From *Collaborative Financial Infrastructure Protection: Tools, Abstractions, and Middleware*, Baldoni R and Chockler G, Eds., Springer, January 2012

Chapter 2. Modeling and Risk Analysis of Information Sharing in the Financial Infrastructure

Authors: Walter E Beyeler and Robert J Glass (Sandia National Laboratories) Giorgia Lodi (Univ. of Roma)

Abstract: This chapter defines the community of banks as a Complex Adaptive System of Systems or CASoS and analyzes the value of information sharing as a general policy to protect the community against cyber attacks. We develop a model of interacting banks that have networks of business relations with a possible overlay network of shared information for cyber security. If a bank suffers a cyber attack it incurs losses and there is some probability that its infection will spread through the business network, imposing costs on its neighbors. Losses arising from financial system compromise continue until the problem is detected and remediated. The information sharing system allows detection events to be broadcast, and also increases the probability of detecting the experimental probes that might precede the actual attack. Shared information is a public good: one institution's agreeing to share information speeds responses at other institutions, reducing their probability of initial compromise. Information sharing participation carries with it costs which need to be balanced by direct expected gain or to be subsidized in order to have a critical number of banks to agree to share information and to discourage free riding. The analysis described in this chapter examines the incentives motivating banks to participate in information sharing, the benefits to the financial system that arise from their participation, and the ways banks' incentives might be shaped by policy to achieve a beneficial outcome for the system as a whole.

Contents:

1. The Community of Banks as a Complex Adaptive System of Systems

The community of banks is a quintessential Complex Adaptive System of Systems, or CASoS. Each institutional bank is composed of a number of components (e.g., branches, ATM machines) and acts as a system, composed of technology, policy and protocol, and manned by people who both make and implement decisions. In this way, adaptation occurs at multiple scales driven by the ability to provide for the people who are their customers. Banks are linked together into a community through the movement of funds from one to another (e.g. a payment system) and customers who can choose to move from bank to bank based on their view of the security of their funds and the cost of services. As such the community of banks constitutes an adaptive system of systems. While this adaptive system of systems is indeed complicated, more importantly, it is complex. It is composed of a number of similarly

behaving components that are linked together in similar ways, many of which implement binary decisions, for instance, do I keep my funds with this bank or move to another? And these binary decisions are often influenced by limited assessment of security and, frequently, reaction on the part of customers who may mimic the behavior of others. As such, emergent behavior ensues that is very difficult to predict except in stylized form and through underlying statistical analysis.

Designing solutions within a CASoS, or CASoS Engineering, is a new and growing discipline (Glass et al., 2008, 2011 and see http://www.sandia.gov/CasosEngineering/). Because many CASoS are high-consequence systems that contain people, CASoS Engineering requires that we first recognize the system for what it is and develop a model which reflects that system in a way that we can test design alternatives. Examples of the application of CASoS Engineering include financial payment system operation and cascading (Soramaki et al, 2007, Beyeler et al., 2007, Renault et al., 2008), the design of mitigation strategies for pandemic influenza (Davey et al, 2008, Perlroth et al., 2010) and analysis of the global energy system to foster multi scale security (Glass et al., 2008) to name a few. In this chapter, we begin our process at a very high level and consider the value of information sharing to the individual and community of banks in context of cyber attacks.

2. Cyber Security as a Common Good; Balancing the Risks and Rewards of Sharing Information

Designs, such as the CoMiFin system, counter the increasing threats that financial institutions face from cyber attacks. The perpetrators of these attacks, whether motivated by the prospect of financial gain or because they see such attacks as a means of garnering publicity or otherwise pursuing a political cause, benefit from sharing technology and other information among themselves. Banks can also use information sharing to help speed their detection of, and response to, such attacks. The benefits of having a robust, secure information sharing platform are clear, both for the system as a whole and for the individual financial institutions composing the system.

While we can stipulate that a financial system in which all participants share information related to cyber security will be more secure, with increased likelihood of attack detection, it is less clear that there is an incremental path from the current situation, in which no institutions share information or do so informally, to the more secure condition in which all institutions share information. The path to complete participation may be blocked by the particular incentives to join that each institution faces.

Agreeing to share information entails some cost to the participating institution: these include costs of acquiring and maintaining equipment, training staff to use it, and integrating it into existing business practices. Although these expenses may be small in relation to existing operating costs, non-monetary costs might have some influence on the decision to participate.

Incorporating any innovation in a functioning operation may be resisted on principle by inherently conservative institutions. This resistance would be especially important during the initial stages of adoption, when the favorable experience of peers is unavailable. Sharing information of any kind with competitors may be seen as a cost in that it might be exploited in some way for competitive advantage.

Designs, such as CoMiFin, incorporate many kinds of protection against this prospect, but anxiety on this account might still be a deterrent to adoption.

The benefits of information sharing include a decreased probability that a particular attack will be successful, and an increased rate of detection and recovery should an attack succeed. A large part of this benefit is naturally seen by the institutions in the system; however there may be important positive externalities as well. If an attack on a particular bank is successful, it may create problems for other banks in the system. If the attack introduces propagating malware, for example, it might be spread to other institutions through business or social communications. Operational disruptions at one institution might impose costs on other banks by preventing clearing and settlement of interbank transactions or customer transactions. There are also possible reputation costs to the system as a whole arising from a successful attack on a single institution. Such externalities create an incentive for each institution in the system to see other banks join in an information sharing arrangement. Conversely some of the benefits created by a joining institution are experienced by other institutions in the system, whether or not they themselves participate.

The information sharing benefit to the system as a whole will depend on the number of institutions choosing to participate. We might see large benefits from a small fraction, or it may be that nearly all banks must join before substantial collective returns are achieved.

3. Viewing the Problem as a Dynamic Multi-player Game

The interaction between a bank's decision to participate in information sharing, and the subsequent costs and benefits confronted by other banks as a consequence of this decision, suggests that game theory (Straffin 1993) would provide an effective lens on the problem of fostering beneficial systemic outcomes. The problem can be posed as a multi-player game in which the benefits to each bank of joining the information sharing system are a function of the decisions of all other banks in the system. This provides a good conceptual picture of the problem but is not a realistic way of solving it. First, the number of players and variety of inter-player interactions make it impractical to calculate the benefits to each player for each possible pattern of participation. Second, the astronomical complexity of this payoff matrix means that banks would not use it as a basis for making their decisions. Instead we treat the problem as an iterated game, in which banks repeatedly update their information about their environment, and assess their decision using this updated information.

At each step institutions evaluate the anticipated costs and benefits of joining, and move accordingly. These moves change the cost/reward structure for the other banks in the system, which then revisit *their* decisions on the next iteration. The game concludes when no bank has an incentive to change its decision. In some situations the game may not have an equilibrium endpoint of this kind: oscillatory solutions are possible, particularly if the calculation of costs and benefits involves lagged observations of the performance of the system.

4. Model Formulation

4.1 Purpose and Scope

The model is intended to provide an understanding of how the various factors that influence adoption of information sharing by individual banks can interact to influence the overall level of participation, to identify any obstacles to achieving a degree of participation leading to a globally beneficial condition, and to assess the effectiveness of policies designed to overcome those obstacles by shaping the incentives confronted by the participants in the financial system.

The model represents the way cyber attacks impose costs on the system, and on how a successful information sharing system can reduce those costs by increasing the chance that they will be detected and effectively blocked or mitigated. The particular nature of the attack and of the mitigation are not specified, but we assume the probability that an attack will be detected by the information sharing system is a function of the number of banks that participate in the system and that are targeted by the attack. This dependence reflects the essential benefit conferred by information sharing.

4.2 Limitations

The model is a highly stylized representation of the complex processes and relationships that compose a modern banking system. It is designed to focus on the most critical factors and processes controlling adoption of a protective technology having substantial externalities within a system having both cooperative and competitive interrelationships among its components. It might be refined in many details, for example attacks might be differentiated based on their expected cost and the efficacy of information sharing in detecting them; or the social networks interconnecting bank customers might be considered in modeling the evolution of customers' risk perceptions. However, not all refinements are likely to produce new insights regarding the question of concern: how adoption can be fostered by policy.

The model does not include adaptive moves by the adversary, which might be very important. Adoption of the information sharing platform might cause adversaries to adjust their mode of attack in order to evade its detection algorithms. The platform itself might become a target of attack either with the direct goal of causing operational problems at a number of banks, or as a preliminary step in a campaign directed at financial institutions themselves. Conversely the presence of an integrated system for information sharing might make defensive responses more adaptable, in that defenses against novel attack strategies can be implemented for all participants at once.

4.3 Model Elements and Governing Equations

Figure 1 illustrates the components of the system that we model and describes the basic processes we consider.

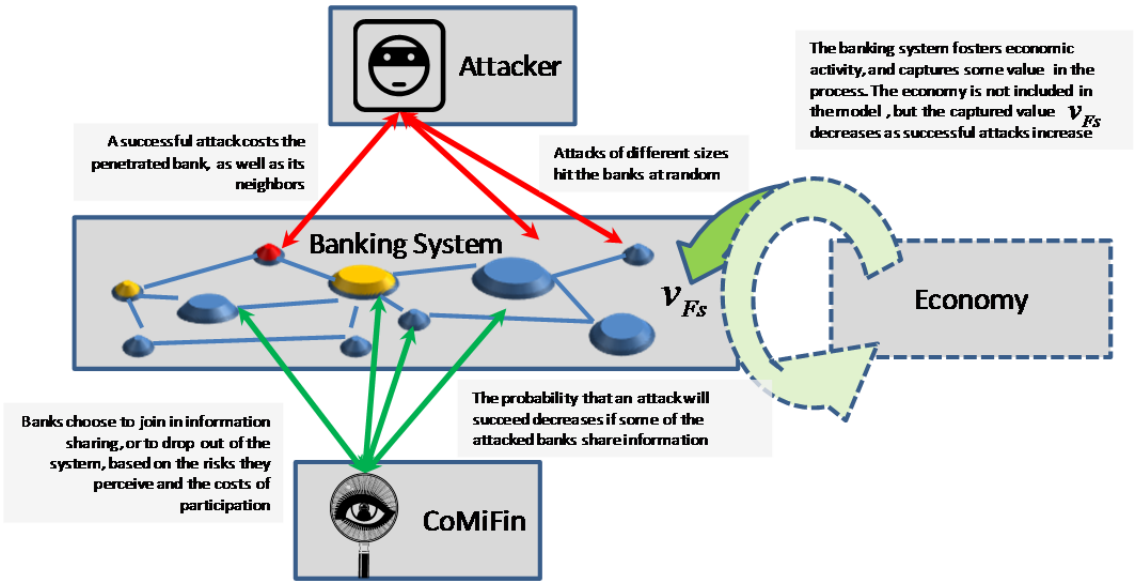

Figure 1 - Schematic of Model Components and Processes

The Banking System contains N interacting financial institutions b_i. Each bank has an overall size parameter s_i that influences its role in the system as described below. These institutions provide financial services in some economic environment, which is not represented directly in the model; however, the economic value captured by the banking system as a whole is represented as a net flow of v_{Fs} from the economy to the banks. This flow is relevant to the present problem because it may depend on the perceived risk that economic actors impute to the Banking System as a whole, arising from successful cyber attacks. We model this dependency as a logistical function of the current probability of successful attack at any institution $\hat{p}_a$:

$$v_{Fs} = \frac{v_{\max}}{1 + e^{(\hat{p}_a - p_c)s_a}} \tag{1}$$

The logistical function reflects the assumption that there is some critical level of attack probability p_c at which users begin to regard the system as risky, and to impose risk premiums in some way that degrades the economic value available to the financial system. At high levels of attack probability, the economic value is driven toward zero. Sensitivity of value to attack probability around the critical point is determined by s_a .

Banks compete for shares of the total value v_{Fs} , and the probability that a particular institution will be attacked is assumed to play some role in its success in the competition for market share. Successful attacks on the system will change risk perceptions, tending to shift business away from banks seen as riskier and towards those seen as less risky.

We model this process through a network of business interactions among banks. The links in this network represent the extent of interactions between pairs of banks. These interactions include

relationships between the banks themselves as well as business transactions between the sets of customers served by the banks. Strong business ties among financial institutions, represented by these links, have several implications in the model.

First, linked financial institutions are assumed to be potential competitors for each other's customers, with some part of the relative attractiveness of a bank related to the relative risk of attack. In part, links between banks reflect business relationships among their customers, and these relationships allow customers of different banks to compare information about the services and performance of their banks. Differences in risk across links are assumed to drive business from the more risky to the less risky bank to some extent. A link in the network between banks b_i and b_j represents business connections with a strength s_{ij}. The proportion of this business assigned to bank b_i is:

$$w_{ij} = \frac{s_{ij}}{1 + \hat{p}_i / \hat{p}_j} \tag{2}$$

where $\hat{p}_i$ is the current estimate of the probability of successful attack on bank b_i. These probability estimates, and the estimated probability of attack at any bank $\hat{p}_a$, are based on the history of successful attacks that the overall system experiences. These estimates are formed over a period τ_m which represents a memory time over which banks and their customers base their expectations about performance:

$$\hat{p}_i(t_{n+1}) = \hat{p}_i(t_n) + \frac{\Delta t}{\tau_m}\left(Ia_i - \hat{p}_i(t_n)\right) \tag{3}$$

where Ia_i is an indicator of successful attack against bank b_i during the period $t_n \rightarrow t_{n+1}$.

Attacks on the system are modeled as a random process that occurs at some specified frequency f_a. A particular attack is directed at some number n_a of the banks in the system. Some number of these banks, n_{sa}, will be participants in the information sharing system at the time of the attack. The probability that an attack will be detected by the information sharing system is assumed to increase as the number of attacked banks that share information increases, and to have a maximum detection probability of $p_{\max}$ that is achieved when all banks participate:

$$p_{\det} = p_{\max}\left(\frac{n_{sa} - 1}{n_a - 1}\right)^{1/\gamma} \tag{4}$$

The parameter γ describes the "power" of information sharing by controlling the fraction of participation needed to approach the maximum detection probability. A value of 1 defines a linear increase in probability with participation. Values larger than one cause a large increase in detection

probability with relatively small participation; values less than one require a large fractional participation before the maximum benefit is seen.

A successful attack is assumed to create some cost α at each of the banks attacked. In addition the attack imposes a cost β on neighboring banks in the financial network. These costs can be due to any of several mechanisms, such as operational interruptions or delays at the attacked bank that inconvenience the neighboring bank or its customers, or propagation of the direct effects of the attack through the business linkages between the banks. Both the cost to the successfully attacked bank and costs to its neighbors are relative to the bank sizes.

A bank's decision to participate in information sharing is based on the costs they have experienced as a consequence of attack in comparison to the various costs associated with participation in the information sharing system. The model includes three components of cost: initial capital costs associated with installing the system; recurring costs of maintaining the system; and perceived costs due to business risks that might arise in consequence of participation. These last costs include uncertainty about the impact of incorporating the new system in a bank's existing IT processes and concerns about leakage of business intelligence through the shared information. Although the CoMiFin system design includes extensive protection against such leakage, a lack of experience with the successful performance of this protection may create an important barrier to adoption.

Accumulated experience with the information sharing system, whether by the bank considering adoption or by other banks in the financial system, is assumed to provide public information about the operational impacts of using the system and to alleviate concerns regarding possible loss of sensitive information. This effect is modeled by assuming that the leakage cost that a bank considers when judging whether to participate decays as experience with information sharing accumulates within the financial system:

$$c_l = c_{l\max} e^{-h_s/(N\tau_m)} \tag{5}$$

where h_s is the accumulated experience, up to the current time t_n with the information sharing system by all financial institutions:

$$h_s = \sum_{k=0}^{n} n_s(t_i) \tag{6}$$

with $n_s(t_i)$ being the number of banks participating at time t_i .

4.4 Evolution of System State

The information sharing system evolves through a series of time steps. At each step, the economic value captured by the financial system as a whole, v_{Fs} , is first calculated using Equation (1) based on the system's recent experience with attacks, reflected in the current probability of attack $\hat{p}_a$. This total value

is then allocated among the individual institutions based on their total size, and on the shifts in business allocation due to contrasts in comparative risk of attack across pairs of banks, using Equation (2).

Next, an attack may be attempted based on the specified attack frequency, with the number of banks attacked randomly chosen between 1 and a specified maximum size. The attacked banks are chosen at random from the entire population. Some of the attacked banks may be participants in the information sharing system. The number of sharing banks attacked determines the probability that the attack will be detected and effectively blocked via Equation (4). If the attack is not blocked, it imposes costs on the attacked banks as well as on all neighbors of the attacked banks in the Banking System.

Each bank then updates its expected probability of attack, and the costs it incurs because of attack, using its experience during the step. The overall probability of attack on any bank, $\hat{p}_a$, which controls the rate of value capture by the Banking System, is also updated depending on the success of an attack at any bank. These probabilities retain information about experiences over a time period of τ_m. This represents a forgetting time for decision makers in the Banking System and the economy.

At the end of the time step, each bank evaluates its participation in the information sharing system. Its recent experience with attacks is used to calculate an expected cost per unit time arising from successful attacks. This cost determines whether the bank joins the information sharing system, if it currently doesn't participate, or drops out of the system if it does. Banks outside the sharing system decide to join if the expected attack cost is greater than the sum of the (amortized) fixed cost, maintenance cost, and cost of prospective loss of business information cost. Banks currently in the system continue to participate as long as the expected attack cost is greater than the maintenance cost and cost of information loss.

4.5 Output Metrics

Our primary interest is in identifying conditions that maximize benefit to the system. The direct measure of this benefit is the aggregate value generation rate v_{Fs}. This is a function of the frequency of successful attacks on the system, which in turn depends on the degree of participation in the information sharing system. The number of participants n_s is therefore an additional important metric. Participation responds to the costs confronted by banks, and so these costs are the natural targets of policies designed to foster information sharing system benefit.

We examine a financial network consisting of 100 banks. Each bank has links to a number of other banks in the system (uniformly distributed between 2 and 10), representing strong business relationships between the banks or their customers. The size of each bank is proportional to the number of its links to other banks. Each link is assumed to represent a unit of potential economic value created by the banking system. The scale for measuring costs and benefits in the information sharing system is set by the maximum economic value created by the system as a whole ($v_{\max}$ in Equation (1)), which is given by the total number of links between banks.

Using the nominal parameter values for the cost of attack and for the information sharing countermeasures, the total economic value captured by the banking system varies in a narrow band around a value of 245 (Figure 2). The model configuration gives a maximum limit of 300, so that roughly 15% of the potential value to the banking system is lost to the direct and indirect costs of cyber attack. Despite the motivation reflected by this loss, no banks join the information sharing system in this simulation due to the cost values assigned for installation, maintenance, and operational uncertainty.

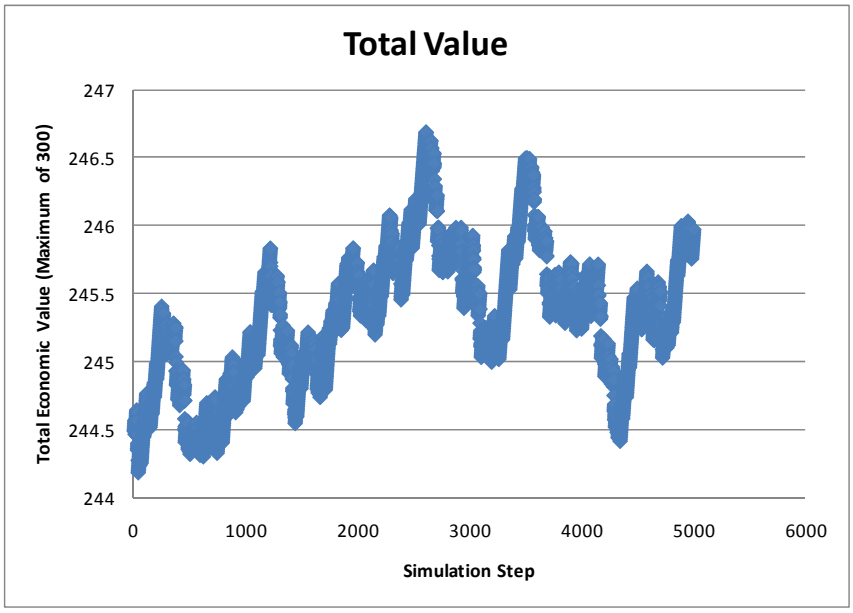

Figure 2 - Total Economic Value in a System with High Sharing Cost and No Participation

Reducing the information sharing system maintenance cost by an order of magnitude tips the decision in favor of adopting information sharing. Figure 3 shows the resulting trajectory of total economic value captured by the banking system, along with the value function calculated with high cost (Figure 2) for comparison. In this case the total value captured increases from the reference case by 35 units, or more than 10%. The total maintenance cost in the original simulation if *all* banks had participated would only have totaled 0.1 units. While these particular results are based on notional parameter values and are not necessarily reflective of the real system, they do illustrate the potential for small changes in *local* costs seen by individual banks to create large *systemic* benefits.

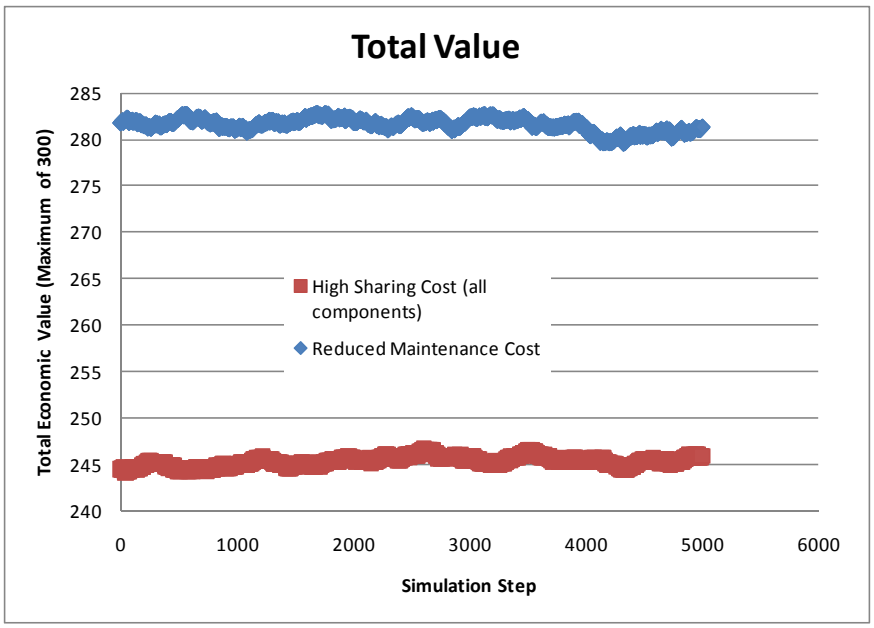

Figure 3 - Total Economic Value in a System with Lowered Maintenance Cost

The number of banks participating in the system has an interesting trajectory in this case, as shown in Figure 4. The number of banks does not reach a fixed plateau, but instead fluctuates around a participation level of approximately 80%.

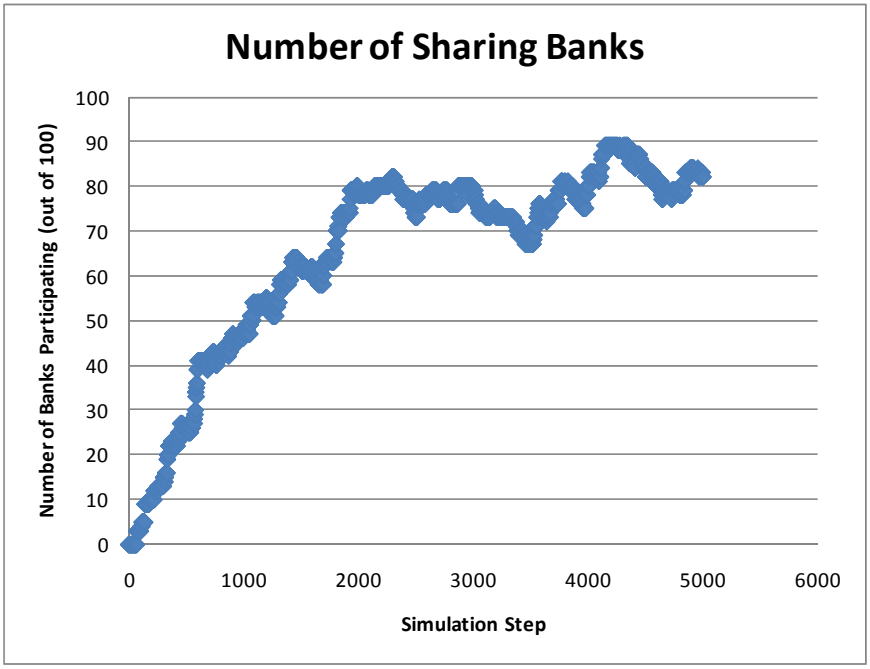

Figure 4 - Number of Banks Participating in the Information Sharing System with Reduced Maintenance Costs

This behavior arises from the way in which banks estimate the costs of successful attack based on their experience of such attacks in the past. As the information sharing system suppresses losses at a bank, that bank reduces its estimate of the potential losses associated with attack, and therefore is led to drop out of the system to reduce cost. This switching behavior occurs in spite of the fact that fixed costs are regarded as sunk, and the participation decision is predicated on recurring costs and leakage costs alone. When banks drop out of the sharing system, the frequency of successful attacks increases, leading to renewed awareness of attack costs, and therefore to increased participation in the sharing system.

5. Conclusions: Insights for policy – motivating robust collective security

The model developed here includes some important basic factors that bear on banks' decisions to adopt new protective technology. Even with a high level of abstraction, the model can provide some insights regarding barriers to good cooperative solutions, and possible means of surmounting them. First, if the cost of the system is (or is seen to be) high then that cost can create a barrier to obtaining an improved *global* state. Some mechanism for spurring coordination is necessary in this condition. With the notional parameter values used here, a reduction in maintenance cost (or equivalently a subsidy offsetting those costs) serves to foster adoption, which subsequently spreads through a large portion of the system. This strategy deters free riding by making the ticket less expensive. Third, the effectiveness of information sharing in reducing the perceived costs of attack (by reducing their success rate) may create a kind of complacency regarding the need to maintain participation. Mechanisms for offsetting a tendency for the financial system to "forget," such as prominent communication of the number of intercepted attacks, might be an especially important element of the information sharing environment.

References

Beyeler, Walter E., R.J. Glass, M.L. Bech, and K. Soramäki, Congestion and cascades in payment systems, *Physica A,* 15 Oct. 2007; v.384, no.2, p.693-718.

Davey, Victoria J., R.J. Glass, H.J. Min, W.E. Beyeler and L.M. Glass, Effective, Robust Design of Community Mitigation for Pandemic Influenza: A Systematic Examination of Proposed U.S. Guidance, *PLoSOne*, 3(7): e2606. doi:10.1371/journal.pone.0002606, *highly accessed*, 2008.

Glass, Robert J., A.L. Ames, W.A. Stubblefield, S.H. Conrad , S.L. Maffitt, L.A. Malczynski, D.G. Wilson, J.J. Carlson, G.A. Backus, M.A. Ehlen, K.B. Vanderveen, and D. Engi, Sandia National Laboratories: A Roadmap for the Complex Adaptive Systems of Systems CASoS) Engineering Initiative, Sandia National Laboratories SAND 2008-4651, September 2008

Glass, Robert J., A.L. Ames, W.E. Beyeler, B. Zak, D.A. Schoenwald, S.A. McKenna, S.H. Conrad , S. L. Maffitt, A General Engineering Framework for the Definition, Design, Testing and Actualization of Solutions within Complex Adaptive Systems of Systems (CASoS) with Application to the Global Energy System (GES), Sandia National Laboratories SAND 2008-7952, December 2008

Glass, Robert J., A.L. Ames, T.J. Brown, S.L. Maffitt, W.E. Beyeler, P.D. Finley, T.W. Moore, J.M. Linebarger, N.S. Brodsky, S.J. Verzi, A.V. Outkin, and A.A. Zagonel, Complex Adaptive Systems of Systems (CASoS) Engineering: Mapping Aspirations to Problem Solutions, Proceedings of the 8th International Conference on Complex Systems, Quincy, MA, 26 June – 1 July 2011, Sandia National Laboratories SAND 2011-3354, May 2011

Perlroth, Daniella J., R.J. Glass, V.J. Davey, A.M. Garber, D.K. Owens, Health Outcomes and Costs of Community Mitigation Strategies for an Influenza Pandemic in the U.S, *Clinical Infectious Diseases*, , 50:165–174, DOI: 10.1086/649867, *expedited publication*, 2010.

Renault, Fabian, W.E. Beyeler, R.J. Glass, K. Soramäki and M.L. Bech, Performance and resilience to liquidity disruptions in interdependent RTGS payment systems, Proceedings of the Joint Banque de France / European Central Bank conference on *Liquidity in interdependent transfer systems*, Paris, 9-10 June 2008.

Soramäki Kimmo, M.L. Bech, J. Arnold, R.J. Glass, and W.E. Beyeler, The topology of interbank payment flows, *Physica A,* 1 June 2007; vol.379, no.1, p.317-33.

Straffin, Philip D., Game Theory and Strategy, 1993, Mathematical Association of America