

Can spending on information security be justified?

Evaluating the security spending decision from the perspective of a rational actor

Andrew Stewart¹

February, 2009

Abstract. We investigate the optimality of various strategies for spending on information security. Being able to understand the strengths and weaknesses of spending strategies is useful to organizations. Our analysis begins with a whole-systems view of the security spending decision that encompasses people, technology, and economics. We then present a taxonomy of justifications for spending on information security. Each justification within the taxonomy is discussed, with that analysis used to examine the apparent rationality of a number of common spending strategies. A model is constructed that can be used in a practical manner to enable an organization to select a rational approach to spending on information security. That model is based on the finding that many of the common justifications for spending are unconvincing in the absence of actuarial data. Also incorporated within the model is the observation that a number of resulting pressures push companies towards inefficiency in their spending.

Keywords and phrases: Information security; Information systems; Spending strategies; Efficiency; Incentives; Psychology; Economics

1 Introduction

When faced with capital budgeting decisions, managers are bound by fiduciary duty to identify those investments that will maximize shareholder value. As such, decisions about spending must be carefully considered and evaluated in rational economic terms. The *efficiency* of spending on information security matters to companies. The *effectiveness* of that spending affects the company, its business partners, and ultimately its customers.

In this paper we consider the security spending decision to comprise of three questions:

- A. Should I spend? B. How much should I spend? C. What should I buy? (1)

These three questions should ideally be evaluated left to right. In other words, there is an order of operations. If an organization doesn't need to spend, it shouldn't jump straight to thinking about how much it *will* spend. Similarly, if an organization only needs to spend a little and jumps straight to thinking about what products are available on the market and what it will purchase, then this can lead to overspending. More formally, if the answer to question A can be expressed as a value between 0 and 1,

¹email: andrew.j.stewart@mac.com

then if an organization skips question A and jumps to question B , then the likelihood that it will spend inefficiently increases as A approaches 0. If an organization skips both questions A and B and jumps to C , its likelihood of spending inefficiently increases as A or B approach 0.

We are concerned in this paper with the spending decision at the unit of an entire company. As such, our focus is primarily on question A : whether to spend or not, and to a lesser extent on question B : how much to spend. Question C is best answered by considering the range of techniques, technologies, products, and services available in the market.

Today there are an array of justifications employed for spending on information security. There are also a number of well-known approaches used for attempting to determine the “right” amount to spend. We will examine whether data exist to support these common justifications and approaches. A key question is in determining the extent to which the value of security efforts can be measured. Cost-benefit analysis is appealing, and there is much talk in the popular information security press about “calculating the return on investment for security.” A number of factors make such a calculation difficult. One notable hurdle is that as security efforts become more successful, their value-added typically becomes less visible and therefore harder to measure. Many security technologies are designed and deployed in an attempt to avoid some undesirable event, such as a security breach. If the security initiative is successful and a breach does not occur, the challenge then is in proving that the breach did not occur as a result of the initiative. Perhaps the outcome was just luck? Or perhaps a breach wouldn’t have occurred even if the security measures were not deployed?

When considering any spending decision it is important to consider opportunity costs. Funds *not* spent on information security could be used for some other purpose. For example, spending \$3 million on security means forgoing the use of those funds to develop a new line of business or launch a marketing campaign. By recognizing opportunity costs, our examination of the security spending decision will be economic in nature and not one purely of accounting.

The main contributions of this paper are twofold. First, in identifying that a number of weaknesses in common justifications for spending on information security result in the selection of spending strategies that create inefficiency in the round. Second, in the presentation of two spending strategies that compensate for those weaknesses. These spending strategies are designed to be practical and highly straightforward for organizations and practitioners to use. This is in contrast to more complicated models that employ techniques such as Monte Carlo methods, regression analysis, and such like.

The remainder of the paper is organized as follows. We begin in §2 with a whole-systems view of the security spending decision that notes a number of structural hurdles. In §3 we present a simple taxonomy that describes a range of justifications for spending on information security. Those justifications are further described in §4. We observe in §5 that a number of pressures push the average company towards inefficiency in its security spending. In §6 a practical model is presented that can be used by an organization to guide its approach to the security spending decision.

2 A whole-systems view of factors that can influence the security spending decision

There are a number of factors that can influence the decision-making process when the security spending decision is considered. Some of these factors exist because of the nature of information security. Others are common where decisions are taken by humans, and these can be better understood through the study of psychology and economics. Below, we broadly consider the influence of people, technology, and economics on security spending.

The topic of security requires the consideration of risks. As such, there are a number of psychological effects that can occur. Viscusi (1997) shows that when presented with a number of descriptions of a particular risk scenario, the majority of people will remember the most alarming description. In other words, human beings tend to react to risk in an emotional manner and we are typically most influenced by the high-water mark for the severity of the possible outcome. An executive that reads the bad press concerning security vulnerabilities in wireless networking might be so worried about those risks that they select a policy of disabling the wireless access cards in the laptops the organization provides to its employees. The opportunity cost forgone is the productivity gain that the employees would receive from being able to roam wirelessly.

A ‘big event’ such as a widespread virus infection can create a risk averse mind-set. Depending on the impact of the incident, an organization might swing from having little or no desire to spend on security to an overwhelming desire to spend. The concern here is that money might be thrown at the problem in a reactionary manner. Slovic (2000) shows that people find it hard to think rationally about risks that carry heavy costs.

A major security incident can have such an emotional impact that it creates a desire to “fight the last war.” If a computer is compromised because of a missing patch, then the limited resources of the security team might be swung around to focus on applying patches. In reality there might be more risk in the software running *on* the computers and so the focus is misdirected. This phenomenon is known as the *availability heuristic* in which a person bases their prediction regarding the frequency of an event on the example that comes most easily to mind.

A manager responsible for the security spending decision within an organization has certain incentives to spend on security measures. Security projects create “visible achievements” within the organization: security software installed, security training carried out, all of which create positive exposure for the manager, potentially advancing their career goals. This is a *principle-agent problem* where the goals of the security manager do not necessarily reflect the goals of the shareholders.

Principle-agent problems crop up in several areas where they might affect the security spending decision. An auditor who is auditing the information security of a company has an incentive to point out a tremendously long list of problems. That thoroughness might be partially based on a desire to avoid liability in the future if a problem occurs because of an issue that was not identified in their report.

Agency problems are not unique to information security and in the above text we have not performed a comprehensive survey of the fields of economics and psychology in order to identify the range of their potential influence over the security spending decision. Our goal has been to frame the subsequent content with the overarching idea that a whole-systems view is important to consider. Spending driven by emotion or perverse incentives is very likely inefficient.

3 A taxonomy of candidate justifications for spending on information security

Below, we describe a number of justifications used in the current day for spending on information security. The term *candidate* is used so as not to imply that these justifications are proven. To identify the various candidate justifications we performed a review of the literature and surveyed the popular information security press. A number of interviews were also carried out with practitioners working as security consultants and employees within the information security groups at a number of US and UK-based organizations.

Each candidate justification claims to describe a tangible benefit that can be received by an organization that spends on information security. Other classes of justification exist. For example, the Fall 2008 issue of the *Disaster Recovery Journal* contains an article arguing that spending on business continuity is a moral duty.

To organize the various candidate justifications we employ a simple hierarchical representation, shown in Figure 1. This enables us to broadly define and categorize the information so that it can be discussed.

Candidate justifications for spending on information security

- a. *Avoid future losses*
 - i. *Cost of break/fix repairs*
 - ii. *Damage to reputation*
 - iii. *Litigation costs*
 - iv. *By creating a deterrent effect*
- b. *Create future gains*
 - i. *Enable a business process*
 - ii. *Broaden customer base*
 - iii. *Meet a required standard*
 - iv. *Create competitive advantage*
 - v. *Achieve new functional capabilities*
- c. *Create or increase efficiencies*
 - i. *Create operational efficiencies*
 - ii. *Provide multi-purpose controls*

Figure 1: Simple taxonomy of candidate justifications for spending on information security.

The taxonomy is organized from the perspective of the potential benefit received by the organization. The root node is the single classification of candidate justifications for spending on information security. Through inheritance, every node within the tree is also only a *potentially* valid justification for spending. The second level is a tri-part classification between those justifications that claim to *avoid future losses*, those that claim to *create future gains* (or generate other positive effects), and those that claim to *create or increase efficiencies*.

4 Analysis of candidate justifications

We now describe and analyze each candidate justification. Where it makes sense we group justifications into logical themes.

(a.i. through a.iii.) “*To reduce or eliminate the costs associated with a possible future security incident.*”

A typical model in this category compares the predicted costs of a future security incident with the cost of the security measure that would prevent that incident. Adkins (2004) proposes an enhanced model in which predicted costs are represented by a range of possible values and an accompanying probability distribution function. The Adkins model also incorporates the time value of money and the possibility that more than one security incident might occur within the specified planning window. Bodin (2008) propose a methodology that allows decision-makers to combine a number of risk measures into a single composite metric. Below, we develop a simple model that is representative of the more common approach. The model employs ALE (Annualized Loss Expectancy), typically a central component in models of this type.

Generally, ALE is defined as the probability of a potential future security incident (P_i) multiplied by the cost of the event (C_i). A security incident costing \$1M with a probability of 0.4 has an ALE of \$400K:

$$\text{ALE} = P_i \times C_i \tag{2}$$

From (2), ALE can also be calculated as the sum of multiple potential future security incidents:

$$\text{ALE} = \sum_{i=1}^n P_i \times C_i \tag{3}$$

Let Q represent the total cost of the security measure expected to nullify the ALE. The net benefit of the security measure can then be expressed as:

$$\text{Net benefit of security measure} = \text{ALE} - |Q| \tag{4}$$

If the net benefit of the security measure is positive, i.e. if $\text{ALE} - |Q| \geq 0$, then the ALE methodology indicates that the security measure should be pursued. ALE models enable a type of cost-benefit analysis where the “benefit” is expressed in terms of avoiding loss. This potential justification of eliminating or reducing future costs is attractive perhaps because common security technologies such as firewalls, anti-virus software, and intrusion detection systems all function in a defensive, preventative role.

The challenge with models of this type is in estimating the probability that a future security incident will occur (P_i), the magnitude of the associated costs (C_i), and the effectiveness of the security measure. The ALE methodology itself has also been criticized for conflating high probability events that have a low impact with low probability events that have a high impact. ALE-based models can also become overly complex where they attempt to identify and describe all potential future security incidents. Models that are more sophisticated than the one we have described would also require additional variables such as estimates of upper and lower-bounds for potential losses, the number of potential security incidents and their probability, etc.

There is a strong contrast between the ability to gather data regarding security incidents and the data gathering and modeling that takes place in conventional insurance markets. To calculate the risks to a house, an insurance company studies the actuarial data relating to the type of house and its location. This includes whether the house is constructed primarily from brick or wood, whether it sits in a floodplain or tornado alley, and so on. Analyzing historical data and extrapolating those trends

is a scientific approach to the task of calculating the probability that a house might be destroyed. Anderson (2008) notes that many of the data regarding information security failures are, quote: “not only poor but are collected by parties such as security vendors or law enforcement agencies that have a vested interest in under or over-reporting.” Ryan (2003) analyzes fourteen security surveys that were widely publicized from 1995 to 2000, finding those surveys to be replete with design errors in the areas of sample selection, the form of questions asked, and underlying methodologies. Seven of the fourteen surveys ask if the organization at which the respondent works has a security policy. The results range from 19% in one survey to 83.4% in another. Eight surveys ask whether there had been unauthorized access to systems. The results range from 4% in one survey to 58% in another.

Only recently have data on security incidents begun to flow into the public domain in a structured manner. This has occurred primarily because of mandatory breach disclosure laws such as California Senate Bill 1386 (see Shostack, 2008). Prior to mandatory breach reporting laws, businesses had no reason to publicize the fact that their security had been compromised. In the absence of mandatory reporting or the belief that reporting is the right thing to do, it was perhaps predictable that companies would not publicize security breaches because of worries about reputation.

A complicating factor however, is that the technology landscape is changing much more rapidly than the techniques of home building. A house can be expected to have a wood, stone, or steel structure and there are a relatively small number of choices for how the ceilings and interiors can be constructed. The risks to houses are also well known. Insurance companies only need to track a relatively small number of variables. In contrast, the technological landscape that affects the ease or difficulty of attacking or defending computer systems fluctuates as new technologies are introduced, as technologies are configured and reconfigured, and as the security of computer systems decays over time (see Muffet, 1995). Breach data might become stale where the technology involved in the breach incident becomes obsolete.

Parts a.i. through a.iii. of our taxonomy describes three components of potential future losses. These are the cost of ‘break/fix’ repairs, damage to reputation, and litigation costs. We will describe each in turn.

Break/fix repairs represent the costs to rebuild or replace computer systems damaged (or made untrustworthy) in a security incident. Given the type and degree of an incident, the number of computers affected, and the amount of person hours required for recovery efforts, it is reasonable to think that these costs could be estimated.

The literature has an evolving view of the effect of security breaches on the market value of companies that have suffered a breach. The work carried out to-date has predominantly been event studies. An event study measures the price of a stock surrounding an event in order to capture the effects of the event. If a company suffers a breach then the potential for losses such as fines, penalties, and damage to reputation should become reflected in its stock price.

Using an event study, Cavusoglu (2002) identifies a negative effect on CAR (Cumulative Abnormal Return) for companies that suffer a security breach. Campbell (2003) finds that only security breaches that affect confidentiality have significance, with the impact of breaches not relating to confidentiality having no statistical significance. A study of DoS (Denial of Service) type attacks by Hovava (2003) finds that such attacks do not create any significant loss in value for the companies who are targeted. Acquisti (2006) compile a broad data set where PII (Personally Identifying Information) was comprised due to the failure of a security mechanism such as in a computer hacking incident or because of stolen or lost equipment. They show that

such privacy breaches have a negative effect on the stock market value of NYSE and NASDAQ-listed firms, but that the effect is relatively small and short-lived.

The third potential area where costs might occur as the result of a security incident are litigation costs. Litigation costs are those that would be incurred if a security incident caused a customer or some other party such as a government body to pursue litigation against the organization that suffered the breach. In the 1945 case *United States v. Carroll Towing Co.*, Judge Learned Hand defined a precedent-setting formula for determining liability in cases of negligence:

“...to provide against resulting injuries is a function of three variables: (1) The probability that [the harmful event will occur]; (2) the gravity of the resulting injury (3) the burden of adequate precautions. Possibly it serves to bring this notion into relief to state it in algebraic terms: if the probability be called P ; the injury, L ; and the burden, B ; liability depends upon whether B is less than L multiplied by P : i.e., whether B less than PL .”

Note the similarity between the formula described by Judge Learned Hand and (2) – the generalized formula we have described for ALE. Crucially, both rely on the ability to quantify the probability of the event and the magnitude (or cost) of the “injury.”

Given this, a manager considering the security spending decision might arrive at the conclusion that even if there is no evidence that a particular security measure is effective, but the general consensus is that the use of the security measure represents “adequate precautions,” then it might make sense to deploy the security measure for that reason alone. We will speak in §6 about how this idea can be used to create a ceiling on spending as part of a spending strategy.

To summarize, reducing or eliminating the costs associated with a possible future security incident is a common justification for spending on information security. Recent work has identified that the damage to reputation caused by the disclosure of a security breach is small or short-lived. There are a lack of data that could be used to calculate the probability that a security incident might occur. The rapidly changing nature of the technological landscape also poses challenges of data collection and raises questions of how useful such data would remain over time.

(a.iv.) *“To create a deterrent against attacks.”*

Sometimes a business chooses to spend money on a security measure as a visible deterrent, such as a guard at the door. This idea doesn’t carry over to the online world very well. The cost of tracking down and prosecuting an attacker is high. It is a relatively simple task to detain and search an individual as they leave a retail store. The same thing can’t be done with an online transaction – at least not cheaply. The high cost stems from the extensive investigative effort required, support for law enforcement activity, and the difficulty in tracking down an individual who may have laundered their connection through many jurisdictions and countries. A deterrent effect depends on the likelihood and severity of punishment. In the case of online crime, the likelihood of prosecution is very low.

Drawing attention to security measures with the hope of creating a deterrent might conceivably create the *opposite* effect. A boastful computer hacker might be attracted to the “badge of honor” they would gain by compromising the security of a company that is reputed to be well-protected.

(b.i. and b.ii.) “*To enable a new business process.*”

Some business processes carry such obvious risks that without security they could never be launched. An example is online banking. In these cases, security measures might also act as a signal to help attract prospective customers. (Consider the advertising of biotech firms involved in the bioengineering of crops and the “green” advertising of oil companies in this regard.)

Let B represent such pure benefits enabled by the security measure. (4) now becomes:

$$\text{Net benefit of security measure} = (\text{ALE} - |Q|) + B \quad (5)$$

As with the components of the ALE calculation, the challenge is in determining how to quantify B . One approach would be to say that if the new business process generates \$1M per annum and the spend on the security component comprises 10% of the \$500K spent to enable and operate the business process, then the “return” from the security component is: $0.1 \times \$1\text{M} = \100K ; \$50K net. But this approach is incorrect where a business process generates \$1M but \$50K has been *unnecessarily* spent on security. Referring back to (1), the first consideration should be whether *any* spending on security is required.

It does seem however, that there are certain types of business activities that require security. We will explore this idea further in §6 as we construct various models.

(b.iii.) “*To meet a standard required to carry out business.*”

Some large organizations require that their business partners provide evidence that their processes include certain security activities, or that certain security technologies are deployed before allowing them to connect into their environment. More common are externally imposed standards such as PCI – the Payment Card Industry standards. Under PCI, if a company wishes to accept credit card payments then it must comply with the PCI standard.

A fast-growing thicket of laws concerning security, privacy, and operations regulates corporate behavior. Complying with the law has been one of the fastest-growing reasons to spend on security in recent years. Some of these laws impose fines and even jail time for noncompliance, so compliance spending has been viewed in some cases as “spending to keep the CEO out of jail.”

Information security law in the United States is mostly sectoral, covering finance and health. The Sarbanes-Oxley act, often simply called SOX, is the broadest of these laws. SOX imposes new levels of due care in record keeping by public companies.

Kim (2006) reports a study in which 25 measures of performance were reviewed across a variety of organizations. Organizations were then examined that followed some of the 63 practices defined by COBIT – a common framework used by organizations to structure and guide their SOX compliance programs. Kim found that 21 of those 63 practices defined by COBIT had an effect on performance with 4 being key. Therefore, a full two thirds of the practices recommended within COBIT provide no measurable benefit beyond enabling the stamp of SOX certification for the organization. An organization should not expect that such standards will necessarily provide any benefit beyond their ability to enable certain types of business activity.

(b.iv.) “*To create competitive advantage.*”

Most organizations do not develop their own security technologies in-house. The

result of this “buy rather than build” philosophy is that most organizations employ nearly identical sets of security technologies in almost identical ways. Consequently, nothing stops competing organizations from acquiring the same functionality, either from the same vendor or from another vendor whose product has similar features. Security technologies are widely available for purchase on the open market.

Without some unique way of using a security technology, it seems unlikely that a company would be able to create a sustainable competitive advantage using that technology. Even with a unique technology or way of using a technology, it may be challenging to explain the benefits to prospective customers. In comparison, a sustainable price advantage needs no explanation. See Carr (2004) for a ranging discussion of commoditization in IT.

(b.v.) “*To achieve new functional capabilities.*”

When security technologies are implemented they often provide a new level of visibility into the technology environment and control over the environment. Security technology that protects against DoS attacks has valid applications in other areas of IT such as bandwidth monitoring and shaping, traffic monitoring, and service-level monitoring. Cavusoglu (2005) show that an organization that deploys an IDS (Intrusion Detection System) benefits from being able to use the IDS for assisting investigations, irrespective of whether the organization realizes a positive or negative value from the IDS.

Security capabilities can be used to enhance IT operations and vice-versa. Inventory management is useful for understanding what needs to be defended and for prioritizing security efforts. Without change management it is difficult to determine if a change to a system was authorized or malicious. Configuration management can be used to ensure that systems are configured according to their functional and security needs. Operational and security concerns are an integral part of each of these activities.

When an organization builds and standardizes mechanisms and processes that address security concerns, it increases its ability to react to change. For example, the joining of two IT infrastructures as the result of a merger or acquisition is easier and quicker if established security mechanisms are available that are well understood and easy to deploy. In this way, security controls promote flexibility.

When a company spends money on security, it can often reap both security and other functional benefits. But those benefits can be difficult to untangle from each other. The question exists as to whether such functional benefits are generated by the security measures themselves, or whether they emerge from the policy and pursuit of standardization.

(c.i.) “*To improve operational efficiency.*”

A security bug in an application is essentially a software bug with a security *implication*. In that sense, the quality of security in applications is dependent on the quality of their engineering. Soohoo (2001) measures the effort required to fix a security bug in the design phase of a software project compared to later phases such as implementation and testing. The results are shown in Figure 2.

The earlier that security bugs are found in the development process, the less costly they are to fix. This is a tenet of the software testing field – that it becomes progressively more expensive to fix bugs in a software project as it moves closer to release. As such, the finding by Soohoo is a reaffirmation of established software

<i>phase of software lifecycle</i>	<i>ROI</i>
design	21%
implementation	15%
post-implementation testing	12%

Figure 2: Percentage ROI by phase of testing for detection of software security bugs.

engineering principles.

@stake (2001) perform a study in which they measure the performance of a web server running on a particular computer. The default configuration of the web server software was used, as was a default installation of the underlying computer operating system. The number of web pages that could be served within a specific time period was measured. A number of improvements to the security configuration of the web server and the computer were then made, and the number of web pages that could be served within the same time period was again measured. Several variations of security configuration were tested and it was found that improving the security configuration increased performance in all cases, up to a 3.28% gain in web server performance.

This finding can be explained due to the specific nature of the security configuration changes made in the study. To reduce the ‘attack surface’ of the underlying computer, the number of network services were reduced to only those that were required for the web server to function. In a similar manner, web server functionality that was present by default but was not being used was removed. These techniques benefit security because they reduce opportunity for an attacker. It can not be an unexpected result however, that removing functionality makes a computer or web server run faster.

5 Resulting spending strategies that create inefficiency

In the preceding text we have described a number of weaknesses in common justifications for spending on information security. Many of those weaknesses have their root in an absence of actuarial-like data that would allow the value of security measures to be quantified. Such data would have to describe the type, quantity, and details of security incidents that occur. An understanding would also have to be gained as to which security measures are effective and to what degree. Principle-agent problems, psychological effects, and other structural challenges also complicate the security spending decision.

The most efficient situation is where every organization understands whether they should spend on information security or not, and from there how much they should spend and on what security measures. Because that situation does not exist, every organization today experiences some degree of inefficiency in its spending.

Faced with an inability to determine the need to spend or to quantify how much to spend, organizations have turned to the advice of “experts” and to herd-like spending strategies. Gartner, a well-known information technology research and advisory firm make the recommendation in their 2005 report *What Your Organization Should Be Spending for Information Security* that 3-6% of corporate IT budget should be spent on security. Another advisory firm, The Meta Group, recommend in their Diamond

report #2856 (*Recommendations on how much of IT budget should be allocated to security spend*) that banks should spend 8% of IT budget on security and that manufacturing companies should spend 3%. It is unclear how these numbers were arrived at.

A second strategy that allows an organization to offload its security spending decision onto an external party is the pursuit of “best practices.” Organizations should consider where “best practices” come from: they are dictated by consultants, vendors, and the security industry as a whole. Each of these groups has a vested interest in the spending decisions made by organizations, and anyone can call their advice a “best practice.” A quick search shows that the top Google results for the phrase “security best practice” include a small consultancy whose web site is “copyright 2003”; the SANS institute; Cisco; three general IT industry magazines; two university web sites with grab bags of ideas organized quite differently; and a Yahoo! page that includes advice such as “Ensure you’re patched.” “Best practices” are designed to be vague enough to apply in the general case. From the perspective of an individual organization they are inefficient by their very nature because they are unlikely to match the specifics of the technology environment, business goals, and practices of the organization.

The third resulting spending strategy that we will note is where an organization hires a CISO (Chief Information Security Officer) or CSO (Chief Security Officer) and relies on that individual to own the security spending decision. This approach exacerbates the potential agency problems described in §2. Handing decision-making ability regarding spending to the individual with the greatest incentive to spend increases the possibility of inefficiency.

Spending strategies such as setting levels of spending according to some percentage of IT budget, relying on “best practices,” or by trusting in a CSO push organizations towards similarity in their levels of spending (“7% of IT budget”). Such strategies therefore have the perverse effect of *creating* inefficiency. This must be true unless all organizations have an identical optimal level of spending. Figure 3. shows a conceptual model in which companies are mapped against their spending on information security. The distribution is shown as Gaussian, or normal. The variance of the distribution is low, representing the effect of pressures exerted on firms towards uniformity in their spending (i.e. towards the mid point μ).

The more uniformity driven by spending strategies such as the pursuit of “best practices” and “spend 7% of IT budget on security”, the greater the probability that any single company within the distribution is spending inefficiently. For this reason, we refer to the mid point of the distribution as the “unenviable middle.”

6 Practical spending strategies

“Half the money I spend on advertising is wasted; the trouble is I don’t know which half.” – Attributed to John Wanamaker

To compensate for the lack of data that would enable an organization to determine whether it is over-spending or under-spending, a number of models have been published that incorporate uncertainty through the use of sensitivity analysis, probability distributions, confidence intervals, regression analysis, simulations, Monte Carlo methods, and other techniques. Many of these models require graduate-level understanding (or higher) of the mathematics that underpins them. The scope of such models can also be quite narrow and not sufficiently broad to address the security

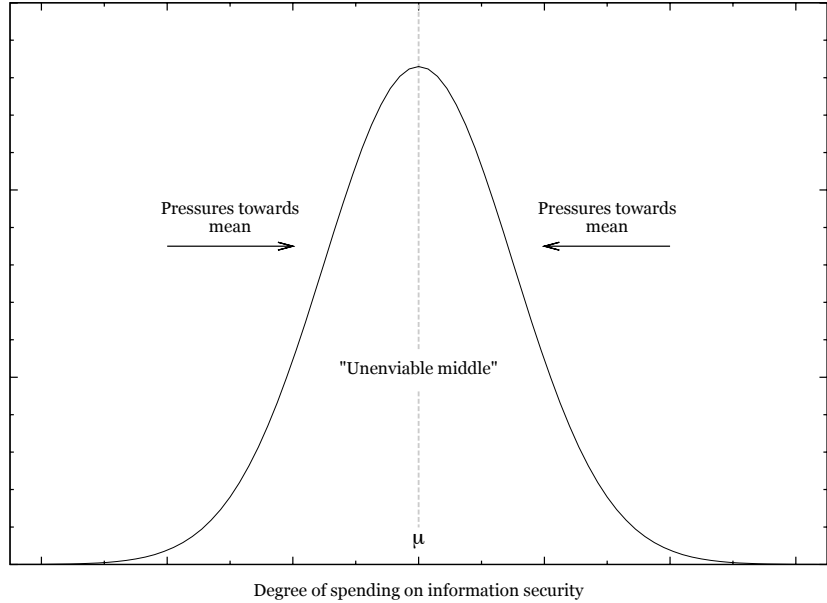


Figure 3: Normal distribution of spending levels.

spending decision at the unit of an entire company. As such, practitioners and organizations might not find them to be especially *practical*, and this likely hurts their broad adoption. The observation that these models are generally rather complicated is also made by Willemson (2006). Below, we propose two simple spending strategies that can be used in a straightforward and practical manner by an organization or practitioner. Both strategies take into consideration the earlier observations within this paper.

The logic of the first strategy is as follows: until data that would support decision-making are more widely available, or until practical models are available that enhance decision-making in the absence of such data, the rational approach is a “wait and see” strategy. When this strategy is employed, no spending on security takes place until it becomes self-evident to the organization that it is required. This strategy can also be referred to as a “minimal spend” strategy.

Since no spending occurs until it is known that it is necessary, the strategy reduces the potential for wasteful spending. It might also help to address the principle-agent problem where the individual responsible for the security spending decision has the incentive to pursue a “visible achievements” (high spend) strategy. In this way, a minimal spend strategy can act as a control system to dampen the individual interests of the security officer. Gordon (2003) describes how control systems for investments in information security will likely result in a net benefit to the organization.

As described by (4), a common justification for security spending is that it minimizes the possible future costs that would be felt if a security incident were to occur. In that respect, security spending can be considered a type of insurance or anticipation cost, with the hope that the net effect is positive to the organization in the long run. With no data available that could be used to predict these possible future costs, the rational strategy is to either not pay such anticipation costs or to seek to minimize them. The minimal spend strategy therefore prescribes that the spending

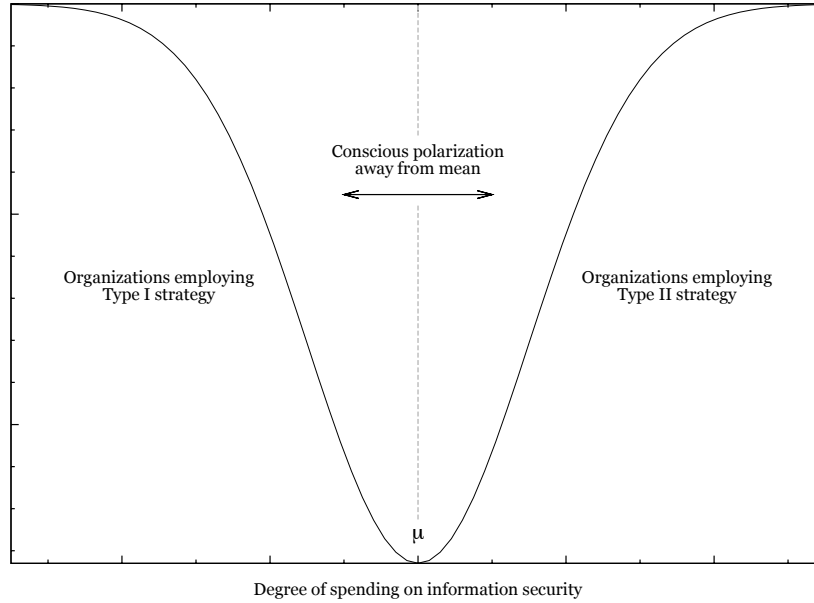


Figure 4: Spending under Type I and Type II strategies.

that the organization knows it *must* carry out - such as to meet legal regulations or to achieve a certification required to do business - should be carried out with the goal of spending the minimum required. That minimum represents a ceiling or maximum on security spending within the organization.

This strategy need not be rigid; it can be employed in a manner that evolves over time. As new data become available for decision-making, they can be used to develop practical models and rules of thumb that can be put to use, loosening the minimal spend approach. For example, the model presented by Konsynki (2008) for identifying and managing perverse incentives can be applied to the asymmetric information problem that exists between the CSO and the CFO who acts as the gatekeeper for spending within the organization.

In §4 we noted that there do appear to be certain types of business activities that require security in order for them to operate successfully, the canonical example being online banking. Given this, we propose a second spending strategy. We refer to this second strategy as a Type II strategy with the previously described strategy being referred to as Type I. An organization first considers whether a Type I strategy is applicable, given its unique circumstances. If the organization determines that a Type I strategy is not appropriate, it approaches the security spending decision with the goal of optimizing security spending for the specific needs of the organization. A Type II strategy requires consideration for whether a Type I strategy can be employed, thereby ensuring that the ‘order of operations’ from (1) is enforced. The use of a Type II strategy also requires a conspicuous rejection of the aforementioned spending strategies such as “spend 7% of IT budget on security.”

Figure 4. shows the pattern of spending on information security when both strategies are employed. The right hand portion of the curve represents the distribution of organizations pursuing a Type II strategy. The left hand portion represents organizations pursuing a Type I (minimal spend) strategy. In concert, these two strategies

create a conscious polarization away from the mean level of spending; a reversion *away* from the mean. Across both strategies, the mean is considered to represent an inefficient, sub-optimal level of spending for the organization.

7 Conclusion

Spending decisions either create or destroy shareholder value. As such, the security spending decision must be carefully considered.

We have described two spending strategies intended to be simple and straightforward for an organization to employ in a practical manner. These strategies account for a number of weaknesses in common justifications for spending on information security. They also take into consideration the observation that a number of pressures push companies towards inefficiency in their spending.

An organization faced with the security spending decision should first consider the necessity of spending. Only then should it determine *how much* and *on what* to spend. If this ‘order of operations’ is not strictly followed, inefficient spending is the probable result.

The Type I spending strategy is a minimal spend strategy. It employs an initial ceiling that represents the spending that the organization knows it must carry out, such as to meet compliance with an externally imposed standard. The Type II spending strategy first requires an organization to identify that a minimal spend strategy is not appropriate. The organization then attempts to optimize its pattern of spending for the specific combination of regulatory, technological, strategic, and other circumstance it faces. This requires the conspicuous rejection of strategies that create similarity in spending between organizations, such as “spend 7% of IT budget on security.”

This idea of a minimal spend strategy will likely create a good deal of cognitive dissonance in the minds of security professionals who hold spending on security as convergent with the best interests of the organization. But for the reasons we have described, if spending on security cannot be justified then it cannot be considered *rational*.

8 Acknowledgements

Many thanks to Professor Benn Konsynski of Emory University. Thanks also to Nick Nedostup and Atholl Stewart for their early input on the topic, and to Professor Steve Stuck of Emory University for assistance with the Gaussian curves.

9 References

Much of the work referenced in this paper was first presented at WEIS, the *Workshop on the Economics of Information Security*. See: <http://econinfosec.org>

@stake, “Defined Security Creates Efficiencies,” *Secure Business Quarterly*, Vol. 1, Issue 2, 2001.

Alessandro Acquisti, Allan Friedman, and Rahul Telang, “Is there a cost to privacy breaches? An event study,” proceedings of the *Workshop on the Economics of*

Information Security, Robinson College, University of Cambridge, England, June 26-28, 2006.

- Roger Adkins, "An Insurance Style Model for Determining the Appropriate Investment Level against Maximum Loss arising from an Information Security Breach," proceedings of the *Workshop on the Economics of Information Security*, University of Minnesota Digital Technology Center, May 13-14, 2004.
- Ross Anderson and Tyler Moore, "Information Security Economics - and Beyond," survey paper, available: <http://www.cl.cam.ac.uk/~rja14/>
- Lawrence D. Bodin, Lawrence A. Gordon, Martin P. Loeb, "Information Security and Risk Management," *Communications of the ACM*, Vol. 51, no. 4, pp. 64-68, April 2008.
- Katherine Campbell, Lawrence A. Gordon, Martin P. Loeb, and Lei Zhou, "The economic cost of publicly announced information security breaches: empirical evidence from the stock market," *Journal of Computer Security*, 11, pp. 431-448, 2003.
- Nicholas Carr, "Does IT Matter? Information Technology and the Corrosion of Competitive Advantage," *Harvard Business School Press*, 2004.
- Huseyin Cavusoglu, Birendra Mishra, and Srinivasan Raghunathan, "The effect of internet security breach announcements on market value of breached firms and internet security developers," *International Journal of Electronic Commerce*, volume 9, 2002.
- Huseyin Cavusoglu, Birendra Mishra, Srinivasan Raghunathan, "The Value of Intrusion Detection Systems in Information Technology Security Architecture," *Information Systems Research*, Vol. 16, No. 1, pp. 28-46, March 2005.
- Anat Hovava and John D'Arcy, "The Impact of Denial-of-Service Attack Announcements on the Market Value of Firms," *The American Risk and Insurance Review*, Volume 6, Issue 2, Oct 2003.
- Gene Kim, "Prioritizing Processes and Controls for Effective and Measurable Security," CERIAS Security Seminar, September 20, 2006.
- Benn Konsynski, Fariborz Farahmand, and Mikhail Atallah, "Incentives and Perceptions of Information Security Risks," proceedings of the *Twenty Ninth International Conference on Information Systems*, Paris, 2008.
- Alec Muffet, "WAN-hacking with AutoHack – Auditing security behind the firewall," proceedings of the *5th USENIX Unix Security Symposium*, 6th June, 1995.
- Julie J.C.H. Ryan and Theresa I. Jefferson, "The Use, Misuse and Abuse of Statistics in Information Security Research," proceedings of the *2003 ASEM National Conference*, St. Louis, Missouri, 2003.
- Adam Shostack and Andrew Stewart, "The New School of Information Security," Addison Wesley, 2008.

Paul Slovic, "The Perception of Risk," Earthscan Publications, 2000.

Kevin Soohoo, Andrew Sudbury, and Andrew Jaquith, "Tangible ROI Through Secure Software Engineering," *Secure Business Quarterly*, Volume 1, Issue 2001.

W. Kip Viscusi, "Alarmist decisions with divergent risk information," *The Economic Journal*, 107, November 1997.

Jan Willemson, "On the Gordon & Loeb Model for Information Security Investment," proceedings of the *Workshop on the Economics of Information Security*, Robinson College, University of Cambridge, England, June 26-28, 2006.