

The Statistical Value of Information

Luther Martin

Voltage Security
1070 Arastradero Rd Suite 100
Palo Alto, CA 95123, USA
martin@voltage.com

Abstract. Accurate estimates of the value of business information are difficult to obtain, but we may get a reasonable estimate for it by using a methodology similar to the value of statistical lives methodology that is commonly used to assess the efficiency of government regulations. Estimates obtained in this way provide insight into the information security market and may explain why some technologies that could be used to protect the information infrastructure are not widely deployed.

1 Introduction

There is a striking parallel between the problem that governments face in deciding which health and safety regulations to enact and the problem that businesses face in deciding which information security technologies to deploy. In each case scarce resources need to be allocated to competing programs, each of which provide different benefits, but also require costs to be paid to realize the benefits. In each case it is reasonable to require that the benefits gained outweigh the costs that they require. The problem of measuring the efficiency of government regulations has been widely studied, and we should expect that the insights gained from that analysis to provide useful insights into the tradeoffs involved with information security technologies.

2 Justifying the Cost of Government Regulations

The analyses that governments perform to establish the effectiveness of health and safety regulations are typically based on the “statistical lives” that will be saved by the regulation, because while it is possible to predict the average number of lives that will be saved in a large population, it is typically impossible to identify the individual lives that will be saved. Historically, this was generally done by measuring the “cost of death” by adding the cost of medical expenditures associated with an incident and the income loss associated with the death. In 1982, the US Occupational Safety and Health Administration (OSHA) proposed a hazard communication regulation (warning labels, etc.) that was initially rejected by the Office of Management and Budget because the costs associated with the regulation exceeded its benefits. In the

appeals process for this regulation, an alternative to the cost of death methodology was adopted. One study estimated that the cost of death methodology underestimated the actual costs associated with regulations by roughly a factor of 10 [1].

In this alternative methodology, the benefit of a regulation is defined by the value of lives saved instead of the cost of deaths avoided. US federal agencies routinely assess the benefits of regulations using statistical value of life data, but the results of such analyses do not always dictate government policy, and regulations still get implemented that have extremely high costs per life saved. Much of the subsequent work in this area has used the value of lives framework for their analyses.

To estimate the value of a statistical life (VSL) using this methodology, we estimate the value of a life by the willingness of people to pay to avoid a death. In a hypothetical example, suppose that a new regulation reduces the annual risk of death by 1 in 10,000, so that we expect to save one life in a population of 10,000. If the average person in this population is willing to pay \$500 to avoid this risk, then a population of 10,000 is willing to pay \$5 million to save this life. In this case, we say that the VSL is \$5 million. We can also calculate the VSL by dividing the willingness to pay to avoid the risk by the probability of the risk, or \$500 divided by 0.0001 in this example.

Studies have shown that reliable estimates for a VSL to range from \$4 million to \$9 million [2], so government regulations that cause costs of less than that rough range per life saved are easier to justify than regulations that cause costs of more than that approximate range per life saved. And even though people may be hesitant to place a monetary value on a human life, we can still infer the value they place on a life through market mechanisms like the additional salary required to accept a chance of death in the workplace, the higher prices that safer cars command, or the lower values of houses near sites known to contain toxic wastes. The VSL approach provides a framework to interpret the choices that people make in situations like these.

3 Justifying the Cost of Information Security

The difference between the cost of death and value of life methodologies for determining the cost-benefit value of a government is somewhat analogous to the difference between justifying the cost of deploying an information security technology based on either the cost of a security breach occurring or the value of the data that is being protected by the security technology. The cost of death methodology corresponds roughly to justifying information security technology based on the cost of the security breaches that technology would prevent. This is the basis for the Annualized Loss Expectancy (ALE) methodology [3] that is currently widely used in risk management. The value of life methodology, on the other hand, corresponds roughly to justifying information security technology based on the value of the data that the technology would protect. Because the value of life methodology has proven to be more useful than the cost of death methodology in understanding the effects of government regulations, we might expect an analogous methodology to provide useful insights into the choices made when selecting information security

technologies to deploy. We call the model developed in this way the “statistical value of information,” (SVI). To find the SVI for a particular information security technology we divide the amount that businesses demonstrate that they are willing to pay for the technology by the probability that the technology reduces the chance of information being disclosed.

In a hypothetical example of the SVI, suppose that deploying an information security technology reduces the annual rate of a particular kind of security breach by 1 in 10, or a probability of 0.1. If the total cost of ownership (TCO) of this information security technology is \$100 per user per year, then we will feel justified in deploying this technology if the value of the data that it will protect is \$1,000 (the TCO of \$100 divided by the probability 0.1) or more. Often, however, we do not know the value of information. In this case we can infer the SVI if we know how effective a technology is at eliminating security breaches as well as the TCO of the product. From this information we can estimate the SVI, much like we can estimate the VSL from observing the choices that people make. Much like government regulations, some information security technologies seem fairly easy to justify in this way while others seem to require an extremely high value for the protected information for their justification.

An example of the SVI that we can easily estimate is that of the data on a laptop computer. It is possible to purchase insurance that covers both the theft and damage to a laptop computer for approximately 6 percent of the cost of the computer itself,¹ and the chances of a laptop computer being stolen and the data on it compromised is probably close to 3 percent per year.² Anecdotal evidence suggests that the TCO of a whole-disk encryption product to be approximately \$150 per user per year. This gives us an estimate of about \$5,000 (\$150 divided by 3 percent, or a probability of 0.03) for the SVI that whole-disk encryption reflects. This is much lower than the much-higher (close to \$1 million) estimates than some informal surveys return [4], but it reflects the anecdotal evidence that the data on a laptop computer is worth more than the computer itself. Both the relatively low rate of laptop theft and the low rate of businesses deploying whole-disk encryption to protect the data on laptop computers indicate that the SVI estimate of \$5,000 is probably more accurate than the estimates from surveying users. If the SVI of data on laptop computers were actually closer to \$1 million, we would expect laptop theft to be rampant or whole-disk encryption to be ubiquitous, neither of which we observe today.

A more interesting example is that of e-mail encryption. Most e-mail encryption products rely on the digital certificates created by a public-key infrastructure (PKI) of some sort, and the TCO of the PKI keeps the cost of implementing most e-mail encryption technologies fairly high, and most of the cost of deploying and supporting e-mail encryption is due to the use of PKI. Because of this, we will assume that there are minimal additional costs for e-mail encryption beyond the costs for PKI. This will give us a low estimate for the SVI for e-mail, but one that is good enough to explain the adoption pattern for e-mail encryption.

¹ Computer loss and damage insurance is offered by Safeware (<http://www.safeware.com/>), for example.

² Although they do not give detailed data on computer damage and theft rates, Safeware claims that accidental damage causes more loss than theft of computers does.

Estimates for the TCO of user-operated PKI range from \$157 per user per year [5] to \$220 per certificate [6]; using \$200 per user per year is probably a reasonable estimate for the TCO of e-mail encryption that typical organizations will experience. The other information needed to estimate the SVI that e-mail encryption reflects is difficult to obtain. It is extremely difficult to get an accurate estimate for the chances of the information in unencrypted e-mail being compromised, but we can get an order-of-magnitude estimate for this probability by comparing the chances of routine causes of death with the chances of eavesdropping on e-mail.

Usability issues with e-mail encryption have led to most e-mail being encrypted at a gateway appliance instead at a user's desktop, so in its typical use, encryption of e-mail messages only protects against eavesdropping outside corporate networks, as messages are passed through the Internet, where the actual chances of a message being observed and read are fairly small. Table 1 lists a selection of probabilities of death due to injury from various causes. Based on this data, an estimate of roughly 1 in 10,000 for the chances of e-mail messages being subjected to eavesdropping attacks seems to be fairly reasonable, although it might actually be rarer. An informal survey of security experts revealed that most could list a friend or acquaintance who had either died from either intentional self-harm or as a car occupant, while none could list a single case of eavesdropping of e-mail outside corporate networks. Based on this, an estimate of 1 in 20,000 may be more appropriate.

Table 1. Chances of Death Due to Injury, United States, 2002 [7]

Type of accident or manner of injury	Annual probability
Chance of death due to any injury	1 in 1,755
Intentional self-harm	1 in 9,096
Assault	1 in 16,325
Car occupant	1 in 17,625
Complications from medical care	1 in 101,281

Using 1 in 10,000 as the estimate for the probability of a compromise with the estimate of \$200 per user per year for the cost e-mail encryption, we get an estimate of roughly \$2 million (the TCO of \$200 divided by the probability 0.0001) for the SVI that e-mail encryption reflects. While it is possible that the data on a laptop computer is worth \$5,000 to a business, it seems unlikely that the information contained in typical business e-mail is worth \$2 million, and the slow adoption of e-mail encryption by businesses may reflect this. The cost of encrypting e-mail may be much too high, particularly when we consider the relatively small chance of eavesdropping in typical situations.

Research has established that VSL is not constant throughout heterogeneous populations. In particular, VSL has been shown to be related to income, marital status and age [8]. People with higher incomes and those that are married tend to place a higher value on a statistical life than those with lower incomes or those that are unmarried. The relationship between VSL and age for adults tends to follow an inverted U-shape, increasing up to roughly age 40 and decreasing thereafter, while the VSL for children is typically higher than that for adults. VSL estimates usually reflect the preferences of workers with an average age of roughly 40. Some regulations tend

to reduce risks for the very old or the very young, so it may be appropriate to use different VSL estimates for different regulations.

Similarly, the SVI may vary by industry or by the position of the owner of information. Businesses in the financial services or health care industries may deal with information that has a higher SVI than information in other industries. Although the information that they have may not be exceptionally valuable, government organizations tend to be risk-averse. In these cases, a higher willingness to pay to protect information is reflected in a higher SVI. Thus financial services, health care and government are more likely to invest in information security technology. Executives and human resources workers may deal with information that has a higher SVI than other information within their organizations, so we should also expect to see investment in information technology for these workers before we see widespread deployment of the technology.

4 Risk-Risk Tradeoffs

There are additional parallels between efficient regulation and the efficient use of information security technology. In particular, in many cases, implementing a government health and safety regulation may reduce one type of risk while increasing another. When performing a cost-benefit analysis for government regulations, the risk-risk tradeoffs for regulations may be significant, and the unintended consequences of reducing one risk may actually increase the exposure of the overall population to risk.

There are four types of risk-risk tradeoffs that regulations may introduce [2], and each of them has an analogous situation in information security. First, a regulation may reduce one risk while increasing a different risk. Banning saccharine, for example, may have reduced some health risks due to the exposure to the saccharine, but may also have increased health risks to others due to obesity caused by substituting sugar for saccharine in some diets. Similarly, using a particular information security product may introduce new vulnerabilities even as others are reduced. For example, requiring that all information security products be Common Criteria certified and operating in an evaluated configuration may decrease some security risks while increasing others. This happens because the inflexibility of the Common Criteria does not allow users of certified products to install patches or software updates and stay in an evaluated configuration. This leaves deployed systems exploitable by any new vulnerabilities that are discovered since the completion of the Common Criteria certification. Or an information security product may fulfill its role perfectly, but also introduce exploitable buffer overflow vulnerabilities.

Another risk-risk tradeoff comes from individuals increasing their risky behavior because they feel safer when protected by risk-reducing technology or the effects of risk-reducing regulations. Studies have shown, for example, that requiring drivers to wear seatbelts or drive cars with air bags does not decrease the fatalities from automobile accidents because drivers tend to feel safer and thus engage in behavior that is more risky than they would in the absence of the seat belts [9] or air bags [10].

Similarly, it may be the case that computer users who are protected by anti-virus software feel safe from computer viruses and tend to be less careful with dangerous attachments to e-mail than they would in the absence of anti-virus software.

Implementing ways to reduce risk may also result in activities that increase risks more than the original risk is reduced. Regulations that require new construction are an example of this, because the activity of construction is often more dangerous than the risk that is reduced by the results of the construction. Deploying information security technologies can also introduce new vulnerabilities in a similar way. Giving consultants or other contractors access to corporate networks carries the risk that they will use their access to carry out malicious activity or to otherwise subvert the networks to which they have temporary access, for example.

Finally, investing limited budgets to reduce risk in one area means that the same funds are not invested in reducing the risk in other areas, even ones that provide a greater reduction in risk. In the case of information security, it appears that deploying e-mail encryption technology before deploying whole-disk encryption may do this, as will investing in information security technologies when there are other areas where a limited budget can be spent to reduce other risks.

5 Overcoming the High SVI Required by Some Technologies

The requirements that businesses have for software are different than the requirements of home users. Businesses require interoperability with complex legacy systems, while home users typically only look for compatibility with a small number of simple operating environments. Businesses also require around-the-clock availability of technical support, both before and after a sale, and the detailed testing process through which many business information security purchases need to pass often require vendor support at each step of the way. On-site visits by vendor representatives to customer sites are often required as part of the sales process of business software, which require the time of skilled technical staff as well as travel to customer sites. The result of these additional costs is that there is probably a lower bound for the price at which business software can be profitably sold and a corresponding lower bound on the TCO for it. For a particular probability of loss, this gives a lower bound for the SVI that information security products designed for business use can efficiently protect.

In a hypothetical example, suppose that the costs of developing, selling and marketing business software put a lower bound of \$50 per user per year on its TCO. If the use of this technology reduces the probability of a loss by 1 in 10,000 per year, this gives a lower bound of \$50,000 (the TCO of \$50 divided by the probability 0.0001) for the SVI that this technology can efficiently protect. If the SVI is greater than \$50,000 then using this technology is justified. If the SVI is less than \$50,000, then using the technology will require costs that exceed its benefits, and we should not use it. If the SVI is exactly \$50,000, then we should be indifferent to using the technology or not using it. We call this point the “breakeven SVI,” and from above we estimate that the breakeven SVI for whole-disk encryption is roughly \$5,000 and the breakeven SVI for e-mail encryption is roughly \$2 million. Due to the lower limit

of the TCO for business software, there are probably cases where the breakeven SVI for some technologies will always be too high to justify their widespread use.

If the SVI turns out to be too low to justify the use of existing information security technologies, then increasing the SVI will make using the existing technology appealing. One way to do this is to levy fines against organizations that disclose information that should be protected. The recent changes in the regulatory environment of business seem to have adopted this strategy, including those included in the Financial Modernization Act of 1999 (otherwise known as the Gramm-Leach-Bliley Act), the Health Insurance Portability and Accountability Act of 1996, and California Senate Bill 1386. So we can understand the information security elements of these regulations as increasing the SVI in order to increase the adoption of information security technology where market forces have not yet led to its adoption.

It is also possible to change the breakeven SVI by either decreasing the TCO of security technology or by increasing the amount by which the technology protects against compromises. We usually see small incremental changes in the TCO of security products over time as vendors try to use the lowest TCO as a selling feature of their products. This gradually decreases the breakeven SVI, which in turn makes using the technology more attractive. Alternatively, incremental improvements in information security technologies often result in small incremental increases in the effectiveness of the technology. So progress in security technology may eventually decrease the breakeven SVI for many information security technologies. Similarly, an increased threat can also decrease the breakeven SVI. If the probability of e-mail being intercepted increases to 0.1, for example, and the TCO for e-mail encryption remains roughly \$200, then the breakeven SVI will decrease to \$2,000 (the TCO of \$200 divided by the probability 0.1). Of all the ways to decrease the breakeven SVI and significantly affect future information spending patterns, this one seems the most likely.

6 References

- [1] W. Viscusi, "Analysis of OMB and OSHA Evaluations of the Hazard Communication Proposal," March 1982.
- [2] W. Viscusi and J. Aldy, "The Value of a Statistical Life: A Critical Review of Market Estimates throughout the World," Related Publication 03-2, AEI-Brookings Joint Center for Regulatory Studies, January 2003.
- [3] National Bureau of Standards, "Guideline for Automatic Data Processing and Risk Analysis," Federal Information Processing Standard 65, August 1979.
- [4] M. Slocombe, "Symantec: Average Laptop Contents Are Worth Half A Million Quid!," January 30, 2006, http://www.digital-lifestyles.info/display_page.asp?section=cm&id=2960.
- [5] VeriSign, Inc., "Total Cost of Ownership for Public Key Infrastructure," <http://www.verisign.com/static/005321.pdf>, March 2005.
- [6] US Government Accounting Office, "Status of Federal Public Key Infrastructure Activities at Major Federal Agencies and Departments," GAO-04-157, December 2003.
- [7] R. Anderson and B. Smith, "Deaths: Leading Causes for 2002," *National Vital Statistics Reports*, 53(17), March 7, 2005.

- [8] D. Kenkel, "Using Estimates of the Value of a Statistical Life in Evaluating Regulatory Effects," in F. Kuchler (ed.), *Valuing the Benefits of Food Safety*, USDA Miscellaneous Publication Number 1570, April 2001.
- [9] S. Peltzman, "The Effects of Automobile Regulation," *Journal of Political Economy* 83(4), pp. 677-726, 1975.
- [10] S. Peterson, G. Hoffer, and E. Millner, "Are Drivers of Air-Bag-Equipped Cars More Agressive? A Test of the Offsetting Behavior Hypothesis," *Journal of Law & Economics* 38, pp. 251-64, Oct. 1995.