Software Diversity for Information Security

Pei-yu Chen[2], Gaurav Kataria[1] and Ramayya Krishnan[1,3]
[1]Heinz School, [2]Tepper School and [3]Cylab
Carnegie Mellon University

Abstract:
In this paper we analyze a software diversification-based strategy to achieve information security. The notion of using diversity to limit correlated risks is a widely accepted strategy in many fields. Various risk management approaches strive to minimize the variance of losses faced by individuals by either risk pooling, as in insurance, or diversification, as in portfolio management. However, these approaches are advantageous only for risk-averse agents as the expected loss remains unchanged. Exploiting externalities unique to information systems, we show that diversification can not only reduce loss variance but also minimize expected loss. We formulate the optimal amount of diversity investment by a firm taking into account both the negative network externalities accruing from attacks as well as positive network effects that accrue from uniformity and interoperability.

Keywords: Computer security, Information security, Software security, Risk management

I. INTRODUCTION

Network effects have been the driving force underlying a firm's decisions on technology adoption i.e. whether to adopt, what to adopt and when to adopt (Katz and Shapiro 1985 and 1986; Brynjolfsson and Kemerer, 1996). In the case of software adoption, firms often find it more valuable to adopt software which has large market share. By making a choice compatible to others' firms enjoy positive network effects stemming from greater benefits of compatibility and interoperability both within and outside the organization (Rohlfs 1974). As a result, markets with network effects are usually "tippy" i.e. tipping in favor of one product (Farrell and Klemperer 2001). The rise of MS Windows as the most popular choice for desktop operating system can be mainly attributed to this very fact (Economides 2001).

However, often ignored is the negative network externality associated with consuming popular software. This negative network externality has become increasingly important more recently as more and more security attacks take place, and at the same time, firms realize that their ability to stay secure is somehow dependent on the actions (e.g. patching) of others that use the same software. More specifically, a popular software may attract considerably more attacks due to its high market share. And, by using popular software to interconnect with many partners, firms risk being attacked and affected by the breaches at their partners (Kunreuther and Heal 2003). Therefore, by joining a larger network (e.g., sharing a software with more users) a firm may face higher risk. This observation has lately gained more traction since the recent string of some fairly devastating worms like MS-Blaster (CERT CA-2003-20) and Sasser (Symantec 2004). These worms have exploited the vulnerabilities present in Microsoft Windows operating system to propagate from computer to computer, eventually targeting most of the un-patched machines connected to Internet. Unfortunately, this meant most of all the world's machines were affected due to the fact that over 90% of all client-side computers use Microsoft Windows[1] (Geer et al. 2004). Whether the lower quality of Microsoft Windows

[1] According to market researcher OneStat.com, Windows now controls 97.46 percent of the global desktop operating system market, compared to just 1.43 percent for Apple Macintosh and 0.26 percent for Linux.

or the size of Microsoft's market share is the cause for large numbers of attacks against it is unknown. Some observers have cited economies of scale as the main reason why most attackers choose to attack Windows machines (Honeynet Project 2004, Symantec 2004).

It may be that in considering positive network externalities alone and disregarding negative externalities firms have over invested in homogeneous systems. Some observers have even argued that this has led to market failure in the case of OS market (Geer et al. 2004). In this research we aim to address the following research questions: Can a firm benefit from investing in a different IT infrastructure than other firms it does business with? Can a firm benefit from maintaining a diversity of systems? Can society benefit from such diversity? What is the social cost of lacking diversity? And what is the optimal level of diversity at the firm and society level?

We show that diversification not only reduces loss variance but also minimizes expected loss. We also provide a framework for determining optimal diversification strategy for a firm. In section II, we formulate the problem of a firm that needs to purchase software as a budget constrained decision problem. In sections III and IV, we discuss the benefit of homogeneity and reduction in expected loss via diversity, respectively. In section V, we discuss the optimal diversification strategy taking into account both the positive and negative externalities. Finally, we conclude by discussing our results and describing our larger research agenda to incorporate the role of industry, government and market forces for achieving socially optimal software diversity.

II. IT BUDGET

Expenditure on IT is a budget constrained problem that takes into account switching cost, training, migration, interoperability, and integration (Figure 1). Security is a newly

emerging and increasingly important constraint in the IT budget allocation process[2]. When selecting its software, a firm decides how much uniformity it wants with the external world and to what extent it is willing to stay distant (or diverse). We adopt the terminology of *homogeneity* (or *diversity*) to indicate on a scale of 0 to 1 the extent to which a firm is similar (or dissimilar) in its choice of software with the rest of the world. By staying homogeneous internally and externally it expects greater benefits of interoperability while risking catastrophic consequences via simultaneous failure of all its systems as has been the case with some recent worms[3].

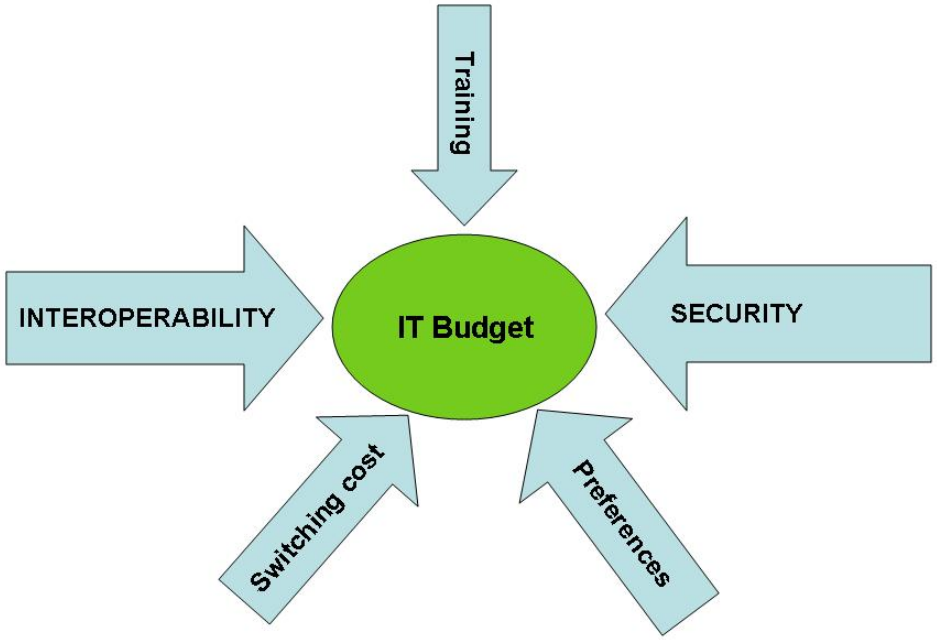

Figure 1: A firm has to choose between homogeneity and diversity with Interoperability and Security being the primary tradeoff.

We consider firm technology acquisition strategy, i.e., whether they should acquire technology of the same type to ensure maximum interoperability or they should "mix-

[2] A Worldwide Study Conducted by CIO Magazine and PricewaterhouseCoopers in 2003 said "Looking ahead to 2004, security will become more strategic as companies invest greater resources in developing strategy, defining architecture and risk assessment." http://www.csoonline.com/csoresearch/report64.html
[3] More than half of Korea's Internet backbone went down during SQL Slammer worm just due to denial of service (DOS) problem.

and-match" different technologies to reduce security loss. The strategic decision variable is the level of diversity, with interoperability and security risk as the main tradeoff. Other factors that can potentially impact the level of diversity are switching cost when firms have installed base from incumbent software. Without loss of generality, in our analysis we consider a firms software environment to consist of two technologies – an incumbent technology and a competing technology. The firm may choose to have x_1 proportion of its systems on incumbent software 1, while having the remaining $1-x_1$ on the competing software 2. Then, assuming a risk-neutral firm, the net utility derived can be written as follows:

$$E[U(x_1)] = \text{Benefit}(x_1) - E[\text{Loss}(x_1)] \quad (1)$$

Where,

Loss is stochastic and benefit is deterministic.

Within this framework we discuss benefits of homogeneity in Section III and benefits of diversity in Section IV.

III. BENEFIT OF HOMOGENEITY

The advantages realized by compatibility and integration have been discussed in the IS and standardization literature (Klemperer 1987, Varian and Shapiro 1999). In choosing software that has large market share, a firm ensures easy connectivity with its partners and suppliers, while at the same time, having all its internal systems operating on identical software ensures seamless interconnectivity. The advantages of such a setup have been widely discussed in data warehousing and process integration literature (IBM 2004).

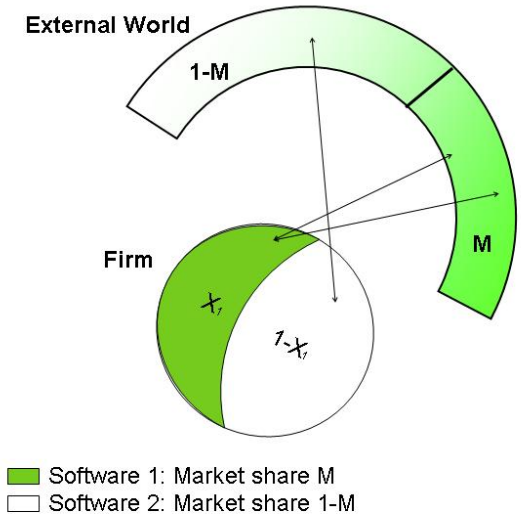

Figure 2: A firm's software choice affects its connectivity both internally and externally.

The benefit of homogeneity is primarily the benefit of interconnectivity. In a standardized environment this benefit should be independent of x_1, the level of diversity in a firm. Unfortunately, software interfaces today are not standardized, which means that having software from same vendor on all computers provides an "extra" benefit to the firms.

$$\text{Benefit(of having a combination)} = K_S[(x_1N)^2 + (x_2N)^2 + x_1N * M * E + x_2N * (1-M) * E] \quad (2)$$

Where,

K_S = Standardization coefficient

The standardization coefficient is a scaling factor to denote the benefit of interconnection of two computers running same software platform. Some standards do exist today (e.g. SOAP and XML), however due to many proprietary extensions and interfaces the full benefit of interconnection is not fully realized. The other variables in (2) are as follows:

N = Number of computers in the firm

x_1 = Proportion of computers running incumbent software 1

$x_2 = (1-x_1)$: Proportion of computers running the alternative software 2

M = Market Share of software 1; (1-M) = Market Share of software 2

E = Number of external computer that the firm may connect to

Solving (2) for optimal software choice gives a binary outcome: a firm would choose all software 1 if it has greater market share, and if software 2 has greater market share then it would choose all software 2, other things being equal, as shown in figure 3 below. We note that in absence of negative externalities a firm would prefer to invest only in the incumbent software (Farrell and Saloner, 1985). May be this is the reason why society today is over-invested in Windows. In the following section we discuss the negative network externalities and show that reduction in losses via diversity can possibly compensate for the positive network effects with homogeneity.

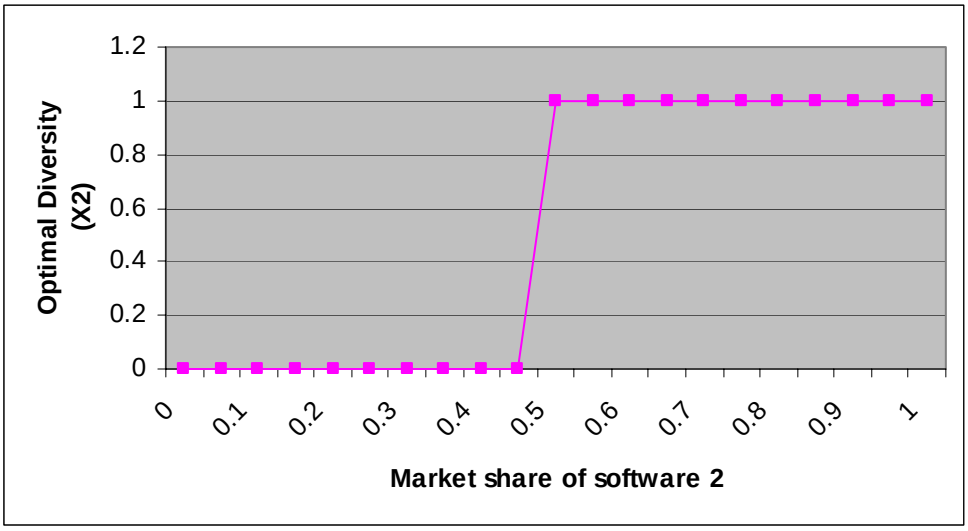

Figure 3: Firms prefers all software 1 when market share of 1 is over 50%, when market share of 2 increases above 50% it prefers all software 2.

IV. LOSS REDUCTION VIA DIVERSITY

By virtue of always being connected and tightly integrated in online business processes, it is widely accepted today that firms receive numerous attacks on their systems. Even without counting targeted attacks the baseline rate of stray virus/worm type of attacks on corporate networks is considerably high (CSI/FBI Survey 2004). There is no doubt that

even after following best practices, systems fall prey to online attacks on a daily basis. Accepting this harsh reality most firms have measures in place to tackle these incidents. Some information security risk management frameworks have been proposed to better understand and tackle this problem (Hoo 2002, Butler 2003). Soo Hoo's approach focuses on Annualized Loss Expectancy (ALE) to choose between security measures, while Butler's approach is more qualitative and based on one-to-one interviews with the management to determine the relative risks of possible outcomes and effectiveness of the risk mitigation approaches. A CMU-LBNL joint study on information security risk quantization states time-to-respond to an incident as the measure of loss incurred (Arora et al. 2004). Recognizing that in addition to scale of the incident, different type of incidents require different attention, they define time to respond as a combination measure of various efforts including diagnostic, repair, legal, public relations, and mandatory reporting; each of which is required in varying proportion to tackle different incidents. For instance, a virus attack may require no more than disk scanning and cleaning, while graffiti on the company's webpage is a public relations nightmare.

Firms employ resources to tackle these scenarios which occur on a daily basis. In our research we model the capability of the firm to respond to such scenarios as a fixed and limited resource. On the other hand, the scale of the incidents varies considerably on a daily basis. In this paper, we consider the number of computers affected by a worm(s) outbreak on a day, as a measure of the seriousness of the incident. In this case the contingency operation of the limited-resource IT department involves patching, backup and/or rebuilding systems. It is perhaps reasonable to assume that the loss to a firm due to computers being affected grows more than linearly to the number of computers affected. Consider the following situation where the loss to the firm may be captured by the total waiting time (or downtime) to bring all the computers back up. If we assume that the IT department services the affected computers in a sequential manner, then the organization faces lost productivity not only for the computer that is currently being serviced but also for all the other affected computers that are waiting in some sort of queue for being serviced by the IT department. Therefore, the damage caused to the firm when y of its computers is affected in a worm outbreak is of the form:

$$L(y) = ky^2 \quad (3)$$

Where,

y is normalized such that it varies from zero to infinity.

By this simple analysis we are trying to depict the loss to an organization as a function of y, the scale of incident (Figure 4). Even though this may not be the most accurate depiction of the loss function, it highlights the non-linear relationship between the loss accrued and the scale of the attack. The important point being that simple attacks are handled with little or no effort (e.g. blocking ports by properly configuring the firewall) while the bigger incidents require much extensive effort and downtime. The total service time increases rapidly with increasing scale of attack.

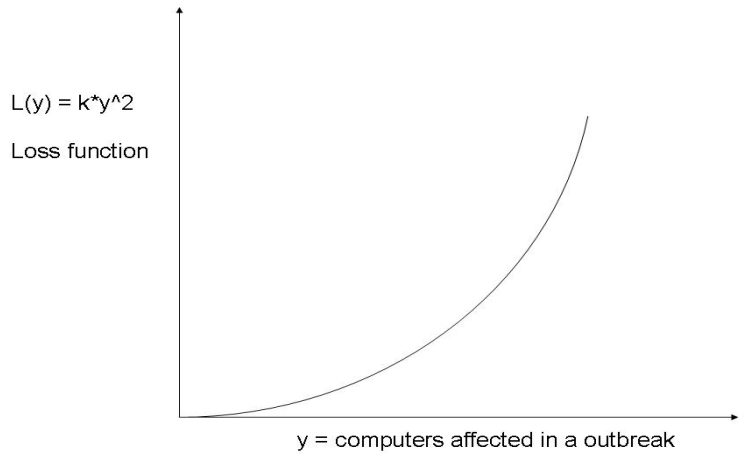

Figure 4: A convex loss function depicting the rapidly rising service effort with increasing scale of incident.

Because we know that the proportion of computers affected on a day is not fixed, it is instead likely to have a distribution where minor incidents happen more frequently and major ones are not so frequent. A minor incident from the perspective of the firm is when just a few computers are affected and a major incident is when a large percentage of computers are down. This assertion is not without factual support. Numerous websites

report vulnerability statistics which show that now-obsolete attacks are still being observed in large numbers and outdated viruses are still in circulation[4]. On the other hand, once every few months we see a major worm which successfully exploits large number of computers worldwide (MS Blaster, Sasser, Code Red, Nimda etc). The LBNL-CMU study and Butler's interviews with security managers offer data to support these statements (Arora 2004, Butler 2003). Based on these observations we model the probability density function for the scale of incident as an exponential distribution,

$$f(y) = \frac{1}{\beta} e^{-y/\beta} \qquad (4)$$

Where,

y = number of computers affected on a day
(Normalized to vary between zero and infinity)

β = mean number of computers affected on a day

The average number of computers affected, β, may depend on many factors including, type of software service, type of industry, inherent security level of software product, market share and/or sentiment against the software product etc. However, as shown in Table 1, Windows which has over 90% market share in client-side operating system market receives 91.6% of attacks at ten most attacked ports. At the same time, the number of vulnerabilities discovered in Windows outnumbers all other operating systems. Market share thus appears to have considerable influence on the attacks.

[4] Symantec, CERT, ISC SANS etc report that many obsolete viruses and variant of already patched worms appear frequently.

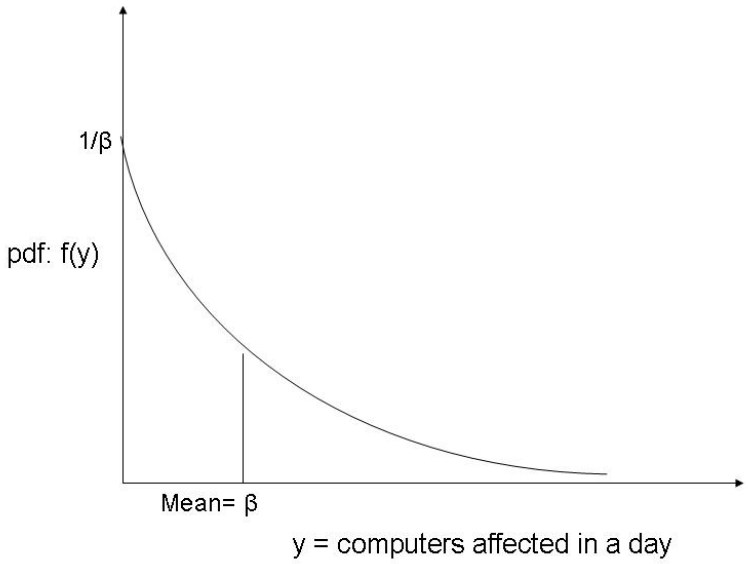

Figure 5: Exponential pdf of computers affected with mean = β. On a given day, it is more likely that fewer computers are affected; the likeliness decreases as the scale of incident increases.

Table 1: Top Ten Most Attacked Ports According to SANS Institute Internet Storm Center

Port/ Protocol	Service	Average number of source IP attacking this port 1/17-2/14	Primary Target OS	Vulnerability Related Information
445/ tcp,udp	microsoft-ds	375026	Windows	CA-2003-19: Exploitation of Vulnerabilities in Microsoft RPC Interface CA-2003-20: W32/Blaster worm CA-2003-23: RPCSS Vulnerabilities in Microsoft Windows
135/ tcp,udp	epmap	159180	Windows	CA-2003-19: Exploitation of Vulnerabilities in Microsoft RPC Interface CA-2003-20: W32/Blaster worm Current Activity 08/18/2003: W32/Welchia Worm
139/ tcp,udp	netbios-ssn	52964	Windows	CA-2003-03: Buffer Overflow in Windows Locator Service CA-2003-19: Exploitation of Vulnerabilities in Microsoft RPC Interface CA-2003-20: W32/Blaster worm CA-2003-23: RPCSS Vulnerabilities in Microsoft Windows
1025/tcp	RPC	60451	Windows	Currently inbound scans are likely RPC and LSA exploit attempts against the Windows
1026/udp	PopUp Messenge	56451	Windows	Typically inbound traffic to this port is Messenger Spam
53/ tcp,udp	dns	23899	Linux/Unix	CA-2002-31: Multiple Vulnerabilities in BIND
80/tcp	http	19790	Windows and Linux	CA-2002-27: Apache/mod_ssl Worm CA-2002-33: Heap Overflow Vulnerability in Microsoft Data Access Components (MDAC) CA-2003-09: Buffer Overflow in Core Microsoft Windows DLL Current Activity 08/18/2003: W32/Welchia Worm
1433/tcp	MS-SQL	9867	MS-SQL Server	Inbound scans are typically looking for Microsoft SQL Server installations with weak password protection and if successful are looking to steal or corrupt data or use some features with SQL Server to compromise the host system.
1027/udp	PopUp Messenge	39057	Windows	Typically inbound traffic to this port is Messenger Spam
137/udp	netbios-ns	24079	Windows	CA-2003-08: Increased Activity Targeting Windows Shares CA-2003-23: RPCSS Vulnerabilities in Microsoft Windows

Source: Data compiled from Internet Storm Center at SANS and US-CERT.

Percentage attacking windows: 752050 / 820764 = 91.6%

Now, given the loss function and the pdf of scale of incident, we can calculate the expected loss to a firm.

$$E(loss) = \int_0^\infty L(y) f(y) dy$$
$$= 2k \int_0^\infty \frac{y^2 e^{-y/\beta}}{2\beta} dy \quad (5)$$
$$= 2k * \beta^2$$

Without loss of generality we assume k =1/2 for rest of the analysis.

IV.1 Diversity As A Means To Security

We have shown earlier that considering compatibility alone leads to homogeneity. However, homogeneity may also lead to higher security risk (Kunreuther and Heal, 2002; Geer et al, 2004). In this section, we would like to examine whether diversity can be an effective way to achieve higher security level, and how much diversity is needed in order to reduce expected loss due to security threats.

Consider, the firm decides to diversify its software use in order to reduce the chances of simultaneous failure of many computers; it may do so by keeping x_1 proportion of its computers on the incumbent software while switching to a competing product for the remaining $(1-x_1)$ portion.

In order to determine the expected loss when a firm decides to use a combination of software, we need to calculate the pdf of total number of computers affected which in this case is the combination of two random variables,

$$y_d = x_1 y_1 + x_2 y_2 \quad (6)$$

Where,

y_d = total computers affected by attacks on both types of software platforms

$x_1 y_1$ = total computers affected by attacks on incumbent software

x_2y_2 = total computers affected by attacks on competing software; $x_2 = 1- x_1$

If both the software products/platforms are nearly identical except that bugs in one are independent of other, then we can assume that $\beta_1 = \beta_2$, which implies that the pdf of y_d is

$$f(y_d) = \frac{1}{\beta(1-2x_1)}(e^{-y_d/\beta(1-x_1)} - e^{-y_d/\beta x_1}) \quad (7)$$

Continuing further we calculate the expected loss, as

$$E(loss) = \int_0^\infty L(y_d) f(y_d) dy_d$$
$$= \beta^2(1 + x_1^2 - x_1) \quad (8)$$

Minimizing the expected loss with respect to x_1, we get the minimum value as $0.75\beta^2$ for x_1=0.5. This result is not surprising as we assumed the two products to have equal β. However, since we normally observe the products are not identical in their characteristics, the number and intensity of attacks faced by two products is usually different leading to different values for β. Therefore we next assume that,

$\beta_1 = \beta$ while $\beta_2 = m*\beta$

Where,

m = is a function of all the factors that cause severity or number of attacks against a product to increase, possibly its larger market share.

Individually the probability density for y_1 and y_2 can be given by,

$$f(y_1) = \frac{1}{\beta} e^{-y_1/\beta} \text{ and } \quad f(y_2) = \frac{1}{m\beta} e^{-y_2/m\beta} \quad (9)$$

The pdf of the y_d can be calculated as

$$f(y_d) = \frac{1}{\beta(m(1-x_1)-x_1)}(e^{-y_d/m\beta(1-x_1)} - e^{-y_d/\beta x_1}) \quad (10)$$

Now calculating the expected loss as we did before in the case of no diversity,

$$E(loss) = \int_0^\infty L(y_d) f(y_d) dy_d$$
$$= \beta^2[m^2(1-x_1)^2 + x_1^2 + mx_1(1-x_1)] \quad (11)$$

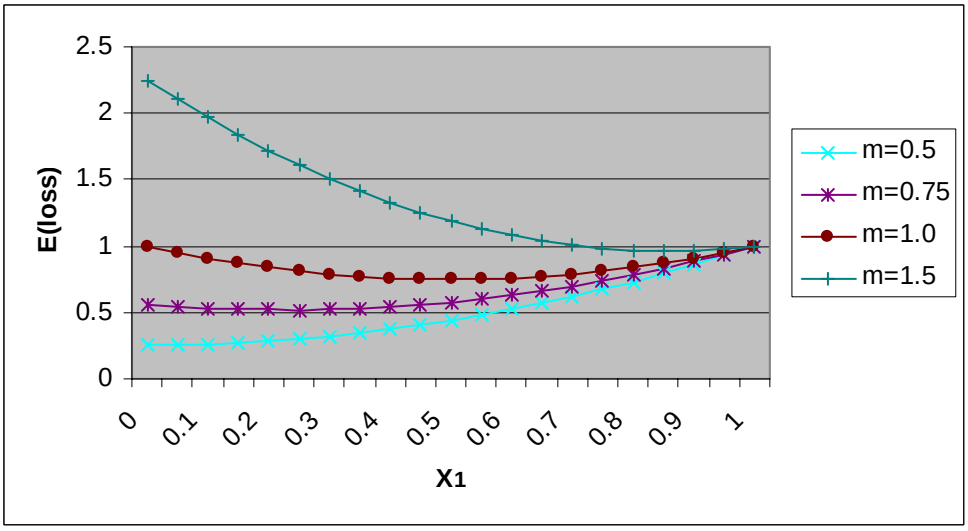

Figure 6: Normalized Expected loss as a function of x_1 for four different values of m. E.g. when two products are almost identical then E(loss) is minimum at $x_1=x_2=0.5$

Differentiating E(loss) with respect to x_1, we see that the minimum loss is realized when,

$$x_1 = \frac{2m^2 - m}{2(m^2 - m + 1)} \quad (12)$$

As an illustration, when m = 0.9 (i.e. when product 2 on average receives 10% fewer computer casualties than product 1), then the optimal amount of diversity is x_1 = 0.396. This means that product 2 which is superior in its security should be used for 60% of the computers. Another interesting observation from this analysis is that diversity is effective

in reducing expected security loss when the ratios of the security levels of non-incumbent software to incumbent software is between 0.5 to 2 (as shown in Figure 6). That is, even though the non-incumbent software is not as secure as the incumbent software, a firm may still benefit from acquiring non-incumbent software for a small set of its machines. The optimal amount of diversity i.e. x_2 (=1-x_1) is plotted against m in Figure 7.

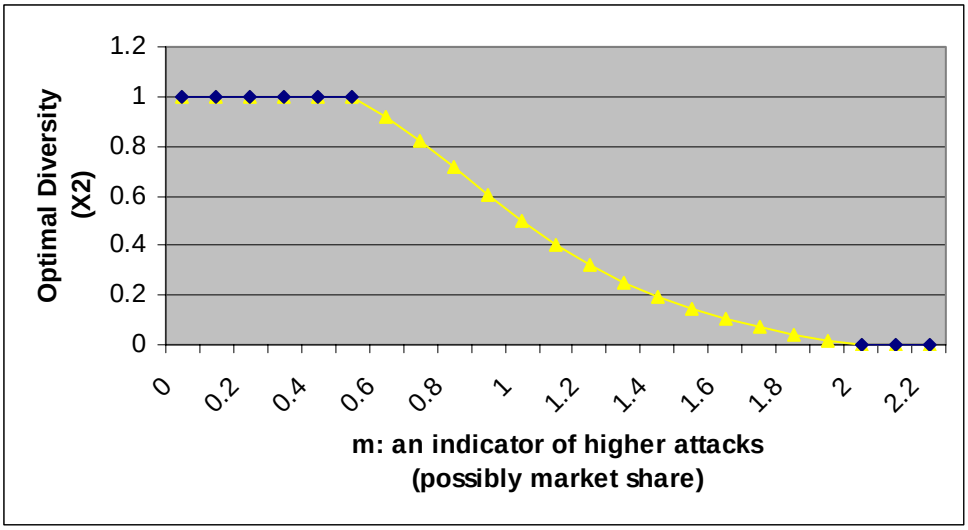

Figure 7: Optimal amount of diversity as a function of relative security m; higher m implies that software 2 receives more attacks vis-à-vis software 1.

It may also be interesting to consider the impact of diversity on the variance of loss, given that a small variance is generally preferred to a large variance. Therefore, extending the above analysis to account for the variance of the loss we have,

$$V(loss) = E(loss^2) - E^2(loss)$$
$$= \int_0^\infty L^2(y_d) f(y_d) dy_d - \left(\int_0^\infty L(y_d) f(y_d) dy_d \right)^2$$
$$= 6\beta^4 \frac{(m^5(1-x_1)^5 - x_1^5)}{m(1-x_1)+x_1} - \left[\beta^2 \frac{(m^3(1-x_1)^3 + x_1^3)}{m(1-x_1)+x_1} \right]^2 \qquad (13)$$

$$= 1.088\beta^4 \qquad evaluated \quad at \quad [m = 0.9, \quad x = 0.396]$$

As compared to variance of loss when only software 1 is used given by,

$$
\begin{aligned}
&V(loss) = E(loss^2) - E^2(loss) \\
&= \int_0^\infty L^2(y) f(y) dy - \left(\int_0^\infty L(y) f(y) dy \right)^2 \\
&= 6\beta^4 - \beta^4 \\
&= 5\beta^4
\end{aligned}
\tag{14}
$$

Thus, in addition to the reduction in expected loss firms also see a five fold reduction in their variance when they switch 60% of their system from an incumbent to a competitor which is 10% safer.

V. OPTIMAL DIVERSITY

Prior literature on technology and software adoption has considered only positive network externalities and has not taken negative network externalities into account. However, recent industry reports (Geer et al. 2003) suggest that negative network externalities exist and take the form of a higher security risk associated with consuming a popular software. This paper is, to the best of our knowledge, the first attempt to analyze security risks and diversification strategies using the lens of positive and negative network externalities. We offer a way to quantify the benefits of diversity and show that diversity can be an effective way to reduce security risk—by reducing expected loss and variance of loss. There are a number of extensions we are working on as part of future work. For instance, if firms make choices as described in our paper, we recognize that market shares of the incumbent and competing software will change and thereby alter their market shares. This calls for an analysis of equilibrium behavior and we are pursuing this within a fulfilled expectations framework (Katz and Shapiro 1985).

The goal of this optimization strategy is to maximize the overall utility realized by the firm. Without more concrete data we cannot precisely estimate the exact amount of diversity required. However, we believe that more careful examination of data can give reliable estimates for the parameters of our model. In the following section we discuss

our larger research agenda and describe how we plan to precisely estimate those parameters.

VI. DISCUSSION AND CONLCUSION

"Recent data from our honeynet sensor grid reveals that the average life expectancy to compromise an unpatched Linux system is 3 months……..data from the Symantec Deepsight Threat Management System indicates a vulnerable Win32 system has life expectancy not measured in months, but merely hours."
-Honeynet Project: Trend Analysis Dec 2004.

Windows has long been the popular choice for desktop computing, but as more alternatives emerge e.g. Linux and Mac OS X, firms may prefer some lack of interconnectivity for reduction in security losses. However, some questions still remain: is maximization of utility a corner solution in favor of homogeneity as has been the case in the past? Can a society benefit from diversity? What's the social cost of lacking diversity?

In our research we aim to estimate the optimal level of diversity for both an individual firm as well as for society. In this paper, we have presented a novel framework to incorporate the benefits of both homogeneity and diversity in software domain. Specifically, we have shown that depending on the characteristics of the firm, industry in general, and type of software, different levels of diversity may be optimal. We are now interested in addressing some of the important questions like: should large firms prefer more diversity as compared to small firms, should government mandate standardization and should government subsidize development of competing software.

We hope to answer these questions by building a more accurate model for lost productivity as a function of scale of incident, and estimating the same using the call center data from CMU computing services.

REFERENCES

Arora, A., D. Hall, C. A. Pinto, D. Ramsey and R. Telang (2004). "Measuring the Risk-Based Value of IT Security Solutions," *IEEE IT Professional Magazine*, 6(6): 35-42.

Brynjolfsson, E. and C. Kemerer (1996). "Network Externalities in Microcomputer Software: An Econometric Analysis of the Spreadsheet Market," *Management Science*, 42(12): 1627-2647.

Butler, S. (2002). "Security Attribute Evaluation Method: A Cost Benefit Approach," International Conference on Software Engineering (ICSE 2002).

Cert (2003). *CERT Advisory CA-2003-20 W32/Blaster Worm*, http://www.cert.org/advisories/CA-2003-20.html

CSI/FBI (2004). Ninth Annual- Computer Crime and Security Survey. http://www.gocsi.com/forms/fbi/csi_fbi_survey.jhtml

Economides, N. (2001). "The Microsoft Antitrust Case," *Journal of Industry, Competition and Trade: From Theory to Policy*, 1(1): 71-79.

Farrell, J. and G. Saloner (1985). "Standardization, compatibility and Innovation," *Rand Journal of Economics*, 16: 70-83.

Farrell, J. and P. Klemperer (2001). "Coordination and Lock-in: Competition with Switching Costs and Network Effects," in M. Armstrong and R. Porter, eds., *Handbook of Industrial Organization*, vol. 3.

Geer, D., R. Bace, P. Gutmann, P. Metzger, C. Pfleeger, J. Quarterman, B. Schneier (2003). "CyberInsecurity: The Cost of Monopoly How the Dominance of Microsoft's Products Poses a Risk to Security," http://www.ccianet.org/papers/cyberinsecurity.pdf

Honeynet Project (2004). "Know Your Enemy: Trends," http://www.honeynet.org/papers/trends/life-linux.pdf

Hoo, K.S. (2002). "How much is enough? A Risk Management Approach to Computer Security," Workshop on Economics and Information Security, University of California, Berkeley.

IBM Inc. (2004). "IBM e-business Technology, Solution, and Design Overview," IBM Redbooks, http://www.redbooks.ibm.com/redbooks/SG246248.html

Katz, M. L. and C. Shapiro (1985). "Network Externalities, Competition, and Compatibility," *American Economic Review*, 75(3): 424-440.

Katz, M. L. and C. Shapiro (1986). "Technology Adoption in the Presence of Network Externalities," *Journal of Political Economy*, 94(4): 822-841.

Klemperer, P. (1987). "Markets with Consumer Switching Costs," *The Quarterly Journal of Economics*, MIT Press, 102(2): 375-94.

Kunreuther, H. and G. Heal (2003). "Interdependent Security," *Journal of Risk and Uncertainty*, Kluwer Academic Publishers, 26(2): 231-249.

Rohlfs, J. (1974). "A Theory of Interdependent Demand for a Communications Service," *Bell Journal of Economics,* 5(1): 16-37.

Shapiro, C. and H. Varian (1999). *Information Rules: A Strategic Guide to the Network Economy*, Boston: Harvard University Press.

Symantec Inc. (2004). "Symantec's Internet Security Threat Report," Volume VI, September 2004.