A simple model for the reliability of an infrastructure system controlled by agents

B. A. Carreras
BACV Solutions, Oak Ridge, TN 37831
bacv@comcast.net

D. E. Newman
Physics Department University of Alaska, Fairbanks, AK 99775
ffden@uaf.edu

I. Dobson
ECE Department, University of Wisconsin, Madison, WI 53706
dobson@engr.wisc.edu

Matthew Zeidenberg
Teachers College, Columbia University, New York, NY 10027
zeidenberg@tc.columbia.edu

Abstract

A simple dynamic model of agent operation of an infrastructure system is presented. This system evolves over a long time scale by a daily increase in consumer demand that raises the overall load on the system and an engineering response to failures that involves upgrading of the components. The system is controlled by adjusting the upgrading rate of the components and the replacement time of the components. Two agents operate the system. Their behavior is characterized by their risk-averse and risk-taking attitudes while operating the system, their response to large events, and the effect of learning time on adapting to new conditions. A risk-averse operation causes a reduction in the frequency of failures and in the number of failures per unit time. However, risk aversion brings an increase in the probability of extreme events.

1. Introduction

Infrastructure systems suffer from rare non-periodic large-scale breakdowns that lead to large economic and other losses by the community and sometimes can cause personal injuries and even loss of lives. The initiating causes of these events are very diverse, ranging from incidents caused by weather, to malfunction of system components, and to willful acts. Whether intentional or not, these events can threaten national security. Some examples of such extreme events are the August 14th, 2003 blackout in Northeastern America, the consequences of the Katrina hurricane in the New Orleans area etc.

The real cause of this behavior is that infrastructures are pushed to their capability limits by the continuously rising demands by consumers and the economical constraints upon them. The infrastructures are normally operating close to a critical point where events of all sizes are possible.

Present infrastructure systems are complex technological systems for which such extreme events are "normal accidents" [1]. As Perrow indicates, such normal accidents are characteristic of these systems and it is not possible to eliminate them. Additionally, these extreme events tend to generate a risk-averse attitude in the people managing and operating infrastructure systems. This change in attitude in turn, modifies the probabilities of occurrence of such events.

Some negative consequences of risk-averse operation on complex systems have been explored by Bhatt et al. [2], who used a model for the propagation of the failures inspired by how forest fires spread [3]. Altmann et al. [4], using a model of human reactions to river floods, have shown that the commonly-employed method of fighting extreme events by changing protection barriers in reaction to them is generally less efficient than the use of constant barriers to contain them. In this paper, we use a simple model of infrastructures, to further explore some of the consequences of such changes in operational attitudes to extreme events. We explore a range of behavior by system operators, varying from risk-taking to risk-averse operation. Initial results of our studies were published in [5].

In studying infrastructures, we can consider them as static systems with external forcing or as dynamical systems. As static system, that is with an external forcing at a given fixed time, they can be tuned to be closer or farther from a critical point where they run into the limit of their operational limits. As a dynamical system under the constant pressure of increase demands from the consumers and economical constraints, they are constantly pushed, in a self-consistent manner, against this critical point.

Although treating the infrastructure as a static system lacks realism, it has the advantage of allowing us to better understand how the system performs as it gets close to this critical point and also how policies that may be reasonable away from critical point can become dangerous as we get closer to the operational limits.

978-0-7695-3450-3/09 $25.00 © 2009 IEEE

We use a very simple model of the cascading process that can lead to extreme events. It is based on the CASCADE model [6, 7, 8]. The CASCADE model is a probabilistic model of load-dependent cascading failure, which was inspired by fiber bundle models [9, 10]. We assume a system with many components, which is under stress. This stress creates a certain distribution of loads among the components. When a component fails, load is transferred to some of the other components. However, in trying to model the behavior of infrastructure systems like the power grid, instead of using conservation of load on the components it was found to be more appropriate to use a constant load transfer which leads to a power law distribution of the size of the events. Here we measure the size of an event by the number of failed components.

Here, we extend the CASCADE model to include a dynamic component that evolves in a manner analogous to the OPA model [11], which we use to model the behavior of the evolving power grid. However, in CASCADE, this evolution does not push the system towards the critical point as occurs in OPA. The position of the system relative to the critical point is instead given by a parameter of the model.

In the generalized CASCADE model, time evolution occurs at two resolutions, or scales. Over a short time scale, which can be taken to be on the order of minutes, small and large cascades can develop. The long time scale corresponds to the evolution of the system over years. The dynamic evolution over the long time scale is governed by a daily increase in consumer demand that raises the overall load on the system and a concurrent engineering response to failures that involves component replacement and upgrades. These components may fail by either overload or random failure. They have a characteristic lifetime; therefore, timely replacement of the components can decrease the probability of failures.

In such a dynamical model the system reaches a steady state in which failures of all sizes may occur at different times.

The system is governed by two parameters, the upgrade rate of the components, μ, and the replacement time of the components, τ_R. Two agents control these parameters and learn how to operate the system in order to maximize a given utility function. Each agent's utility function has a well-defined economic term, which is the profit of selling the services minus the expenses of maintaining and upgrading the system. Each function also includes a penalty for failures, which is not only monetary and in general is proportional to the failure size.

It is through the penalty for failures that we can introduce differing attitudes on the part of the agents. Using basic ideas from prospect theory [12, 13], we can color this penalty and make the attitude of the agents risk-averse or risk-taking. By changing the agents attitudes in such a manner, we can evaluate the impact of the changes on the operation of the system. In particular, we are interested on how the probability of extreme events is affected by such changes in the attitude of the agents.

The rest of the paper is organized as follows. In Section 2, we describe the basic ideas behind the CASCADE model and its extension into a dynamic model. In Section 3, we discuss the parameters controlling the operation of the system and the agents, which are responsible for the decision-making process leading to their determination. The numerical results of the model are presented in Sections 4 and 5. In the former, we discuss the operation of the system with a fixed attitude on part of the agents. In the latter, we discuss the effect of changing the attitude of the agents in response to extreme events. Finally, the conclusions are given in Section 6.

2. A simple model for the infrastructure system

Let us consider a system with N components. In the spirit of the Cascade model [6], we assume that each component i in the system has a load L_i. The loads are distributed uniformly between a minimum load, L_{min}, and a maximum load, L_{max}. The Cascade model assumes that: 1) components fail when their load exceeds a prescribed value L^{fail}; and 2) when a component fails, a fixed amount of load P is transferred to $k < N$ other components [7, 8]. In this new version of the model, the components are also characterized by an average lifetime T_f. This allows for another mechanism for failure. We allow the components to age due to operational stress, which can lead to failure.

The survival function of a component when aging is taken into account is given by [14]

$$S(T) = K^{\frac{T^{\beta+1}}{T_f^{\beta}(1+\beta)}} = \exp\left[-\frac{1}{1+\beta}\left(\frac{T}{T_f}\right)^{\beta+1}\right] \qquad (1)$$

where $\beta > -1$. In what follows, we only consider the case of $\beta = 1$.

Once one or more components fail, we apply the rules of the CASCADE model described above, so a constant load P is transferred to k components. If any of these components then has as a result a load greater than L^{fail}, this component fails in its turn. Thus, the failure of one or more components can possibly lead to a cascading process. The process continues until no more components fail. In what follows we take $L^{fail} = L_{max}$.

To reduce failures, components are replaced periodically. The replacement time of the components is τ_R. The ratio of this replacement time to the failure time T_f. affects the probability of starting a cascade.

In the CASCADE model, all loads can be normalized in the following way:

$$l_i = \frac{L_i - L_{\min}}{L_{\max} - L_{\min}}, \quad p = \frac{P}{L_{\max} - L_{\min}} \quad (2)$$

After such a normalization, we are left with two independent parameters, the normalized load transfer p, and the ratio of the replacement time to the failure time, τ_R / T_f.

The mean size of the cascades increases with the value of p. There is a critical value of p, $p_c = 1/k$, after which the cascade size increases sharply. This is shown in Fig.1, where we have plotted the mean cascade size as a function of p. The figure is the result of calculations for a system with N = 1000, k = 100, and τ_R/T_f = 0.008, so the frequency of cascades in the subcritical regime is low.

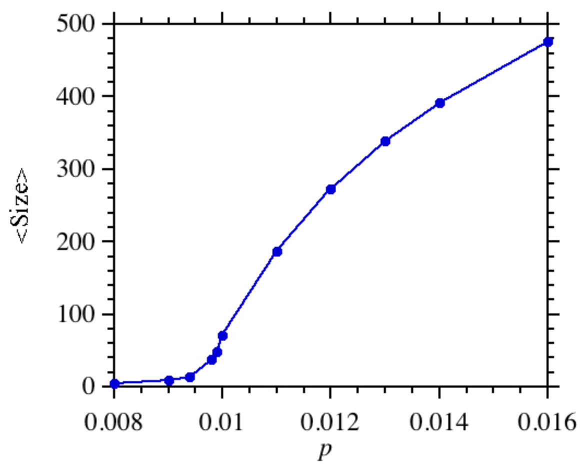

Fig.1 The mean cascade size as a function of *p*

One interesting aspect of the Cascade model is that the distribution of the cascade sizes can be calculated analytically [6]. At the critical point, the distribution has a power tail with a decaying exponent equal to 1.5.

As explained, we have generalized the cascade model by introducing a process involving two time scales. On the fast time scale, which we call minutes, we follow the cascading process as described above. Then, we complete the model by introducing dynamics for the long time scale. On this time scale, the basic unit of time is a day. In modeling the long-term dynamics, we closely follow the OPA model [11], which we have extensively used to study the dynamic behavior of the power grid.

The longer time evolution goes as follows:

1) We introduce an averaged load for the whole system, $L_A(t) = L_0 \exp(\gamma t)$, which is a function of time with a constant rate of growth γ, which represents the rate of increase of the demand. At the beginning of each day, the loads of the components are uniformly distributed around this mean value within a range ΔL_A.

2) At the beginning of each day all components are tested for failure, where their probabilities of failure are distributed as in Eq. (1). If there are any failures, a cascade process starts as described above.

3) At the end of the day all components that are due for replacement or have failed are replaced.

4) After a cascade of a minimum size (greater than 10% of the system size), the maximum load of all components is increased by a factor $\mu > 1$. This represents an upgrade of the system in response to failures.

As in the case of the OPA model [11], this system evolves to a steady state in which failures and replacements are in equilibrium in a time-averaged sense.

3. Basic parameters and agent operation

The operation of the system described in the previous section is controlled by four parameters. One is the rate of increase of the demand, γ, which is taken as a constant. We use γ = 1.00005, which is approximately the average daily rate of increase of the electric power demand in the USA over the last two decades. The second parameter is the load transfer p. This is a parameter that depends on the nature of the infrastructure. With a more detailed model of the power transmission grid infrastructure, such as OPA, p corresponds to the redistribution of power flow. I that model we find that the system dynamically organizes itself to sit right at, or just below, a critical point. Here, we fix its value to be close to but below the critical value, p_c. The motivation comes from analysis of the power systems [15, 16] that shows a probability distribution of blackout sizes with a power tail having a decay index that is somewhat larger than the critical value for the CASCADE model.

The two remaining parameters are the upgrade rate of the components, μ, and the replacement time of the components, τ_R. We use them as the actual control parameters for the operation of the system. There are two independent agents, each controlling one of these parameters, and each trying to optimize a utility function involving that parameter. Agent 1 controls the replacement time of the components, τ_R, and Agent 2 controls the rate of upgrade.

The utility functions for both agents are

$$U_1 = W_p(N - N_F) - W_R N_R - W_{F1} N \Pi\left(q_1, N_F/N\right) \quad (3)$$

$$U_2 = W_p(N - N_F) - W_{U1}(\mu - 1) - W_{U2}(\mu - 1)^2 \Gamma - W_{F2} N \Pi\left(q_2, N_F/N\right) \quad (4)$$

Here, W_p is the price the consumer pays for the "electricity" and is of course the same for both agents. For the first agent, W_R is the cost of replacing a component, and N_R is the number of components replaced by maintenance.

For the second agent, the cost of the upgrade has two components: first, the second term in Eq. (4), which is independent of time and represents the amortization of upgrade cost over time and second, the third term in Eq. (4), which is the cost when an upgrade occurs after a cascade which represents the immediate implementation cost of the upgrade. The multiplier Γ is 1 when there is a cascade and 0 otherwise.

In the last term for both functions, W_{Fi} is the cost (economic, social and political) for agent i for having components fail. N_F is the number of failed components that day. The function Π reflects the perception of how bad a failure is. Therefore, this is the subjective contribution to the utility function. We model it using functions normally used in prospect theory [12, 13]. That is

$$\Pi(q, x) = \frac{x^q}{x^q + (1 - x)^q} \quad (5)$$

In principle, it is possible that each agent has a different perception of the risk. Different values of q correspond to different attitudes of the operator regarding risk: risk aversion ($q < 1$), which heavily penalizes the most frequent events, or risk tolerance, which minimizes the cost of the most frequent events ($q>1$). Examples of these functions for three values of q are shown in Fig. 2. We also introduce a "dynamic" variation of the model which allows q to vary with time, making the agents more risk averse immediately after a large blackout followed by a decaying aversion; this variation is described in Section V.

The operators are designed to make a decision after a given period —we have varied the period between 30 and 100 days—regarding the maintenance schedule and the upgrade amount. This choice allows us to average the utility function over a reasonable number of days and avoid its daily high volatility.

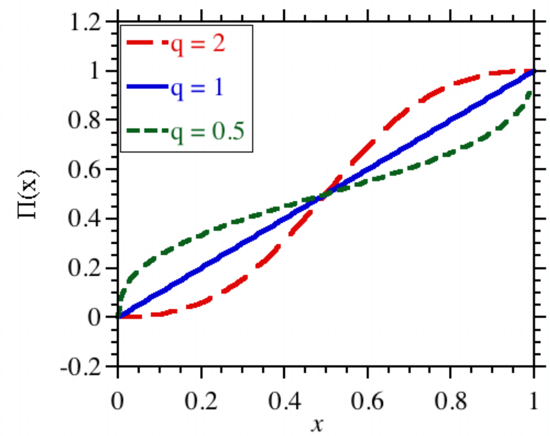

Fig. 2 The function Π characterizes the perception of how bad a failure is as a function of the size (which is roughly the inverse of its frequency) for 4 cases ranging from risk averse q = 0.5 to risk taking q = 4

At the end of each of each period, each operator decides which is the best value to use for replacement time and rate of upgrade for the next period, making this decision by examining:

1) The monthly averaged utility of the last few periods

2) The best averaged utility in the past

The operator (agents) then tries to optimize this by randomly varying the replacement time and the upgrade increment of the system a fixed amount positively or negatively from its best performance in order to find a possible improvement. Although the agents do not communicate among themselves, they clearly interact implicitly. If one of them changes a basic parameter, the other responds by adapting to the new situation. An example of the evolution of the parameters over time is shown in Fig. 3.

4. Impacts of risk-averse policies on the reliability of infrastructures when they operate near a critical point.

Here, we will consider the effect of risk-averse/risk-taking measures for different situations of the infrastructure system in relation to the critical point. The parameter p is the measure of the position of the system relative to the critical point as illustrated in Fig. 1. For the particular case considered here the critical point is at $p = 0.01$.

We have carried out numerical studies using multiple calculations for each set of parameters. The agents do most of their learning at the beginning of the calculation, and this phase of operation is strongly influenced by the random events occurring at those times. Therefore, longer calculations do not help much in improving statistics; it is better to do multiple calculations. In what follows, we use 10 runs for each set of parameters.

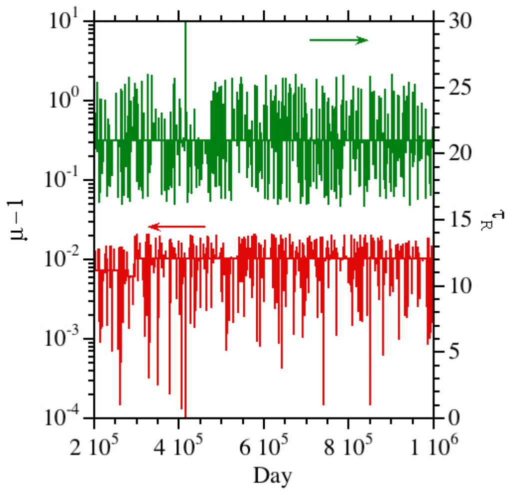

Fig. 3 An example of the time evolution of the agent controlled parameters (τ_R and •) for a typical case

We have carried out calculations for several values of q in the interval $0.5 < q < 2$, ranging from very risk-averse to strongly risk-taking operation. We have also varied the position of the system relative to the critical point by using three different values of p. For each value of q and p, the calculations have been carried out for 10^6 days.

The first agent, the agent controlling the replacement time, reduces the time between replacements as q decreases and moves from risk-taking to risk-averse behavior. The replacement time is reduced more than an order of magnitude as the agent moves from operation with q = 2 to operation with q = 0.5 as shown in Fig. 4. The agent's choice of τ_R is practically independent of the value of p, that is, on how close the system is to critical operation and is perhaps what one would intuitively expect.

In Fig 5, we show the corresponding plot for the choice of the parameter μ. In going from normal operation, q = 1, to risk-averse operation, q = 0.5, the second agent increases the rate of upgrade. This change of parameters is also consistent with the movement toward a risk-averse attitude. It is less clear why we observe the change in optimized behavior of the second agent in going from normal operation to risk-taking operation (q = 2). There is some increase in the rate of upgrade, which occurs in part to compensate for the increased number of failures due to the large replacement times chosen by the other agent. However, the values chosen are sensitive to the proximity to the critical point. Away from the critical, there are no large failures and the agent chooses the cheapest solution, a very low rate of upgrade,

The consequences of the agents' choices of operational parameters are very significant. In going from risk-taking operation to risk-averse operation, the frequency of the cascades is strongly reduced, as shown in Fig. 6. This reduction in frequency leads to a reduction in the number of components that fail per unit time. These results appear to fulfill the goals of a risk-averse operation by minimizing the failures and their occurrences. It should be noted that these improvements actually bear an increased cost; therefore the actual optimal values of the utility functions are reduced as q decreases.

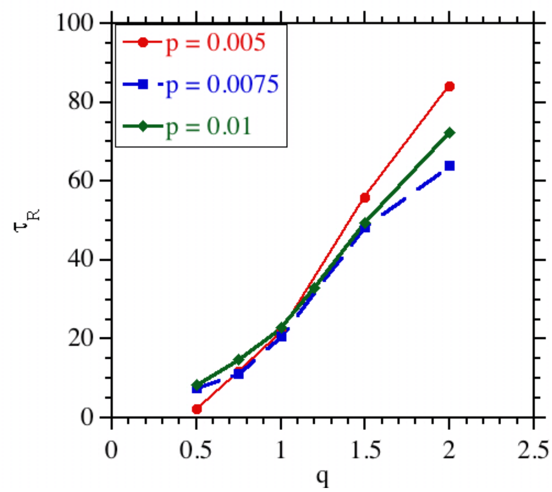

Fig. 4 The equilibrium control parameter τ_R as a function of q, the risk comfort parameter, from strongly risk averse (q=0.5), to strongly risk taking (q = 2).

Despite the increased cost, these results seem to endorse a risk-averse operation of the system, again consistent with what one might intuitively expect. However, there is a problem. Although the averaged number of failing components per unit time is reduced, the number of components per event is not always reduced. This is shown in Fig. 7 where the averaged number of components per event, event size, is plotted as a function of q.

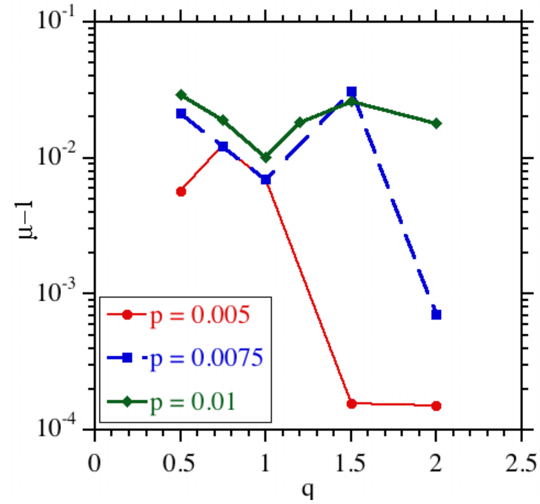

Fig. 5 The equilibrium control parameter μ as a function of q, the risk comfort parameter, from strongly risk averse (q=0.5), to strongly risk taking (q = 2).

In Fig. 7, we see a very distinct behavior when the system is below the critical point or when is just critical. In the first case (p=0.005), when the agent changes from risk-taking to risk-averse there is not only a reduction on the number of failures as seen above, but there is also a reduction on the size of those failures. However, near the critical point this situation reverses and the more risk averse agents become the larger the size of the individual failures become.

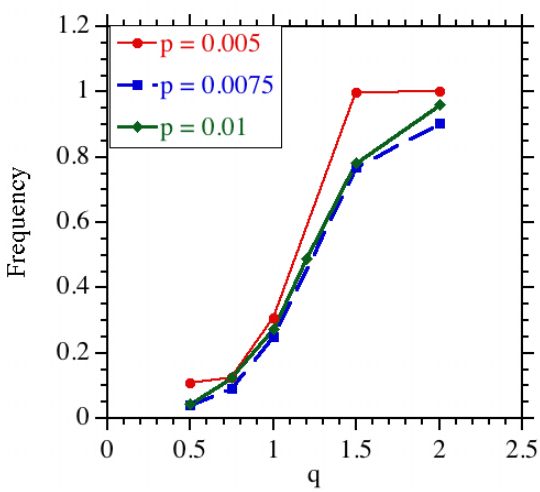

Fig. 6 The frequency of failures increase as the operation moves from risk averse (q=0.5) to risk taking (q=2).

When the system is at the critical point, the probability of an extreme event actually increases when going farther towards a risk-averse attitude. Here, we characterize an extreme event as an event in which more than half of the components of the system fail. In our calculation the system size is 1000: therefore an extreme event corresponds to 500 failing components.

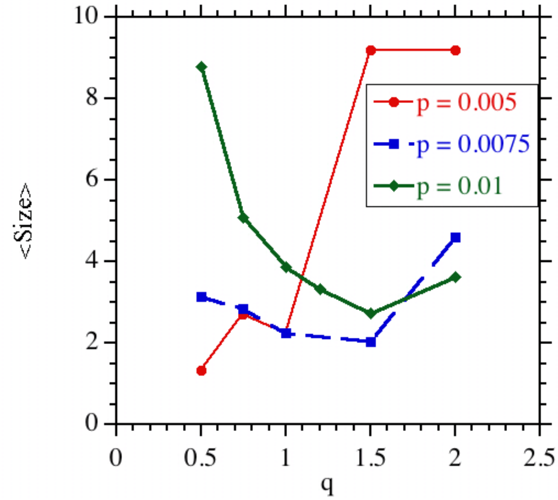

Fig. 7 Averaged number of components failing per event, Size, as a function of q.

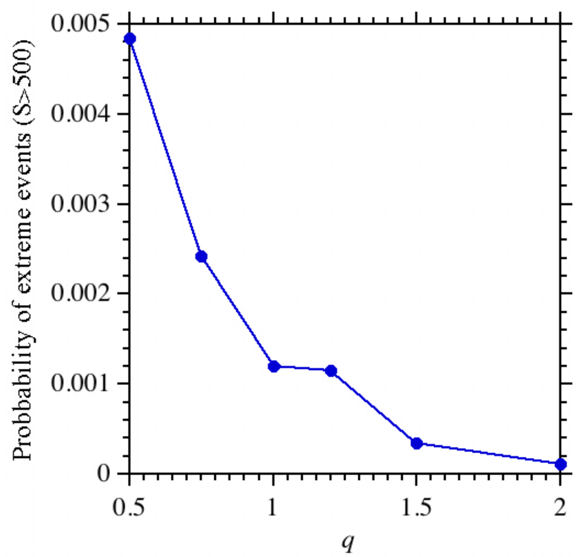

Fig. 8 Probability of large events *decreases* as operation moves from risk averse (q=0.5) to risk taking (q=2).

To quantify this we can evaluate the probability of an event in which more than 500 components fail for each of the values of q in the scan. The result is plotted in Fig. 8. In this we can see a systematic increase in the probability as q decreases. This change in the probability of an extreme event is a result of the change in the probability distribution function of the event size. Once again, we measure the event size by the number of components that fail during a single cascading process. In Fig. 8, we have only plotted the case p = 0.01 because in subcritical operation there are no extreme events.

This can be better understood by looking at the full distribution function of the failure size for the different events below the critical threshold, $p = 0.005$, as shown in Fig. 9a contrasted to the same functions at the critical point, $p = 0.01$, as shown in Fig. 9b

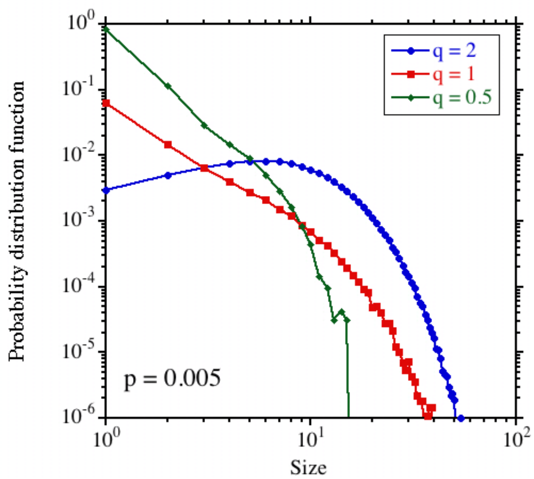

Fig. 9a PDF of failure sizes for subcritical operation.

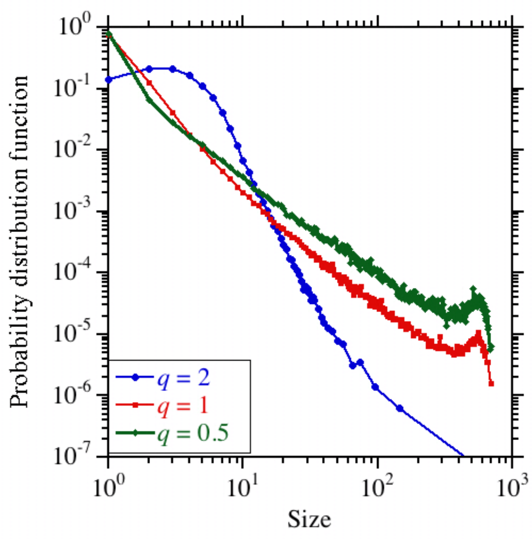

Fig. 9b PDF of failure sizes for supercritical operation.

In going from risk-averse to risk-taking, we see in Fig. 9a that the proportion of large events increases while the small event decrease. This is consistent with the change of the average size. In all 3 of these cases, the pdfs have sharp cut-offs for large events and there are no event greater than 8% of the system size.

However, at the critical point the situation is very different. In these cases, there are events of all sizes with algebraic tails characterizing the pdfs. In this situation we have the opposite result. For risk averse operations the number of system size events increases relative to risk-taking operation. Since the social and economic impact of system size failures are usually very high, the result of the risk averse attitude is, counter-intuitively, the opposite of what is intended.

These changes in the probability distribution function (PDF) are clearly shown in Fig. 9b where we have plotted the corresponding PDF for three different values of q, 2, 1, and 0.5. In this figure, we see that the power tail for q = 2 has disappeared and the PDF has a different functional form from the other two operational regimes. In going from normal operation to risk-averse operation the decay index of the power tail goes from 1.8 to 1.45. Clearly, under risk-averse operation, the system appears to be operating in a more critical state.

In moving from risk-taking to risk-averse operation not only the distribution of failures changes, but also the dynamics of the failures. The clearest manifestation of this is in the change in the long-range time correlations. The risk-taking regime leads to persistence over long time scales while the risk-averse regime leads to antipersistence. The Hurst exponent as operation moves from q = 2 to 0.5 changes from 0.6 to 0.3, as shown in Fig. 10. This is consistent with the idea that the operator in a risk-averse regime is trying to avoid all failures.

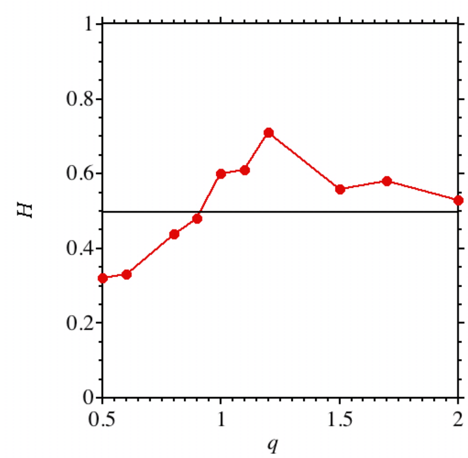

Fig. 10 The Hurst exponent (H) goes from persistent (H > 0.5) to anti-persistent (H < 0.5) as operation is moved from risk taking (q=2) through normal to risk averse (q=0.5).

5. Risk-Averse Operation as a Reaction to an Extreme Event.

Another possible way in which risk aversion may affect the operation of infrastructure systems is when a risk-averse policy is the reaction to some extreme event. In what follows, we incorporate such a possibility into our generalized cascade model by making the exponent q a function of time so the system can react to the extreme event and then forget about it over time..

We use the following formulation of the problem. We start with q = 1.0. If an event of size S, greater than a given threshold S_{th}, occurs at time t_0, then q falls to q_0=0.5, and then evolves in the following way:

$$q(t) = 1 + (q_0 - 1)\exp\left(-(t - t_0)/T_0\right) \qquad (6)$$

We choose q_0 =0.5 because we want to simulate a risk-averse reaction to the large event. The parameter T_0 is the time after which the system (agent) has "forgotten" the event. In the present calculations we take T_0 = 500 days. Later, we will discuss the effect of changing this characteristic memory decay time.

We have carried out a scan over possible values of the event threshold value S_{th}. As before, for each threshold value, we do 10 calculations in order to accumulate reasonable statistics over the various events and agent behavior responses to those events.

The choice of the operational parameter for each different threshold value is consistent with the risk-averse attitude of the agents, as we have seen in the previous section. This is shown in Fig. 11, where we plot the average value of μ and τ_R as a function of the event threshold. Here, the averages are taken over the full length of the simulations, that is 10^6 days, and over the 10 different simulations for each set of parameters.

As the threshold for the event triggering the risk-averse reaction decreases, the system operates in a risk-averse mode of operation for a longer time. Therefore decreasing the threshold lengthens risk-averse operation relative to normal operation. As the risk-averse operation gets longer, we see a reduction in the replacement time and an increase in the rate of upgrade.

Note that when the threshold is above 750 there is no effective threshold at all because there are no events greater than this size. Therefore all such thresholds are equivalent and they serve as the reference case for normal operation.

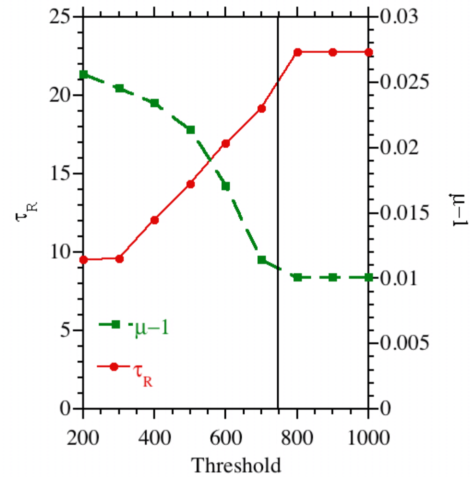

Fig. 11 The average value of μ and τ_R as a function of the "learning" event threshold for the system that forgets about events after a given decay time.

What are the consequences of the change in behavior of the agents? As they operate in a risk-averse mode for longer periods of time, corresponding to lower thresholds, there is a clear reduction in the frequency of the cascades, as we have seen in the previous section. Fig. 12 illustrates that point: as the value of the threshold is reduced, the frequency decreases.

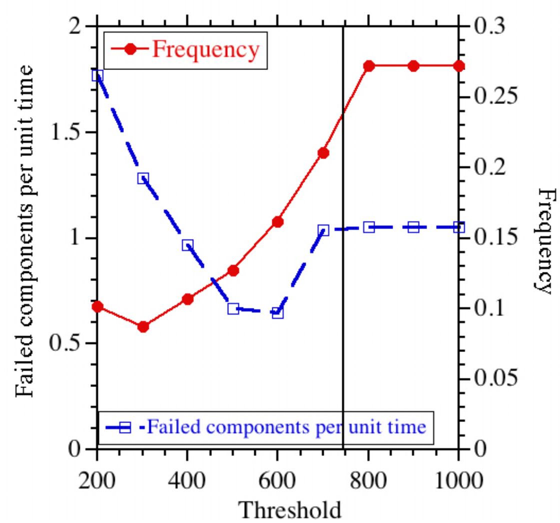

Fig. 12 Frequency of failures and number of failed components and a function of the threshold for risk-averse behavior. This plot suggests there may be an optimal threshold for such behavior which minimizes both the frequency and size of the failures.

The situation is different with regard to the number of failures per unit time, which we have also plotted in Fig. 12. We can see that for large thresholds, for which only a few events trigger the risk-averse attitude, there is a decrease of failures per unit time because of the decrease in frequency of the events. However, as the risk-averse attitude is triggered more often, that is, for smaller thresholds, the increase in the average size of the events compensates for the decrease in frequency and eventually the number of failures per unit time is higher than in the reference case. As we saw in the previous section, risk-averse operation increases the probability of extreme events and this causes an increase in the number of failures if a risk-averse attitude is triggered too often. This suggests that there might be an optimal threshold for the onset of risk-averse operation. In the cases shown here, that optimal threshold would appear to be between 500 and 600.

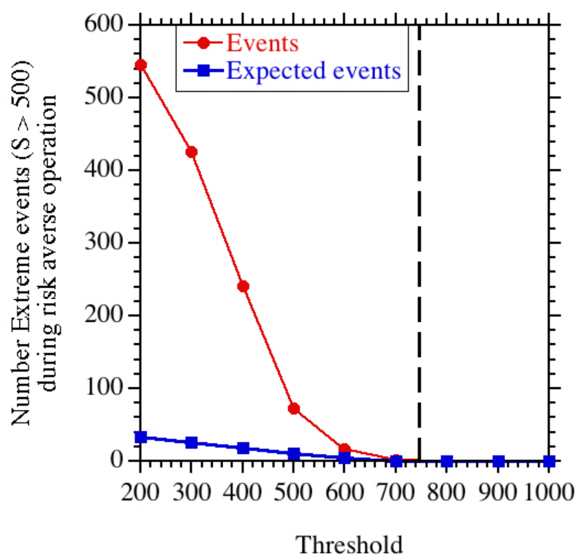

Fig. 13 Number of extreme events (larger then 500) during risk averse operation compared to expected events if no risk averse operation.

We want to examine the effect of risk-averse operation on extreme events in more detail. To do so, we analyze separately the period of time during which the agents are risk-averse, distinct from the period when they are not risk-averse. Using the reference case, we can calculate the probability of an event larger than 500. From this probability, we calculate the expected number of extreme events in the two phases of operation for the different thresholds. These results are compared to the real number of extreme events obtained with those calculations. The results are plotted and compared in Figs. 13 and 14. We can see that in general risk-averse operation leads to a higher number of extreme events during both phases, the risk-averse phase and the one following afterward. From these figures, we see that the increase is clearly more dramatic during the risk-averse phase.

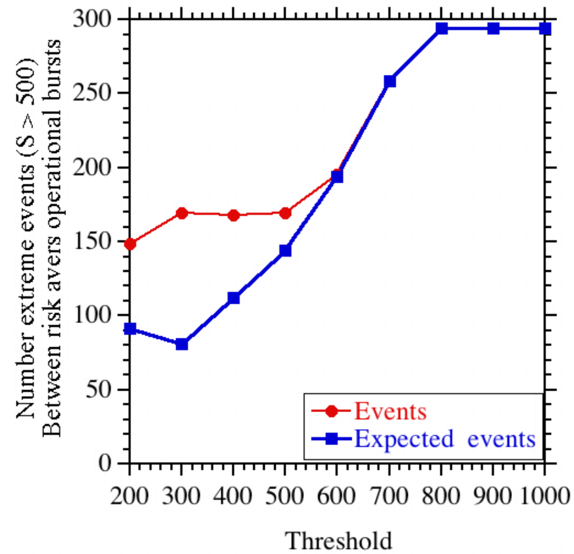

Fig. 14 Number of extreme events (larger then 500) between phases of risk averse operation compared to expected events if no risk averse operation.

6. Conclusions

In this paper, we presented a simple model of the operation of an infrastructure system. This model is a generalization of the CASCADE model, which has been studied in detail in the past. This generalization transforms the CASCADE model from a probabilistic model to a dynamic model. However, its dynamics is limited because the proximity to the operational critical point is not controlled by the dynamics.

The dynamic evolution of the system over a long time scale is governed by a daily increase in consumer demand that raises the overall load on the system and an engineering response to failures that involves upgrading of the components. Two parameters control the operation of the system, the upgrading rate of the components, μ, and the replacement time of the components, τ_R. Two agents operate the system by each selecting one of those parameters.

The utility functions used by the agents to optimize performance incorporate some perceptions of the events that affect the decision-making of the agents. The agents are characterized by three aspects of their (social) behavior:

1. Their risk-averse and risk-taking attitudes while operating the system.

2. Their response to large events, which can trigger a change in their behavior

3. The effect of learning time on adapting to new conditions. These aspects were considered in detail in [5].

These three agent behaviors affect the performance of the infrastructure system.

In going from risk-taking to risk-averse operation there is a reduction in the frequency of failures and in the number of failures per unit time. However, risk aversion brings an increase in the probability of extreme events when the system operates close to its critical point. During risk-averse operation close to critical, the PDF falls off with a smaller exponent than that found in normal operation. This is in general a very unwelcome change because large events are of much higher cost.

When risk-averse operation is triggered in response to extreme events, we obtain similar results as we find in the case of continuous risk-averse operation, but the probability of extreme events can be even higher than in the continuous operation if this reaction is triggered too often, that is if the threshold for entering risk-averse operation is relatively low.

We gratefully acknowledge support in part from NSF grants SES-0623985 and SES-0624361.

7. References

[1] Charles Perrow, Normal Accidents, Princeton University Press, 1999.
[2] U. S. Bhatt, D. E. Newman, B. A. Carreras, and I. Dobson, 38th Hawaii International Conference on System Sciences, Hawaii, January 2005.
[3] B. Drossel, and F. Schwabl, Phys. Rev. Lett., 69, 1629 (1992).
[4] E. G. Altmann, S. Hallerberg, and H. Kantz, Physca A 364, 435 (2006).
[5] B. A. Carreras, D. E. Newman, Ian Dobson, and Matthew Zeidenberg, submitted to Phys. Rev. E.
[6] I. Dobson, B. A. Carreras, D. E. Newman, Probability in the Engineering and Informational Sciences, 19, (1) 15-22 (2005).
[7] I. Dobson, B. A. Carreras, B. Nkei, D. E. Newman, Probability in the Engineering and Informational Sciences, vol 19, no 4, 475-488 (2005).
[8] I. Dobson, B.A. Carreras, D.E. Newman, Probabilistic load-dependent cascading failure with limited component interactions, Circuits and Systems, IEEE International Symposium on Circuits and System, Vancouver Canada, Volume: 5 (2004),
[9] F. T. Peirce, J. Text. Inst. 17, T355 (1926).
[10] H. E. Daniels, Proc. R. Soc. London, Ser. A 183, 405 (1945).
[11] B. A. Carreras, V. E. Lynch, I. Dobson, and D. E. Newman, Chaos, 14, 643 (2004).
[12] D. Kahneman and A. Tversky, Econometrica 47, 263 (1979).
[13] A. Tversky and D. Kahneman, J. Risk Uncertainty, 5, 297 (1992).
[14] T. Bedford and R. Cooke, Probabilistic risk analysis, foundations and methods, Cambridge University Press, 2001.
[15] B. A. Carreras, D. E. Newman, I. Dobson, and A. B. Pool, IEEE Transactions on Circuits and Systems, part I, 51, 1733 (2004).
[16] I. Dobson, B. A. Carreras, V. E. Lynch, D. E., and D. E. Newman, Chaos, vol. 17, 026103, June (2007)