A SIMULATION OF DEFENSE AGAINST ADAPTIVE I.S. ATTACKERS

Norman Pendegraft, College of Business and Economics, University of Idaho, Moscow ID 83843, 208-885-7157, norman@uidaho.edu

ABSTRACT

A universe of two IS targets under attack by an adaptive attacker population is simulated. Various defensive strategies are examined. None is universally optimal. The model suggests that understanding attacker motivations and strategies is very important and that adaptive defensive policies have merit.

Keywords: information security, simulation

INTRODUCTION

This paper considers an information system (IS) with two targets under attack by a single group of similar attackers. Attackers are assumed to allocate their efforts according to one of several simple adaptive rules. Similarly designed defensive allocation strategies are tested. The key question is does it matter what policy the defenders choose?

In general, we assume that both attackers and users are attracted by increased value in the IS and repelled by increased security. The general model is as shown in Figure 1. Figure 2 illustrates user and attacker response to value and security, respectively. Attackers and users may have different inflection points.

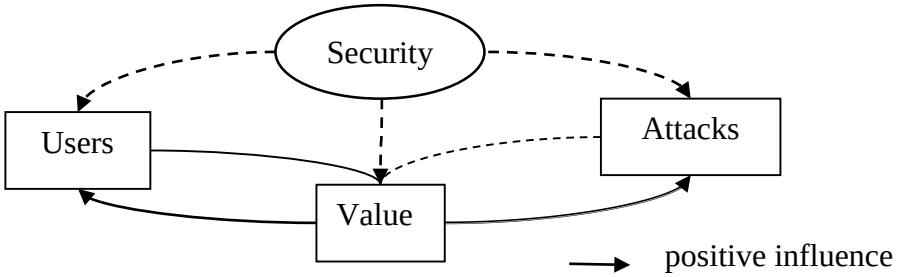

Basic Model
Figure 1

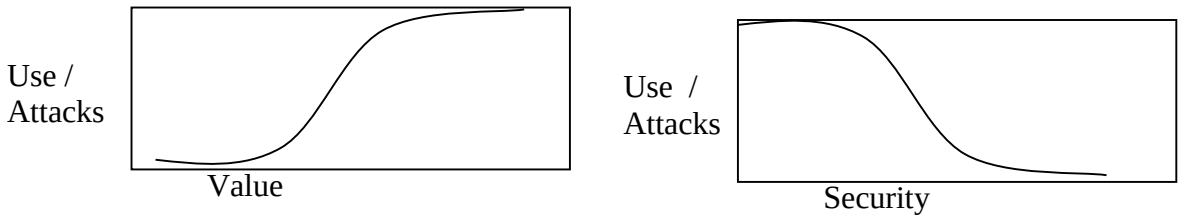

Response to Value & Security
Figure 2

BACKGROUND

The model presented here is an incremental extension to earlier work, but we also removed some complexity. In particular, [7] used variable user and attacker populations. [8] simulated a system with

two types of users. Here we add a second target, thereby introducing the problems of attacker choice. In order to keep the complexity manageable, we included only one user type and kept both attacker and user populations constant.

Much of the relevant background is summarized elsewhere [10] so is only briefly reviewed here. Much of the work in IS security (infosec) focuses on technical issues ([6] for example). Becker [2] argued that criminals behave rationally according to their perceived incentives. Many computer science research problems, outcomes are affected by human behavior [5]. We understand that the human behavior is a critical part of the security problem.

The growing systems dynamics literature on infosec attempts to model the interaction between the components of these systems. ([3], [4], [9]). [8] provides survey evidence for assumptions about user attitudes toward security, and [10] provide experimental evidence for the assumptions about responses to security and to value. The work is also influenced by the game theory literature [12].

The next section describes the model. The third section presents the simulation results. The paper concludes with a discussion of the results and possible extensions.

SIMULATION MODEL

Basic model

The model was built in IThink using a graphical stock and flow modeling interface which is converted into a set of finite difference equations (Appendix 1). The primary construct in the model is VALUE, a generic measure of information system value. The fact that this term is inherently fuzzy is addressed elsewhere [7]. There are three flows representing changes in VALUE. One represents increases due to usage and the second and third represent decreases due to attacks and to other costs. Attackers and Users are deterred by increases in security in similar ways as illustrated in Figure 2, albeit with different inflection points.

We model these S shaped functions with piecewise quadratic functions. $F_x(S)$ ranges from 0 to 1, reflecting no security at one extreme and complete security (i.e. shut down computers) at the other. Users and attackers (x) may have different inflection points, S_x^*. For a target system with security level S, the security factor is

$$\begin{aligned} Fx,s\,(S) &= 1-S_2/2S_x^{*2} && \text{for } 0\le S<S_x^* \\ &(1-S)^2 / 2(1-S_x^*)^2 && \text{for } S_x^*\le S\le 1 \end{aligned} \qquad (1)$$

Because VALUE is unbounded above, the model is piecewise over three ranges.

$$\begin{aligned} F_{x,v}\,(V) &= V^2/2V_x^{*2} && \text{for } 0\le V<V_x^* \\ &1-(2V_U^* -V)^2/(2V_U^2) && \text{for } Vx^*\le V\le 2V^* \\ &1 && \text{for } V>2V^* \end{aligned} \qquad (2)$$

The impact of value and security is multiplicative. That is:

$$F = k\, F_s(S)\, F_v(V)\,. \qquad (3)$$

We included the scale factor to allow for ease of model adjustment and to allow for differences in the magnitude of the impact on value between users and attackers. For users the scale factor is the sensitivity of value to use, i.e. how does use affect value. For attackers, the scale factor is the vulnerability of the system, i.e. how much a given level of attack reduces the value of the system.

There are a other differences between this and previous models. Perhaps the most significant is that here we assume that the number of users remains constant. This is primarily to avoid the substantial complexity that would be introduced not only by the change in user population.

06-Jan-17 2:12 PM

Modeling Attacker Behavior

The model has two targets and one representative set of attackers. The targets could be different servers at a company, say a database server and a web server, or servers at two different companies. Attackers use one of three strategies for allocating attacker effort across the two targets. Our control, or naïve, class of attackers is static. That is they allocate their effort equally to targets 1 and 2 and are, thus, indifferent to both security and value.

In the following V_i is the value of target i, and S_i is the level of security allocated to target i, and Ai is the fraction of attacker effort allocate to target i. Attacker strategies are summarized in Table 1. The first adaptive class of attackers (thieves) is attracted by value. Thus, the fraction of attacker effort allocated to each target is proportional to its value. Thus, $V_1/(V_1+V_2)$ is allocated to target 1 and $V_2/(V_1+V_2)$ is allocated to target 2. We do not assume that this is a deliberate decision by a central authority, but rather reflects the independent decision of all attackers taken in aggregate.

The second adaptive class of attackers (vandals) is assumed to be attracted by "easy" targets regardless of value. We model their behavior as follows. The fraction of attacks directed at target i is $(1/S_i) / (1/S_1 + 1/S_2)$. Thus, less secure systems see more attacks.

Our final group is, in some sense, a combination of these two and reflects the behavior of an attacker who is attracted by a mix of value seeking and security aversion. The fraction of its effort that this group allocates to target is $(Vi/S_i) / (V_1/S_1 + V_2/S_2)$. Again, we emphasize that is all cases, we are attempting to model aggregate behavior not the decision of a central directing authority.

Strategy id	Name	Description
0	Random	A_1 is U[0,1], $A_2=1-A_1$ U is uniform distribution
1	V	$A_1=V_1/(V_1+V_2)$
2	1/S	$A_1=(1/S_1)/(1/S_1 + 1/S_2)$
3	V/S	$A_1=(V_1/S_1) / (V_1/S_1 + V_2/S_2)$

Si = security allocated to target i, Vi= value of target i, Ai =proportion of attack on target i

Attackers Strategies
Table 1

Modeling Defenders' Decisions

Our model assumes that managers can reallocate security resources dynamically. We believe that this is a reasonable assumption. CIO's have the ability to adjust scarce resources in order to protect a multi-target system. For example, the CIO can shift security personnel from one location or project to another. It could be the decision to do an in depth forensic analysis on one compromise rather than another. More time and effort could be spent on analyzing logs of one system vs. another, perhaps uncovering unknown compromises. The shift of personnel could also be used in more expeditiously patching one system vs. another. Scarce personnel resources could also be focused on training users of one system verses another.

Servers and other systems already in operation could also be shifted. The CIO could enable encryption in one target and not the other. Because there is a significant overhead cost to encryption, servers could be shifted from one target to another to maintain throughput. The order in which various systems security is improved can also be adjusted. For example if one target is subject to higher perceived threat, new technology could be installed there before it is installed on other targets.

There can also be differences in the speed at which breaches are mitigated. For example, a more mission critical system could have backups done more often and have the hardware in place to restore them more quickly.

We studied four possible defensive policies. They are summarized in Table 2. Policy 1 is static: whatever the initial allocation, it remains the same. Policy 2 directs defensive efforts based on the values of the targets: more valuable targets get more security resource. Policy 3 allocates security based on effort as a fraction of target value, so that targets experiences attacks high relative to their own value get more resources. Policy 4 allocates resources proportionately to effort: the target with the most attacks get the most defense.

Policy id	Name	Description
1	static	S_1, S_2 are constant
2	W	$Si=Vi/(V_1+V_2)$
3	L/W	$Si=(Li/Vi)/(L_1/V_1 + L_2/V_2)$
4	L	$Si=Li/(L_1+ L_2)$

Si = security allocated to target i $S_1+S_2=1$, Vi= value of target i, Li = loss in value in target i in previous period =(Ai*sensitivity to attack *Attacker population)

Defensive Policies

Table 2

Model Calibration

We experimented with the model parameters and chose a set that allowed us to recreate the behaviors that were observed in [7]. We were particularly interested in systems in with increasing Value and increasing benefit from Security. Since they are, in some sense, normal. Thus, for these experiments we chose parameters that gave us that behavior.

Figure 3 is illustrative of the observed behavior for all attacker strategies. In that figure (and in subsequent graphs) the horizontal axis is time and the vertical axis is the total value of the system. In the base case, the defensive policy is static, and the lines represent, respectively security allocations of 10/90, 30/70, 50/50, 70/30, and 90/10. With a symmetrical system (i.e. identical targets, the optimal security allocation is 50/50.. We emphasize that this is an experimental world: there are other parameter sets that give different results. Thus, our conclusion should not be taken as universal, but rather as demonstrating a possibility.

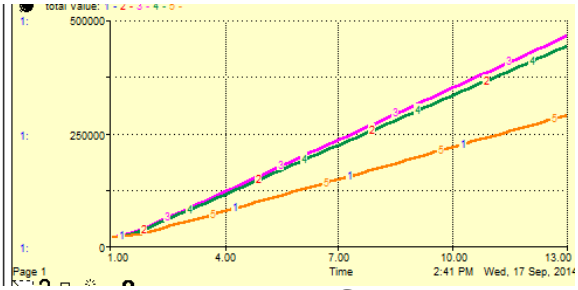

Base Case

Figure 3

EXPERIMENTS AND RESULTS

Symmetric Systems

In our first experiment, we used the same, symmetric system and compared the optimal static policy (50/50) with the three adaptive policies. Figure 4 is illustrative. Regardless of the attackers' strategy, the static policy was as good as the adaptive defensive policies. In the second experiment, we repeated the first, except that we used a non-optimal allocation for the static policy. Figure 5 illustrates those results. It shows that all of the adaptive policies outperform the static policy

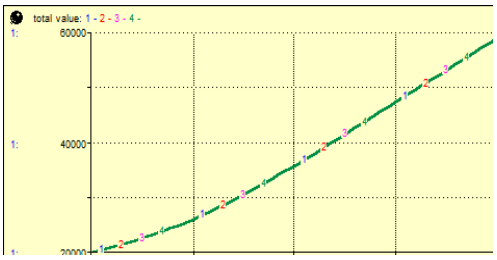

Symmetric System with
Optimal Static Policy
Figure 4

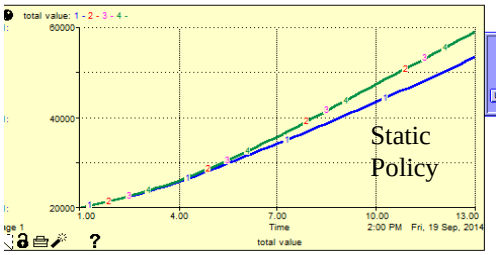

Symmetric Systems: Non-optimal
Static Policy Underperforms Adaptive
Figure 5

Non-symmetric Systems

In subsequent experiments, we used non-symmetric systems. We created the asymmetry by increasing the sensitivity to attack of target 2 by one-half an order of magnitude (from 100 to 300). First, we examined the impact of attacker strategy on the static defensive policy. As illustrated in Figure 6, the optimal static policy depends on the adaptive strategy used by the attacker. The attacker strategy is determined by the rotary switch to the right of each graph. As can be seen, the time evolution of the system value differs for each attacker strategy. In two cases (A and B), there is a clear superior static policy (in A a 70/30 allocation and in B a 30/70 allocation). In the third case static policy 2 (30/70) is just slightly better than policy 3 (50/50).

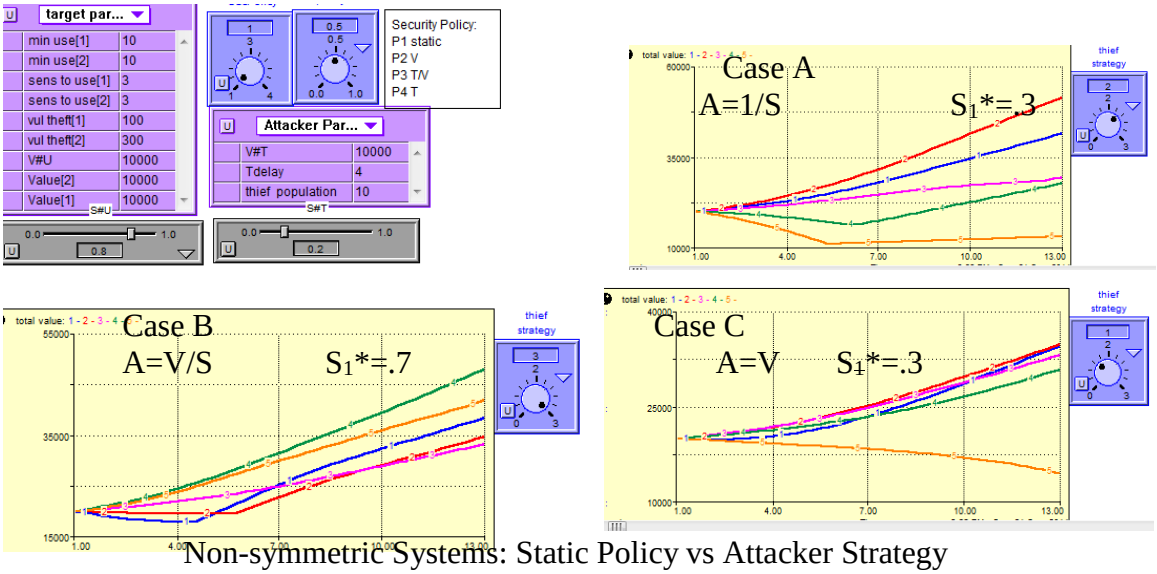

Non-symmetric Systems: Static Policy vs Attacker Strategy
S1 = 0.1,0.3,0.5,0.7,0.9
Figure 6

In our next simulation experiment we looked at the two cases in which there were clearly optimal static policies (cases A and B). However, now we compared the optimal static policy with the various adaptive policies. Figure 7 illustrates that the optimal static policy outperformed the adaptive policies. Finally, we repeated the previous experiment with a non-optimal static policy (50/50 rather than 70/30). As can be seen in Figure 8, in this case two of the adaptive policies are superior to the static policy.

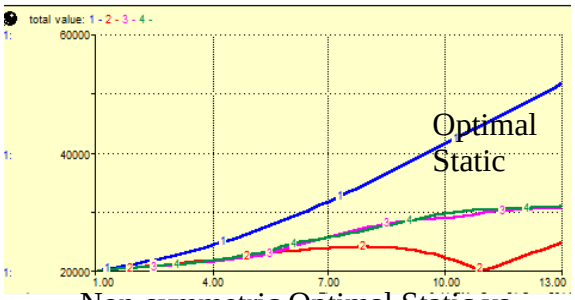

Non-symmetric Optimal Static vs Adaptive
Figure 7

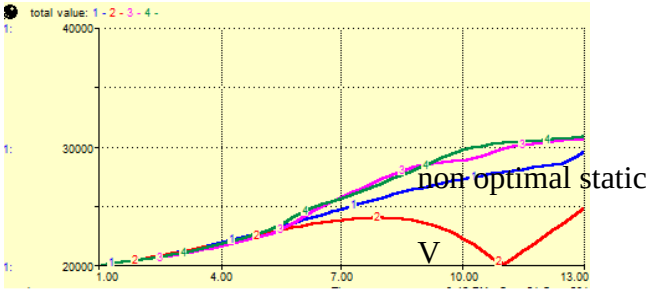

Non-symmetric:
Non-optimal static vs Adaptive
Figure 8

Cooperative Systems

In the next set of experiments, we looked at the impact of security allocation on the values of the individual systems as well as the total value of the systems. Figure 9 illustrates a typical result. In that case static allocations of security maximize value in the overall with an allocation of S_1=0.5. Note that the individual systems' values are maximized with S_1 = 0.1 or S_1 =0.9. The implication is clear and not surprising: cooperative strategies are superior in creating total value.

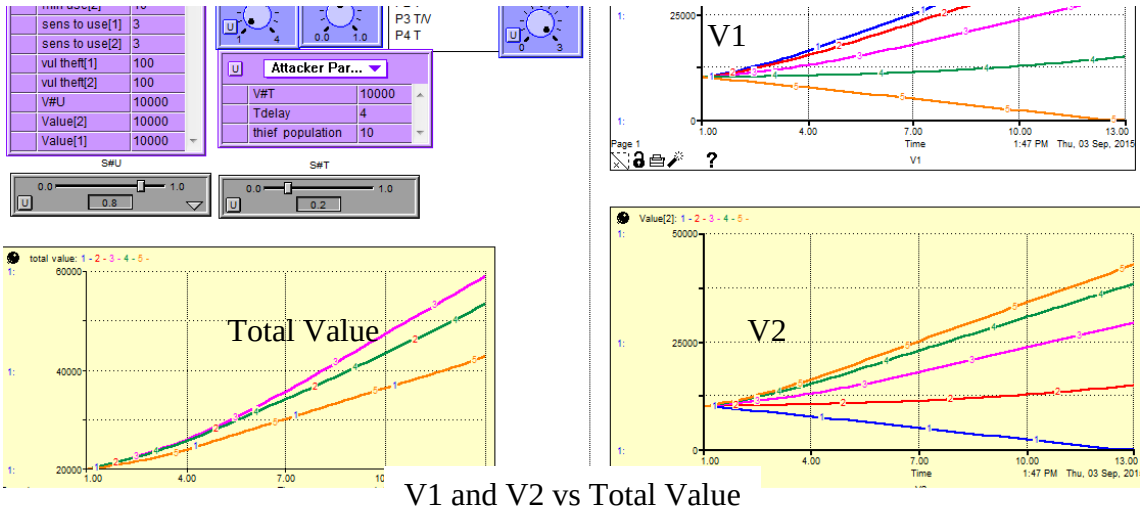

V1 and V2 vs Total Value
Figure 9

DISCUSSION

In answer to the question posed at the beginning of the paper, it is clear that attacker motivation does matter. The optimal static defensive policy changes when the attacker changes strategies. Further, if the defender uses a non-optimal defensive policy, adaptive defensive policies outperform the static policy. These results suggest that understanding attackers motivations is desirable, confirming the assumption implicit in [10].

As an aside, we recall the famous prisoner's dilemma tournament conducted by Axelrod [1]. In that tournament participants entered programs that played iterated prisoner's dilemma. The winning strategy was tit-for-tat which initially cooperates, and then plays as did the opponent in the previous turn. Tit-for-tat was the shortest program entered. This suggests to the authors that complex defensive polices will not necessarily be superior to simple strategies.

The results suggest any number of extensions. As noted earlier, it would be useful to extend the model to include variable populations of users and attackers and multiple user types. Better modeling of the attacker and defender target selection problem would be interesting. Experimental work like that in [10] would be useful in doing so. It would be desirable to develop a better model for budgeting security resources across multiple points of-leverage. These results suggest that this is a difficult problem. The current model assumes that changing defense allocation is cost free: this is probably not the case. It remains to extend the simulation to provide for a cost penalty in changing security allocation. Of course, the major remaining challenge is to calibrate the model to a real system

REFERENCES

[1] Axelrod, R. & Hamilton, W.D., 1981. The Evolution of Cooperation, Science, 211, 1390-1396.

[2] Becker, G. S., 1968. Crime and punishment: An economic approach. *J. of Pol. Econ.* 78, 169-217.

[3] Behara, R.R., Huang, C.D. & Hu., Q., 2010. A Systems Dynamics Model of Information Security Investments, J. of Information System Sec 6(2).

[4] Dutta, A. & Roy, R. 2008. Dynamics of Organizational Information Security, *System Dynamics Review* 24(3) 349-375.

[5] Grossklags, J. (2007). Experimental economics and experimental computer science: a survey, *ExpCS*, San Diego.

[6] Lehtinen, R., Russell, D., Ganemi, G.T., 2006. *Computer Security Basics*, O'Reilly, Sebastapol.

[7] Pendegraft, N. & Rounds,.M., 2007. A simulation model of IS security, *International Jl of Information Security and Privacy, 1*

[8] Pendegraft, N. , 2014. User Attitudes Toward Password Security: Survey and Simulation, *J. of Information System Security* 9#1.

[9] Rosenfeld, S.N., Rus,I., Cukier,M., 2007. Archetypal behavior in computer security. J. *System & Software* 80, 1594-1606.

[10] Rounds,M. J Alves-Foss,J., Pendegraft,N. 2013An Experimental Study to Explore Response to Changes in Security and Reward, *46th Hawaii International Conf. on System Sciences*, 2969-2977

[11] Rounds, M. , N. Pendegraft, & C. Taylor, 2007. The Ecology of IS Security: A Research Agenda. *Information Resources Management Association International Meeting,* Vancouver.

[12] Sandler R. & Lapan, H.E., 1988. Calculus of Dissent: An Analysis of Terrorists Choice of Target, *Synthese* 76#2 pp.245-261.

APPENDIX

Models details on request to author.