SAND2014-4860C

Working Paper

Title: Studying the relationship between system-level and component-level resilience

Authors: Michael D. Mitchell (micmitc@sandia.gov) & Walt Beyeler (webeyel@sandia.gov)

Key Words: Resilience, Diversity, Adaptation, Agent-Based Modeling, Complex Adaptive Systems

Abstract

The capacity to maintain stability in a system relies on the components which make up the system. This study explores the relationship between component-level resilience and system-level resilience. The study aims to identify policies which foster system-level resilience in situations where existing incentives might undermine it. We use an abstract model of interacting specialized resource users and producers, which can be parameterized to represent specific real systems. We want to understand which features of a system, such as input resource stockpiles, demonstrate the efficacy of system-level resilience policies. Systems are subject to perturbations of varying intensity and frequency. Our study focuses on creating a simplified economy which imposes an inventory carrying cost to incentivize smaller inventories. Our study examines how components with varying inventory levels compete in environments subject to periods of resource scarcity. The results show that policies requiring larger inventories foster higher component-level resilience but do not foster higher system-level resilience. Inventory carrying costs reduce production efficiency as inventory sizes increase. JIT inventory strategies improve production efficiency but do not afford any buffer against future uncertainty of resource availability.

Introduction

Resilience is the ability of a system to recover from shocks. Economic turmoil, political instability, and natural disasters are examples of shocks which can stress or destabilize a system. Uncertainty in the future availability of critical resources is a concern to policy makers [Simangunsong 2012, Brown and Lall 2006]. Understanding and fostering the resilience of key systems, such as infrastructures, is a key public policy goal of Presidential Policy Directive 21 (PPD-21) [Obama 2013]. Resilience can be measured using the movement of a system indicator, such as the flow of a key resource through the system [Vugrin et. al. 2014]. The resilience of a system is an emergent property of the resilience of its components because system resilience depends on how the components interact and not simply on their individual resilience. Policies which optimize component resilience may not optimize system resilience, and vice versa.

We explore this trade-off using a simple model system containing two kinds of interdependent entities. Each kind consumes the distinct resource produced by the other. Entities keep inventories of their input and output resource which can be used to buffer periods of scarcity. For one kind of entity, individual members keep different input inventory levels. Larger inventories are costlier to maintain but provide a better buffer against supply disruptions. Competition among individual entities forces a tradeoff between inventory cost and benefit. The outcome of this competition is determined by the frequency of random shocks arriving from the environment.

Agent-based models (ABMs) are an effective means of modeling, understanding, and measuring resilience of complex systems [Datta et. al. 2007, Jennings 2000, Sandholm 1999]. ABMs allow for the simulation of actions and interactions among autonomous agents and their effects on the system as a whole. Agents have adaptive processes and interact with other agents in the system resulting in behavior that is complex and difficult to anticipate. This modeling methodology lends itself well to understanding how complex adaptive systems work, the interdependencies of system components, and how adaptive entities respond to endogenous and exogenous changes [Brown et. al. 2004, citations]. We developed a hybrid model, using system dynamics and agent-based modeling, to represent production and consumption sectors, resource flows, and market exchanges among interacting specialists (entities) in a system. All entities in the system produce and consume resources and trade those resources via a double auction market mechanism.

We conducted a study which explores the relationship between system-level resilience and component-level resilience. The aim of the study is to identify policies that foster system-level resilience in situations where individual incentives might undermine it. We want to understand what features of a system, such as input resource levels, determine the comparative effectiveness of system-level resilience policies. In this study, system-level resilience is defined as the ability for the system to maintain critical resource flows while the system is subjected to disruptions in resource availability [Arrow et al., 1995, Christopher and Peck 2004]. Component-level resilience is defined as the ability for agents in a system to maintain a sufficient local resource buffer to survive periods of resource scarcity. This study utilized the Exchange Model developed at Sandia National Laboratories to investigate complex adaptive systems (CAS) [Beyeler et. al. 2011]. This model provides a framework to abstractly represent a system in which interacting specialists (entities) produce and consume resources that flow among entities via continuous markets.

Model Description

The Exchange Model (ExM) combines system dynamics and agent-based modeling to represent production and consumption sectors, resource flows and market exchanges among interacting specialists (entities) in a system (see Figure 1). All entities have a homeostatic process which maintains 'health' via the consumption of resources. The production of resources is driven by the consumption of resources. Entities store both the resources needed for consumption and production and control those resource levels through interactions with markets. Markets facilitate the exchange of resources by using a double auction algorithm to match bids and offers. Each market manages the exchange of a single resource. The price of a proposal (bid or offer) is reflexive and represents the relative scarcity or abundance experienced by the entity creating the proposal. The environment determines the availability of resources that entities require for survival. Entities use environmental signals to determine the amount of resources to consume and produce.

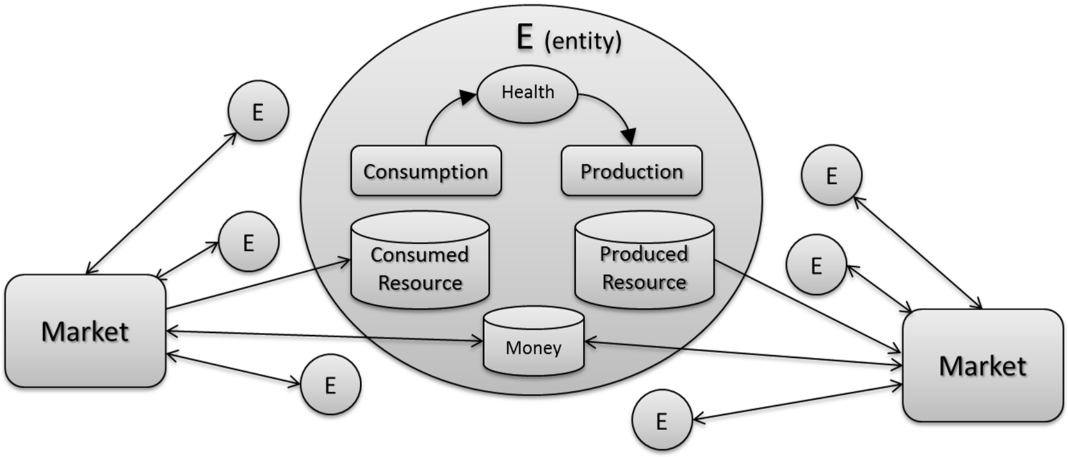

Figure 1 - Illustration of an entity and how entities interact with the system

Entity

Entities maintain their health via a homeostatic process involving the consumption of specific resources. Healthy entities can produce other kinds of resources. Parameters of thefunctions governing consumption, production, and health can be varied to reflect various types of real systems. An entity's behavior is reflexive to its governing equations and control processes. Entities exchange resources among one another via a double auction market mechanism. An entity's health will decline if the entity is unable to obtain the input resources needed for consumption from the market. There are two primary reasons an entity would not be able to obtain the resources it needs from the market. First, the resource is not available in the market. This can be due to resource scarcity or the unavailability of the entity producing the resource. Second, the entity is not able to sell its outputs and generate the money it needs to purchase its inputs.

The environment determines whether entities can sell produced resources to acquire the necessary resource inputs on sustainable terms. The environment is made up of other kinds of entities with complementary requirements. The flow of resources in an environment can be disrupted by shocks to resources stores in some parts of the system resonating through the system via these dependencies. Entities maintain stores of input and output resources to buffer against uncertainty of resource availability.

Production Efficiency

One of an entity's governing state variables describes the function of health on potential production [Beyeler et. al. 2011]. When health exceeds its nominal value of 1, production can increase. Conversely as an entity's health declines, due to a scarcity of input resources, the entity can find a new operating equilibrium by running leaner at a lower health value. If the stress of scarcity becomes too great, then the entity's health value and production rate will decline rapidly.

Simple Model

For the purpose of this study, ExM is configured with two types of entities: A Producer, B Producer; and two types of resources: A, B (see Figure 2). The parameters for an entity type can be constant or have a probability distribution. A Producer produces 1 unit of resource (A) for each 1 unit of resource (B) consumed. B Producer produces 1 unit of resource (B) for each 1 unit of resource (A) consumed. We use this simple structure to explore the relationship between component-level resilience and system-level resilience by focusing on B Producers. We allow B Producers' input buffers to differ from one another, and use their health to indicate the component-level effect of their "decision". System-level performance is measured by the total production of B by all B producers. Our study focuses on component level and system level resilience. We are going to consider the health of B Producers to quantify component level resilience and the availability of resource B as a quantification of the system robustness to various strategies.

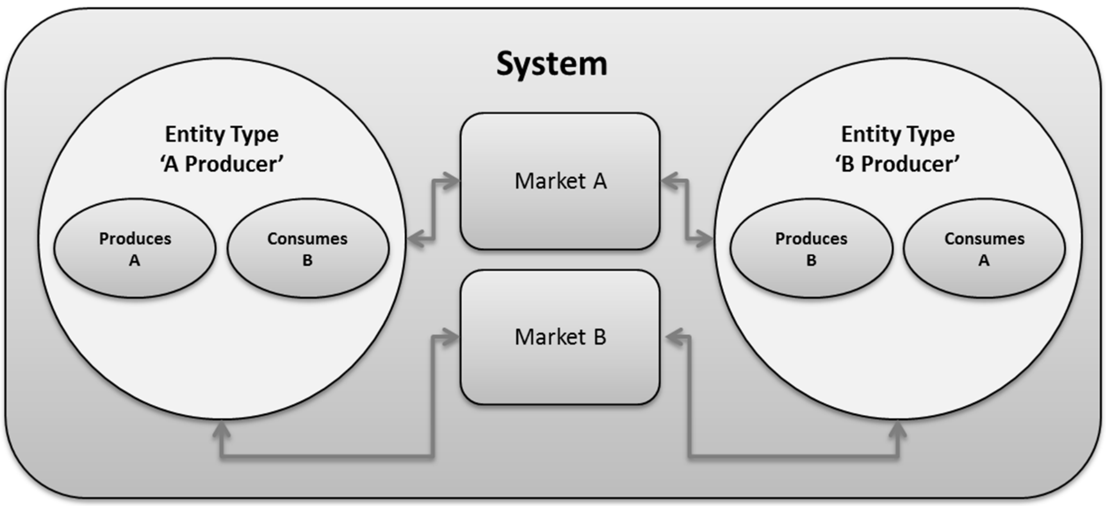

Figure 2 - Illustration of the Simple Model

To study how different configurations of B Producer and environments affect the production of our key resource (B), we configure the B Producer's health to be highly sensitive to its consumption of A. We disrupt the availability of resource A and measure how the scarcity of resource (A) affects the health of B Producers and the availability of resource (B). We will configure A Producers to ensure a constant system demand of the resource (B). B Producers who are not able to obtain the amount of inputs required for consumption either via the market or internal stores will decline in health until they are no longer viable. They are then replaced by another B Producer having a different (randomly chosen) input resource buffer size. In this way the population of B Producers evolves to withstand the particular frequency and intensity of shortfalls in resource (A) that define the environment.

We are comparing two different strategies for B Producers: adaptation and maximum inventory. The adaptation strategy allows for B Producers to be realized with random local inventory levels ranging from small to very large. The system in which B Producers utilize an adaptation strategy will have a population of B Producers with various inventory levels. The system in which B Producers utilize a maximum inventory strategy will have large inventory levels. A large inventory ensures a high degree of component-level resilience for B Producers.

Adaptation

The adaptation strategy requires B Producers to adapt to their environment via a replacement process. An entity that is not viable in the current environment will be replaced by one with different parameters. The parameters we are interested in varying are the parameters which govern an entity's inventory levels. We are varying the parameter tcstore, which controls how much, in terms of time, input resource will be stored. Entities with a small tcstore will have a just-in-time inventory, or minimum inventory, management strategy. Entities with a larger tcstore will have a maximum inventory strategy.

Maximum Inventory

The maximum inventory strategy requires B Producers to maintain a very large inventory compared to their consumption rate. All B Producers in this strategy will have large tcstore values. This strategy will model environments where inventory levels are often large due to either regulation or environmental conditions. Many critical industries, such as energy production, are required to have large inventories of key resources in an effort to provide higher system resilience.

Study

The strategic resource being monitored is the output of the B Producer entities, resource B. There is an inventory cost for B Producers on their input resource inventories, which incentivizes smaller inventories, but the environment is subject to periods of scarcity of resource (A), which incentivizes B Producer entities to have larger inventories. The goal is to compare the two inventory management strategies discussed earlier, just in time and maximum inventory, to determine which strategy conveys the highest component-level resilience and which strategy conveys the highest system-level resilience.

Inventory Carrying Cost

We have configured ExM to impose an inventory carrying cost to B Producers. Inventory carrying cost is modeled as a loss rate proportional to the inventory level. This abstractly represents the carrying costs for maintaining an inventory, such as physical storage, inventory taxes, and expiration of goods. The higher an entity's tcstore the higher the cost that entity pays for having a larger inventory. The inventory carrying cost produces selective pressures for the population of entities to have smaller tcstores.

Scarcity

In order to study the relationship between component-level resilience and system-level resilience, we need to create an environment in which B Producers experience scarcity. We script a disruption in the availability of resource A from A Producers during the simulation. The disruption of resource (A) causes B Producers to rely on their local inventories for consumption of resource (A). When the input resource for B Producers is scarce, there is selective pressure on B Producers to have larger tcstores. We model this scarcity by configuring a perturbation to stop A Producers' production of A at simulation time 10,000 through simulation time 11,000.Time, as defined in our model, is unitless (without a specific measure); we use it as a way to compare changes in state.

Simulation

We ran each simulation for a total simulation time of 20,000. A disruption was scripted at simulation time 10,000 which completely reduced the availability of resource (A); this scarcity lasted until simulation time 11,000. We imposed an inventory carrying cost of 2% per unit time. For each environment, we ran two simulations. One in which B Producers used the adaption inventory strategy, (those entities' inventory size was randomly selected from a distribution of small to large tcstore values) and one in which B Producer used the maximum inventory strategy (those entities only had large inventory sizes).

Each simulation has an initial population of 10 A Producers and 10 B Producers. Each entity consumes 1 input resource and produces 1 output resource. The nominal amount of each input and output resource is 10. A Producers do not have their health influenced by consumption of resource (B) and always maintain their production of resource (A). B Producers' health is highly sensitive to their consumption of resource A. During the simulation, B Producers which are not viable are replaced with new realizations from the B Producer factory's parameter descriptions. This replacement allows us to model adaptation in the selection of tcstore values for various environments.

Table 1 describes the two tcstore value ranges configured for each environment.

Table 1: Tcstore values ranges

Strategy	Tcstore Min	Tcstore Max
Adaption	0.2	5.0
Maximum Inventory	4.0	5.0

For each simulation, we quantify component-level resilience and system-level resilience. Component-level resilience can be quantified by measuring the average health change over time of the B Producers after the disruption. System-level resilience can be quantified by measuring the change in volume of resource (B) flowing through the market, since there is a constant demand for the key resource (B).

Analysis

Time Series

Table 2 describes the periods and simulation times we will use to discuss the results of our analysis.

Table 2 Periods and Simulation Times

Period	Time
Pre-disruption	1,000 - 10,000
Disruption	10,000 - 11,000
Recovery	11,000-20,000

Figure 3 is a time series plot of the B Producer's health trajectories over time. The results are differentiated by the two strategies under consideration. The upper plot illustrates the simulation where entities are realized with a random tcstore. The lower plot illustrates the simulation where entities are realized with only a large tcstore. The color of the entities is graduated based on their inventory size. The darkest entities have the largest inventories.

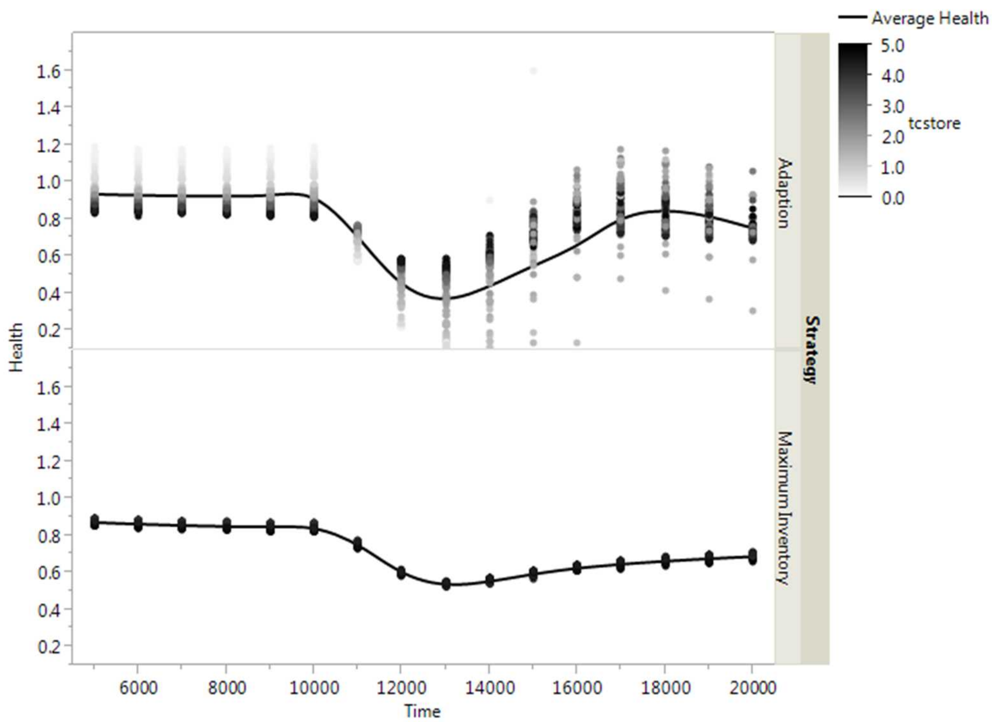

Figure 3 - B Producer Entities' Health over Time

The upper time series shows that in the pre-disruption period of the simulation entities with a smaller tcstore values have a higher health. This is intuitive due to the 2% carrying cost of holding inventory. Entities have to compete with one another for resources leading to less efficient entities having lower health values. Entities with larger tcstores are less efficient due to the higher costs associated with maintaining a larger inventory. In the beginning of the disruption period, entities with higher tcstores are healthier. The period of scarcity causes entities with the lowest tcstore to become unviable and they are replaced with new realizations. A period of recovery follows the period of scarcity where entities with larger tcstores are again less efficient and the larger tcstore entities are outcompeted by entities with smaller tcstores.

The lower time series shows that in the pre-disruption period of the simulation, all entities have similar health values due to the lack of diversity in the entities' parameters. As the period of scarcity begins, the entities' health values move together as they experience the scarcity of their input resource and rely on their local inventory levels to maintain their consumption of resource (A).

Comparing the relative health trajectories of each strategy allows us to talk about the component-level resilience of the strategies. A strategy requiring a larger inventory size lends itself to the B Producers being able to withstand a period of scarcity more effectively than B Producers with smaller inventories; we can see this by noting that the health of B Producers in the lower time series plot have a gentler slope during the period of scarcity compared to B producers in the upper time-series plot. However, the additional inventory size leads to lower overall health of entities in the lower time-series plot than those in the upper time-series plot due to the carrying costs of a larger inventory. Also note that entities which use an adaptive inventory policy have a faster recovery because those entities do not have to rebuild large inventories which were drawn down during the period of scarcity.

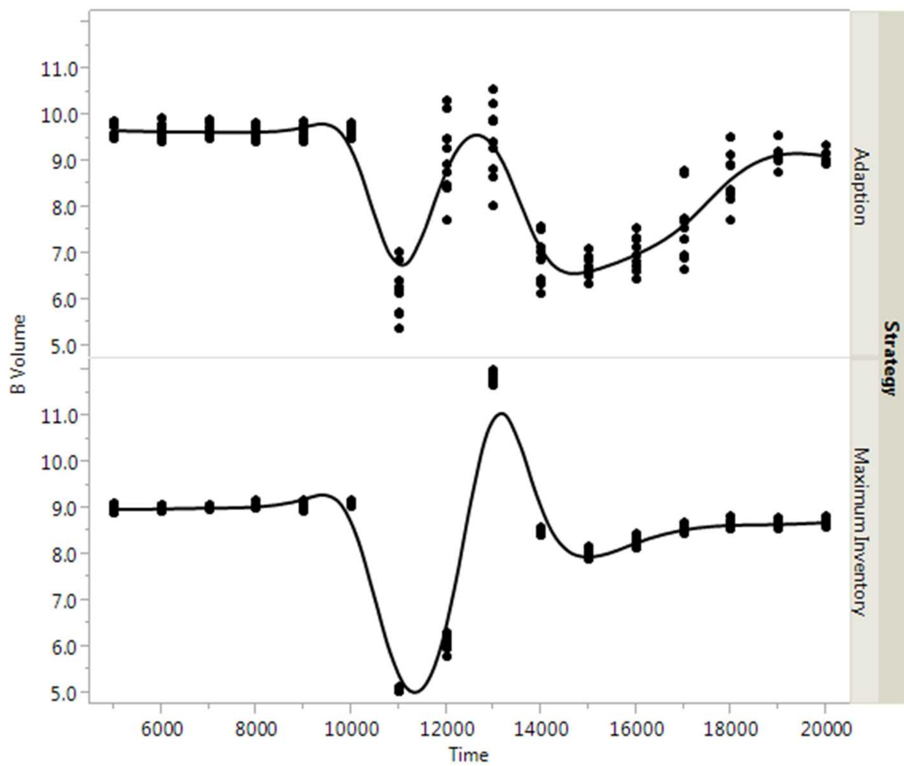

Figure 4 - Resource B Market Volume over Time

The upper plot, which illustrates the volume of resource being produced over time by the adaptive strategy, shows that initially there is a dampened effect on the amount of resources being produced because of the sharp decline in health. After an initial decrease in the production of resource (B), we can see the recovery begin followed by a second decrease in the production of B. Finally, there is a slow gain in production of resource (B) as the entities recover from the period of scarcity. The double-dip recovery is attributed to the adaptation strategy by which failed entities are replaced with new entities. As the first recovery is underway, entities with smaller inventories are unable to compete with entities with larger inventories. This is due to the entities with larger values being healthier after the period of scarcity. The entities with the lowest tcstore values fail and new entities replace the ones that failed. The decline of an entity’s health leads to lower levels of production which causes the second dip in the recovery. Once the entities that have failed are replaced, the amount of resource (B) being produced steadily increases.

The lower plot, which illustrates the volume of resources being produced over time by the maximum-inventory strategy, shows that during the period of disruption and early recovery there is an amplified effect on the amount of resources being produced considering the slight decline in health. The amplified effect on the decline in volume compared to the slight decline in health was puzzling at first. There were a couple of interactions in the model which caused this phenomenon. The carrying cost stressed entities in this strategy driving those entities to operate at higher production efficiencies to offset the cost of maintaining a large inventory. The higher production efficiency causes the entity to operate at a lower health value. These entities are operating at peak efficiency when the period of disruption begins and the shock is lessened by their buffer stocks but the result is the entities are driven to lower production efficiency.

Data

Table 3 shows the health over time for both strategies. The strategies with the highest health values are highlighted. The nominal health value for entities is 1. These values are averages of individual health values of all B Producers for each strategy. Comparing the average heath values by strategy provides us with a comparative analysis of which strategy affords higher component-level resilience.

Table 3 Average Health over Time by Strategy

Time	Adaptation (h)	Maximum Inventory (h)
1000	0.9835	0.9662
2000	0.9767	0.9520
3000	0.9693	0.9427
4000	0.9620	0.9345
5000	0.9567	0.9278
6000	0.9532	0.9221
7000	0.9509	0.9174
8000	0.9491	0.9133
9000	0.9478	0.9097
10000	0.9466	0.9067
11000	0.8488	0.8597
12000	0.7212	0.7795
13000	0.6787	0.7487
14000	0.7057	0.7568
15000	0.7514	0.7744
16000	0.8086	0.7876
17000	0.8261	0.7964
18000	0.8482	0.8033
19000	0.8436	0.8083
20000	0.8167	0.8113

The adaptation strategy has the highest average health values in the pre-disruption period. This finding is consistent with the efficiencies gained from some entities in the system having lower inventories in an environment where there is a 2% inventory carrying cost over time. This finding also explains why the standard deviation of health for B Producers is higher in the adaptive strategy during the pre-disruption period than the maximum inventory strategy. The adaptation strategy has lower health values than the maximum inventory strategy during the disruption and early recovery periods. This is due to the adaptation processes by which entities that were not viable during the disruption die and are replaced with new realizations. This replication process is also the reason the for a more robust recovery than the maximum inventory strategy. The health of an entity influences its production and consumption rates. The adaptation strategy has entities with low health die after the disruption and they are replaced with entities with high health values. These new entities are able to produce more of resource A than entities with lower health values.

The maximum inventory strategy has the highest average health values in the disruption and early recovery periods. This is because these B Producers have larger inventories which buffer against the period of scarcity.

Table 4: Market Volumes over Time by Strategy

Time	Adaptation	Maximum Inventory
1000	8.0738	6.6701
2000	9.0412	7.8178
3000	9.5037	8.4819
4000	9.6563	8.8673
5000	9.6282	8.9882
6000	9.5929	9.0025
7000	9.5779	8.9955
8000	9.6035	9.0283
9000	9.5830	9.0620
10000	9.6036	9.0979
11000	5.7851	5.0655
12000	9.0904	5.9843
13000	9.7515	11.8962
14000	6.9173	8.5282
15000	6.5649	7.9997
16000	7.1437	8.2429
17000	7.4425	8.5685
18000	8.1145	8.6531
19000	8.7804	8.6792
20000	8.9354	8.6874

Table 4 shows the average market volumes over time by strategy for resource (B). The strategy with the highest volume is highlighted. The nominal volume for resource (B) is 10 units. The average volume of resource (B) is an average of the market volumes for all the simulations for a given strategy.

The adaptation strategy has the highest volume in the pre-disruption period. Since entities in this strategy can have lower inventory levels than in the maximum inventory strategy, they are more efficient due to the relatively lower costs of a smaller inventory. What was surprising was that the adaptation strategy has the highest volume in the disruption period. This was puzzling at first until we considered that the replacement of anemic entities with new entities provided a boost in production of resource (B). This phenomenon leads to an overshoot in production and the system has a second dip with reduced production of resource (A). In the late stages of recovery, as the memory of scarcity fades, the adaptation strategy becomes the highest producer of resource (B).

The maximum inventory strategy is designed to be resilient during periods of scarcity. This resilience causes the components to be less efficient in periods without scarcity. This strategy also causes B Producers to have a more severe drop in production of resource (B) during a disruption. As the entities experience scarcity, they utilize their local resource buffers to mitigate the loss of resources from the market. The experience of scarcity causes the entities to reduce consumption of the scarce resource (A) causing a reduction in the production of resource B. For this same reason, the period of recovery is quicker compared to the adaptation strategy because production gradually returns to pre-disruption levels.

Conclusion

There is not a single strategy or policy which is most effective in fostering both high component-level resilience and high system-level resilience. There are tradeoffs to both waiting on adaption to select the most efficient components for a given environment and a policy of maximum inventory which buffers against uncertainty in future resource availability. The adaptation strategy fostered high system-level resilience, but had lower component-level resilience compared to the maximum inventory strategy which fostered high component-level resilience, but had lower system-level resilience.

The adaptation strategy can reduce the impact of carrying costs due to the selection of entities in the strategy having smaller inventories, which leads to more efficient production of resources. During the resource shock, entities with larger inventories were healthier and able to use their local resource buffers to reduce the shock; entities with small inventories health declined and some failed and were replaced with new entities. The diversity of inventory sizes and replacement of failed entities produced lower component-level resilience and higher system-level resilience then the maximum inventory strategy. Adaptation is more efficient in systems where there is low cost of entry, low level of consumer visibility or where the components are not sensitive to consumer confidence. A system in which entities select their own strategies provides less component level resilience overall, but the system is more efficient and periods of scarcity are less severe because all entities do not experience scarcity in the same way and new entities emerge and help the system to recover.

A maximum inventory policy to buffer against uncertainty in future resource availability is very effective at providing a high degree of component-level resilience. Larger inventories are an effective buffer against disruptions in resource availability. Where the maximum inventory did not do well was in its ability to maintain the production of a key resource to the system. The reason for this was that the stress of maintaining larger inventories drove entities in this strategy to a leaner, more efficient production rate in the pre-disruption period at the cost of lower health values for the individual entities. Once the disruption occurred, the scarcity combined with the previous system stress of the carrying costs on a large inventory caused a more severe drop in the production of the key resource than compared to the adaptation strategy.

In industries where there is a high consumer visibility and high customer confidence in the individual institution, then a high level of component-level resilience is necessary and would benefit from a policy requiring larger inventories. However in sectors where consumer confidence is on the availability of a resource and not an entity, then focusing on system-resilience through local adaptation would be a beneficial and cost effective policy.

References

1. K. Arrow, B. Bolin, R. Costanza, P. Dasgupta, C. Folke, C.S. Holling, B.-O. Jansson, S. Levin, K.-G. Mäler, C. Perrings, D. Pimentel. Economic growth, carrying capacity, and the environment. Science, 268 (1995), pp. 520–521.
2. Beyeler, W., Glass, R., Finley, P., Brown, T., Norton, M., Bauer, M., Mitchell, M., Hobbs, J.,. Modeling Systems of Interacting Specialists. 8th International Conference on Complex Systems, June 2011.
3. US Department of Homeland Security (2014) What is Critical Infrastructure? [online] http://www.dhs.gov/what-critical-infrastructure (accessed 9 May 2014).
4. Partha Priya Datta , Martin Christopher & Peter Allen (2007) Agent-based modelling of complex production/distribution systems to improve resilience, International Journal of Logistics Research and Applications: A Leading Journal of Supply Chain Management, 10:3,187-203
5. Jennings, Nicholas R. On agent-based software engineering, Artificial Intelligence, Volme 117 (2)(2000), pp.277-299.
6. T. Sandholm, Distributed rational decision making, G. Weiss (Ed.), Multi-Agent Systems, MIT Press, Cambridge, MA (1999), pp. 201–258
7. Brown, T., Beyeler, W. and Barton, D. (2004) Assessing infrastructure interdependencies: the challenge of risk analysis for complex adaptive systems. International Journal of Critical Infrastructures, 1 (1), 108–117.
8. Obama, B. (2013) Presidential Policy Directive 21 [online] http://www.whitehouse.gov/the-press-office/2013/02/12/presidential-policy-directive-critical-infrastructure-security-and-resil (accessed 2 June 2014).
9. E. Vugrin, M. Baca, M. Mitchell, K Stamber. Evaluating the effect of resource constraints on resilience of bulk power system with an electric power restoration model. Int. J. System of Systems Engineering, Vol. 5, No. 1, 2014.
10. Christopher, M., Peck, H., Building the Resilient Supply Chain. International Journal of Logistics Management, The. Vol. 15 No. 2, pp.1 – 14.

11. Simangunsong, E., Hendry, L.C., and Stevenson, M. Supply-chain uncertainty: a review and theoretical foundation for future research. International Journal of Production Research. Vol. 50, No. 16, 2012.
12. Brown, C., Lall, U., Water and economic development: The role of variability and a framework for resilience. A United Nations Sustainable Development Journal. Vol. 30, No. 4, pp. 306-317. 2006.