrowed for inclusion in another system.

The system hierarchies represented by FractalSys are nested containment hierarchies, though “containment” should not be construed as restricting such hierarchies to physical or spatial ones. Hierarchies can be purely conceptual as necessary, for example, as employees in a corporation. Such system decompositions can be considered hierarchies of “nearly decomposable” systems (Simon, 1962) in which *intra*-component linkages are generally stronger than *inter*-component linkages. Modularization of this sort is common practice in systems engineering. Note as well that Simon’s near decomposability is a convenient way to allow for future expansion and to support incomplete or bounded human expertise in the sense that less-understood or not-yet-finalized systems can remain as “black boxes” while their interfaces solidify. Ultimately, deployed systems must deliver their functions rationally, so solid systems engineering practice must ensure system completeness to the extent required to meet valid requirements. Note, however, that while some “models” used in simulation or analysis can remain incomplete and allow for future expansion, operational systems must be decomposed to the extent required to meet functional specifications. This may seem to be a trivial observation, but it is important in light of the manner in which FractalSys can be used for off-line modeling and analysis, or embedded in an actual operational system.

As shown in Figure 28, each system is composed of other systems and can be defined to assume any of an arbitrary number of states. Further, an arbitrary number of transitions can serve to move the system between its states. Systems at every level are identified by (1) their function as represented by their behavior when operating in specific states (which are frequently a useful proxy for function) and (2) their interfaces—especially those that “actuate” other (external) systems. Hence, system identity is an important concept represented by inter-system arrangement, interfaces, and state. System identity can be defined in (typically) equal parts structure and function. Where experts debate how a system should be *structured*, function can take the lion’s share in providing system identity until consensus is achieved. Where system *function* is debated, structure is usually better understood. Function and structure find themselves in a sort of Heisenberg balance.

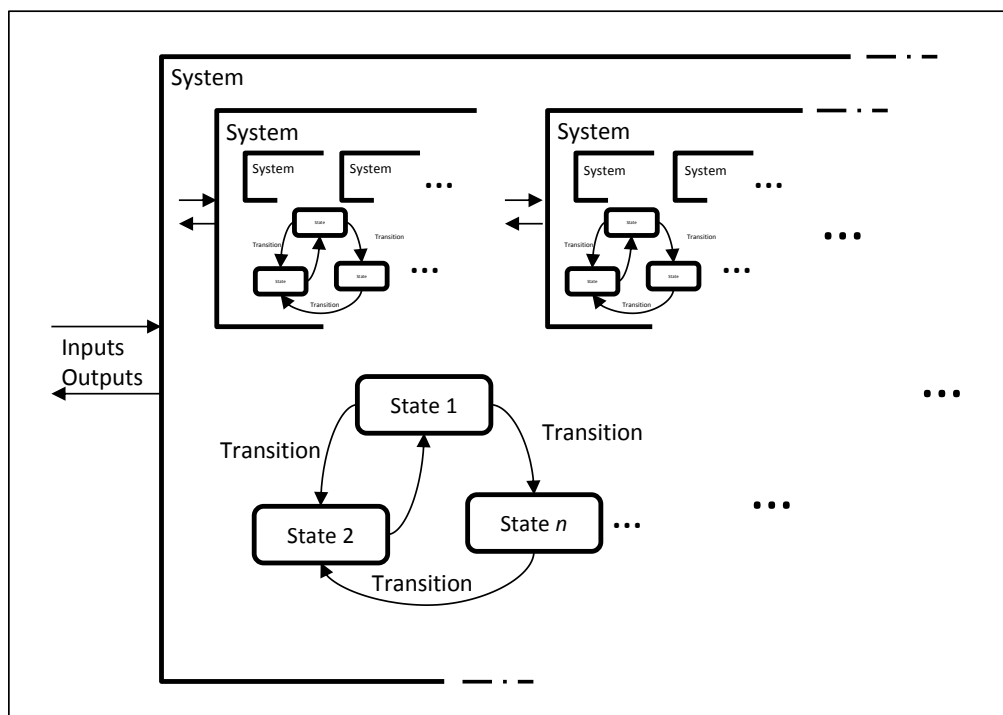


Figure 28. “Fractal” Nature of Systems Modeled by FractalSys

Ashby (1956, p. 25) suggests the state of a system is “any well-defined condition that can be recognized if it occurs again.” More precisely, states are a summarized aggregation of the values of an arbitrary (though generally consensus-based) number of system variables (sometimes referred to as “observables”, “measurables”, or “sensors”) and configuration parameters. Variables and parameters are similar concepts, but are typically differentiated by how fast they change and who (or what) controls them. In general, variables are faster changing and are generally implicated in the fast *microdynamics* of dynamic system theory. Parameters are usually “settings” that broadly control how the system operates and hence are slower to change. They fit better into the slow *macrodynamics* of systems theory and are frequently static over certain system-specific time frames (for a discussion of “comparative statics” see Rapoport, 1986, p. 66ff). In general, states are “named

concepts” superimposed by humans on certain configurations and behaviors of a system. Variables can be external sensor readings or the states of other contained systems.

Transitions are the mechanisms by which a system is moved among its states in either a reactive or anticipatory manner (cf. Rosen, 1985). In fact, the concept of nearly decomposable systems is important in this context since in anticipatory modes, FractalSys can identify missing transitions if it is compelled to do something that is currently undefined. The reactive/anticipatory distinction is important since FractalSys can be used in both control and simulation environments. This is accomplished by manipulating external system actuators or the surrounding environment (which, recall, is just a larger “containing” system). Transitions can be triggered by external data shared among systems or by other transitions within the parent (or containing) system.

As implied by the name, FractalSys supports this hierarchical system-state-transition structure to an arbitrary depth as needed to completely model the system of interest to the level of fidelity required for a specific analysis or control application. FractalSys does not enforce arbitrary rules concerning depth of modeling, and assumes domain experts are employed to ensure appropriate fidelity for particular analysis, simulation, or control tasks.

FractalSys offers a black-box view of systems that ultimately proves to be all that is needed for adequate modeling fidelity. Though a system’s identity is defined by both structure and function, from the outside, a system is understood by what it presents via its interfaces (measurable inputs and outputs). Where internal structure is important to external systems, it can be “presented” as an interface and queried as necessary. Output parameters and variable values serve to define the state of the system as it might be observed by a containing or sibling system. Inputs allow a system to be internally impacted by parent or sibling systems. Changes to the surrounding environment (that is, the containing system) can be modeled as changes in system configuration parameters and variables. Specifically directed impacts (e.g., commands, or forced transitions) are managed through exposed (and controlled) interfaces.

FractalSys considers all activity as a series of system events (driven by data update cycles or a regular update “polling” period). FractalSys monitors all system events and induces state transitions as a function of the changes in variables and parameters (including state variables) in a depth-first manner (which is appropriate for Simonesque “nearly decomposable” systems). The state transition mechanism is extremely important to FractalSys. Sequencing is initiated by, and continued on, the occurrence of events. Each of these events can have one or more transitions associated with it. Once the transition engine has completed the cascade of transitions, the system waits until it needs to act on new incoming data. Note that while the *controlling* model might be quiescent, *controlled* models and attached “real” systems usually continue to be very active performing their simulated or real functions.

When effecting system transitions, FractalSys traverses a recursively defined hierarchy of transition objects using a depth-first search algorithm. This object hierarchy is depicted in Figure 29.

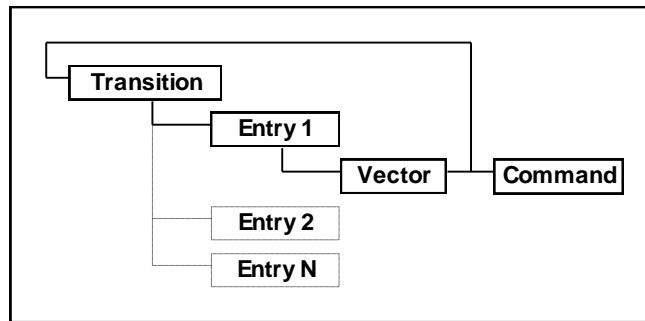


Figure 29. Recursive Definition of Transition Objects

As shown in Figure 29, a transition is comprised of a number of transition entries. Among other things described below, each entry contains a transition vector which is to be executed. This vector can be another transition or a terminal command object that can be used to interface FractalSys to external systems or simulators. In the case of a command, data is transmitted via the command interface to an external system or simulator. When the transition vector is another transition, the transition hierarchy must be recursively traversed to perform the command and/or transition sequences which must be executed in order to satisfy the highest level transition. Since the transition entry is the most important atom of processing (for control or simulation), a brief overview of the standard processing which occurs for each entry is provided. The transition entry object contains the following items:

1. transition vector (another *transition* or a terminal *command*),
2. verification vector (a *state*, comprised of a specified list of system variables and values),
3. timeout (nominally a number of seconds),
4. failure vector (a *transition*, which could be comprised of other entries), and
5. pause (nominally a number of seconds).

When a transition entry is executed, the transition vector is the primary item of actuation. As previously mentioned, in the case of a simple command (which must be defined and instrumented in the command interface database), the command is transmitted via the external commanding interface which will be unique for specific implementations. Thereafter, the transition manager will wait for the verification vector to be satisfied. It will do this by checking the system state until either the requirement is satisfied, or the timeout is exceeded. If the requirement is satisfied prior to a timeout condition, the transition manager will pause for the number of seconds specified before proceeding to the next transition entry (if any). If the timeout condition occurs prior to satisfaction of the verification vector, the failure vector is executed. This failure vector is itself a transition which is managed in the same way. Typically, the failure vector will display error messages and/or execute a safing sequence if FractalSys is connected to hardware.

For safety-critical systems, constraints and restriction processing is effectively built into the transition management routines in two ways. First, many transitions will contain multiple entries, the first of which will verify proper system configuration before the remainder of the entries are executed to change system state. Second, all state transition requests will be “validated” prior to their execution. For example, given a request to transition a system from state A to state B, there will be an internal check to ensure it is a legal (i.e., defined) transition from the current state. Obviously, if no transition is defined, the desired state is unreachable. A violation of such a defined constraint or restriction will generate an alert message.

The transition engine operates from the functional model of the system which defines the systems, states and transitions necessary to successfully simulate or control the system. The system model (function and structure) is developed in the FractalSys software package (Figure 30) and is saved in a database.

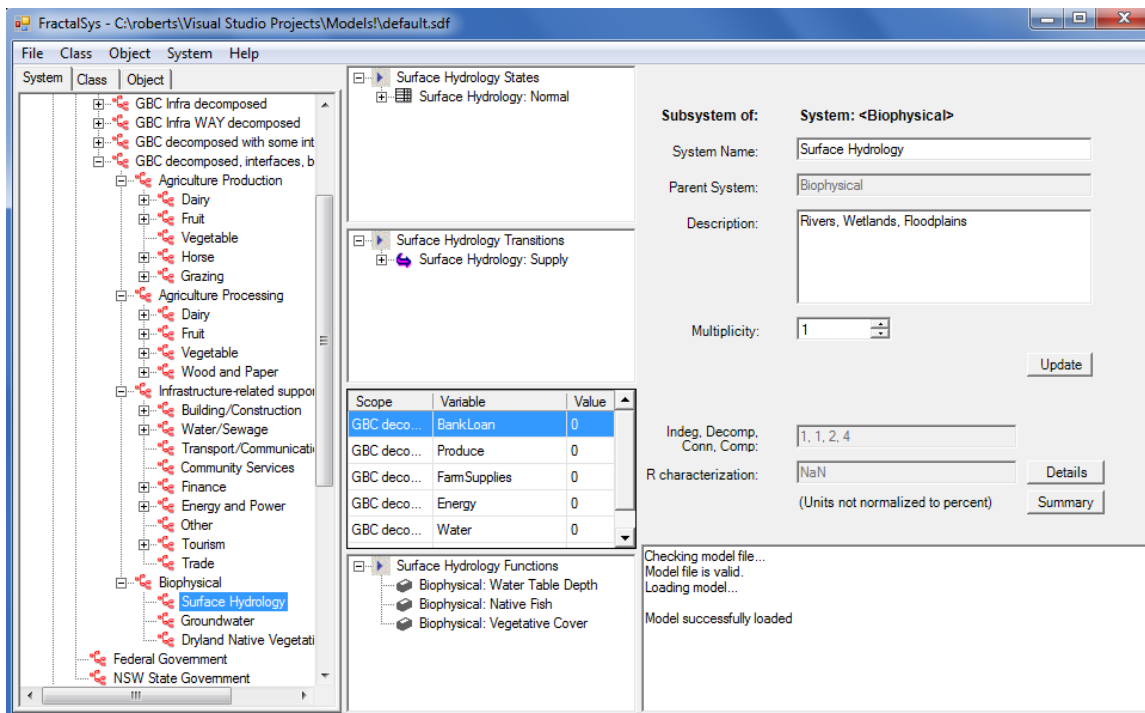


Figure 30. FractalSys User Interface

System “function” occurs while a system resides in a particular state (in general, the model orchestrates the operation of other models). For example, think superficially of a two-state system where it is either “on” or “off”. Clearly, such a system would not be performing its function while in the “off” state. Arguably, function also occurs during transitions, though the best models will limit this to simple “notifications of change.” In this regard, transitions should be viewed as “catastrophic” in the mathematical sense of catastrophe theory (see Rapoport, 1986, p. 67ff). Further, function occurs while a system

is “maintaining” its state—though for most homeostats this is generally invisible to the surrounding or containing systems. Hence, “state” becomes a viable proxy for “function” assuming the granularity of the system definition has been managed well by the experts who define it.

Obviously systems maintain interfaces with other systems. This can be implemented by shared variables or by modeling interfaces as systems themselves. For example, shared variables like “temperature” can be used in state determination by two proximate systems. Such variables allow any data, information, influence, material, etc. that traverses an interface to be modeled as raw data items that are modified by one system and used by another system in determining state. So-called “interface systems,” on the other hand, can be helpful if interfaces are more complicated. For example, to model an RF interface, a “system” can be employed so delays (due to distance) and disturbances (like rainfall) can be specifically modeled. This also allows “interface” systems to have states and transitions. Thus, interfaces are not ontologically “other” and require no special considerations. They are simply systems themselves.

Variables represent measurable items or user-definable items that contribute to definition of system states. Unlike the states they help define (which are fabrications superimposed on systems by human observers), the idea behind variables is that they are more “real” and tethered in the physical world. This is not always the case, but (avoiding the epistemological discussion) the intent is that system states which are based on measurable variables are more grounded in physical reality. For this reason, each “measurable” has a scope or a system to which it is most accurately associated. For example, since an environment temperature might be measurable by many systems, its scope is clearly beyond any one of the systems that measure it, and might properly be best associated with a parent system (like, the environment).

As expected, a variable may contribute to defining many states. For example, “temperature is 100F and humidity is 90% so the weather is *muggy*”, and, “temperature is 70F and humidity is 30% so the weather is *nice*” are both valid state definitions which share variables like temperature and humidity. Variables can be employed in state definitions for any system for which they are within scope.

As shown in Figure 31, it is the extent to which variables are shared by systems that defines the *connectedness* of those systems. Consider a variable that is “controlled” (or “set”) by a certain system X (e.g., a temperature is “set” by an external weather or climate system). If that variable is used to define the state of another system (e.g., Y employs it to determine if it is “hot” outside), there is a “direct connectivity” between X and Y that constitutes a connection. *Setting* a variable that is used by another system increases *connectedness* between systems. *Using* a variable which is set or controlled by another system specifically creates a *dependency* for the using system. Note that if another system (Z) uses the same measurable variable from X in defining one of its states, this constitutes “indirect connectivity” between systems Y and Z (see Figure 31).

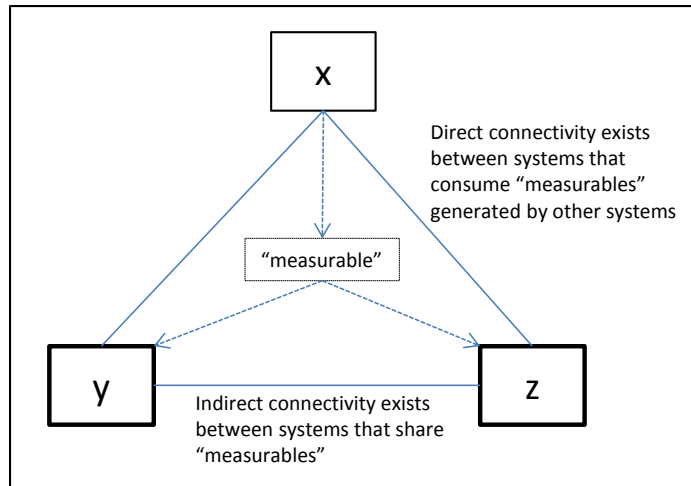


Figure 31. Direct and Indirect Connections Between Systems

FractalSys observes fairly typical scoping and accessibility rules for variable data. As indicated above, any *state* or *transition* that uses a variable defined in an *external* scope (that is, *inherited*, see Figure 32) creates a dependency. FractalSys does not acknowledge directionality of interfaces since there was no driving reason to support such ideas (i.e., all interfaces are bi-directional, though some may only be employed in one direction).

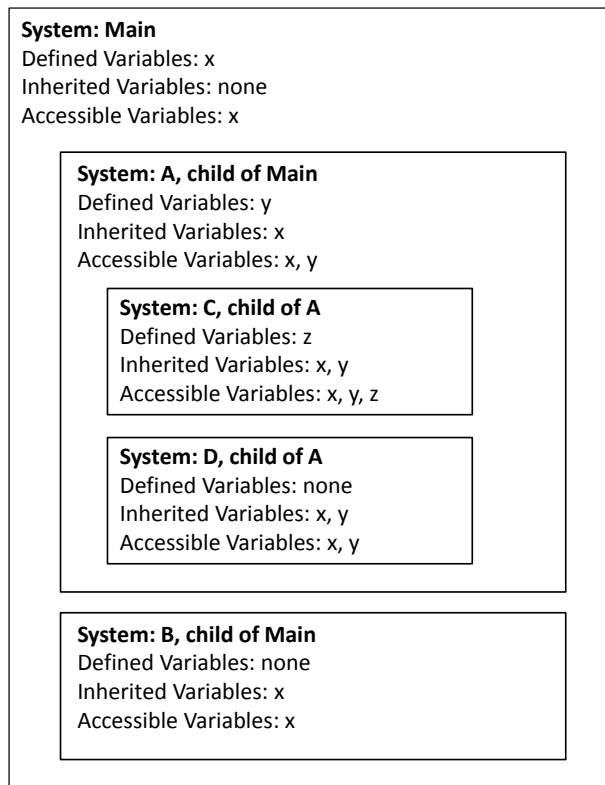


Figure 32. FractalSys Variable Scoping Protocol

Data Model

The FractalSys data model is highly recursive and very dense (see Figure 33).

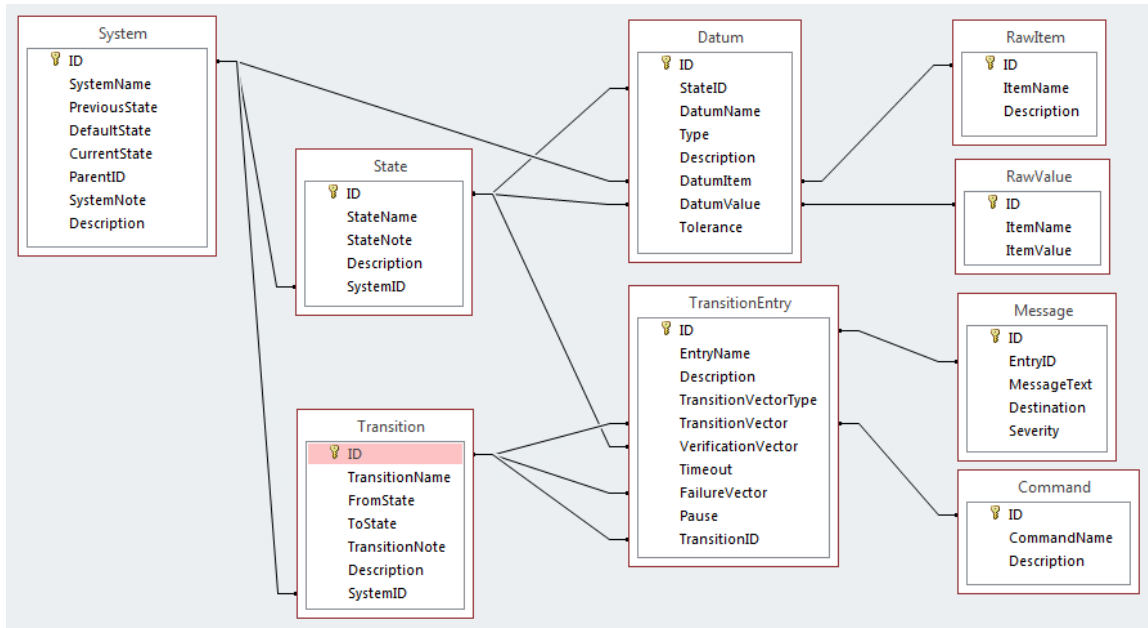


Figure 33. FractalSys Data Model

In general, FractalSys supports an object-oriented physical model of whatever system it models. This physical model of the system is positioned within some reasonable “environment,” then starts at the highest containing level and adds detail as it progresses down the decomposition hierarchy. For example:

*A rocket launch vehicle might have four stages
Stage 1 might have four SRB Castors, Stage 2 has two SRB Castors, etc.
An SRB Castor might have a TVC system
TVC systems have Electrical Control Systems (ECS)
ECS's have.... etc.*

Eventually, the hierarchy decomposes to the level of detail where sensors are available to attach real values to these items. For example:

*A launch vehicle has a Stage 3.
Stage 3 has a Fairing.
The Fairing has a temperature sensor.
The temperature sensor provides an analog value in degrees.*

Thus, the physical model of the system completely defines what can be known (e.g., measured or dictated) about the system.

On this *physical* model of the system, a *topical* model can be superimposed if desired. Where the physical model addresses *completeness*, the topical model addresses *clarity* and *context*. For instance, in the example above, the temperature sensor may indeed be physically located on the stage 3 fairing of the rocket system, but from the perspective of a hypothetical thermal engineer, it may be considered part of the *thermal subsystem* which is a hypothetical subsystem distributed throughout the actual physical configuration. Likewise, while for example there may be an ignition safety device physically located in every rocket castor, from the perspective of an ordnance engineer they all belong to the *ordnance subsystem*. These parallel models support both approaches to the system design and allow a very flexible approach to status monitoring and state determination. Further, the alternative model views add negligible overhead to the system model.

A database containing the latest model values (“measurables” and states) is maintained. This database is accessed by the physical and any topical models of the system. Generally, system state determination is handled by the two (or more) separate models of the system. The physical model (as the most complete representation of the “measurables”) processes data and develops a physical status of the system. For example, system equipment can be in any of a number of states: on, off, standby, etc. While specific details of the physical system will be available, it may be largely ignored in lieu of reviewing the more human-oriented topical model states. These states can be generated by the topical model of the system. Hence, the physical model provides a complete skeleton to which the more “interpreted” topical models can add flesh. However, in practice, system engineers will opt to model the system according to personal or team preference.

Typically, in hardware-in-the-loop systems, state determination is done at predefined intervals during which data is read from the various input data sources and is placed in the database (the “polling” approach). Thereafter, the state models are updated to reflect the current system state based on the physical data sources.

FractalSys forces both the physical and topical models (if developed) of the system to update their states on a regular basis. This is done by messaging the highest level object and telling it to update its state. This message will propagate down the hierarchy (with a depth-first algorithm) until the lowest level object updates its state. As this is done the state propagates up the hierarchy causing all components and subsystems to derive their current states. Finally the top level system state will be derived and in the process all the database variables will be updated.

Absolute tolerances for each individual data item will be established as required. However, each tolerance will have a context (or, multiple contexts) associated with it by a cognizant engineer. For example, if a power supply is ON, then perhaps the voltage

should register 32 ± 6 Volts, however, in a different context, if it is turned OFF for example, then appropriate values for its voltage may be 0 ± 50 mV. These tolerances, defined in the system definition database, will be used to clarify system states. Anomalies will be flagged if the required values are unknown or if they are out of tolerance for a particular context.