

**From *Enterprise Information Security and Privacy*, edited by C. Warren Axelrod,  
Jennifer L. Bayuk, & Dainel Schutzer (2009)**

## **The Economics of Loss**

**by Scott Borg**

### **Security as the Prevention of Loss**

Security is ultimately about preventing losses. In fact, security could be *defined* as any measures people carry out to reduce their losses from illegal or abnormal events.

Economics drives security. It is the tool we use, with varying degrees of success, to quantify security decisions. We decide if a given security expenditure is warranted by comparing it to the scale of the losses it is expected to prevent. If a security measure costs more than the losses it is expected to prevent, then it isn't worth the expenditure. If a security measure costs less than the losses it is expected to prevent, and we have a high degree of confidence in its efficacy, then it probably is worth the expenditure. We choose among different security expenditures by deciding which ones we think will prevent the most losses per dollar spent. We try to achieve the greatest possible reduction in losses per dollar spent.

Anyone who makes reasonable decisions about security is making those decisions on economic grounds. When people fail to recognize this, it is generally because they think economics is limited to discussions of money and markets. But economics is also concerned with production, its costs and its benefits. The estimates of costs and benefits need not be monetary. They might, for example, be stated in terms of lives lost and lives saved. They might involve the quality of life.

Loss of privacy causes a loss of value from an economic standpoint. This is because privacy creates value, both directly and indirectly. It contributes directly to our quality of experience and our ability to make free choices. It contributes indirectly to our financial and personal security, since losing control of personal information makes us vulnerable to other crimes.

While economics does aspire to precision, it can also deal with situations where little precision is possible. The assessments of the quantities involved in the costs and benefits might be vague and intuitive. They might, for example, involve only a rough distinction between massive, unacceptable losses and moderate, acceptable losses. But if security decisions are justified, they will always be founded on estimates that are quantitative and economic. They will involve weighing one set of costs or losses against another set of costs or losses.

### **Quantifying the Risk of Loss**

The key to quantifying security is the concept of expected losses. We can calculate these expected losses:

- 1) by estimating the likelihood of a given type of security event in a given period of time (the Threat),
- 2) by estimating the losses that could arise from that security event (the Consequence), and
- 3) by estimating the extent to which the event will cause those losses, given the existing security measures (the Vulnerability).

These are “risk factors” in a mathematical sense. We can multiply them together to get the expected loss or “risk.” This gives us the classic risk equation:

$$\text{Threat} \times \text{Consequence} \times \text{Vulnerability} = \text{Risk}$$

In this equation, the Threat will generally be stated as a probable number of incidents per year. The Consequence will be stated as the possible cost per incident. The Vulnerability will be stated as the probable extent to which the Consequence will take place, given the defensive or mitigating measures that have been put into effect. The Risk will then be the annualized expected loss.

This basic risk equation can be formulated in a number of other ways that are mathematically and conceptually equivalent. In place of “Vulnerability,” for example, some people like to substitute “(1 minus mitigation),” where the mitigation is the percentage by which the expected loss has been reduced. In place of three basic factors or categories, some people like to make two basic factors by multiplying one of the pairs together, before multiplying it by the third. “Threat” and “Consequence,” for example, are sometimes multiplied together to yield “Hazard,” which is then multiplied by “Vulnerability” to yield “Risk.” Alternatively, “Consequence” and “Vulnerability” can be multiplied together to yield “Expected Loss Per Incident,” which is then multiplied by the “Threat” to yield “Risk.” There are also various systems for allocating subordinate factors between the three categories, so that everything relevant is taken into account somewhere in the equation.

### **Refining the Basic Risk Equation**

The basic risk equation usually needs to be refined a bit before it can be used to make practical budget decisions. This is mainly because in addition to estimating the probable loss, we need to have some idea of how probable it is that our estimate is correct. There are various methods of keeping track of our degree of confidence in our estimates. These methods result in a range of outcomes for each factor in the risk equation, with a probability assigned to each outcome.

There are also various methods of adjusting for the future value of money, so that costs today can be weighed against costs at different times in the future. These adjustments usually take account of the various interest charges that would be required to spend the same sum of money at various times. The interest rates employed in these calculations need to be adjusted for inflation, since it's the future purchasing power that the company cares about, not the nominal quantity of future dollars. Adjustments for the future value of money can also take account of probable changes in the value of money, due to other factors. Some companies, for example, will want to consider the future value of money in the light of future exchange rates between currencies. The exact ways in which a company needs to adjust for the date at which it will spend money will ultimately depend on the nature of that company's cash flows.

Despite these possible refinements, the basic formula for making sense of risk is still essentially the same. It still comes down to an estimate of our expected losses over a given period of time, for different threats, under different policies.

### **The Problem of Quantifying Loss Itself**

The core of the risk equation is the estimate of loss or "consequence." This is the number we most need in order to turn intuitive security assessments into actual economic calculations. Without some idea of how large the consequence would be, it is not worthwhile to talk about how frequently that consequence will occur or how much it can be reduced by security measures.

The key problem, then, is how to quantify this loss of value. Many publications on security economics get the basic risk equation right, although sometimes using different terminology. Once people know their expected losses under various conditions, it is not too hard to calculate the return-on-investment for various security measures. But most of the existing literature on security economics doesn't explain how to calculate the costs or losses correctly. This is why when people try to make the case for increased security expenditures by quoting some amazingly high loss statistic, they get little response: most such claims have no credibility, because hardly anyone trusts the methods used to produce them.

When these calculations of cost go wrong, they tend to go wrong at the very beginning. The mathematics that's employed is fine. When there is relevant empirical data, the statistical analysis of that data is usually done correctly. But the people doing the calculations frequently formulate the problem wrong, look at the wrong data, and have serious misconceptions about what their results will represent.

To avoid these mistakes, it is necessary to begin by talking about loss of value in a relatively comprehensive sense, rather than jumping directly to something like loss of profits, loss of assets, loss of customers, or loss of capitalization. By starting with a more comprehensive picture, we can take account of all the ways in which losses could be

inflicted. Once we have identified the loss in value more broadly, we can then look at how that loss is felt by different parts of a business operation.

### **Confronting the Reality of Hypothetical Actions**

From a business standpoint, any event that prevents a business from actively creating value needs to be counted as causing a loss of value. If a business is expected to create a certain amount of value in the next financial period, and a security event prevents it from doing so, then the loss involved is the reduction in value creation caused by that security event.

Security professionals sometimes complain that any estimates of the gains or losses from security measures need to be based on comparisons of what happened with what didn't happen. How is possible, they say, to quantify what didn't happen? How can security professionals show how much they have contributed, when their entire contribution was simply to prevent something?

But this complaint is caused by a misconception about how business decisions are made. *All* business judgments are based on quantitative comparisons of what happened with what didn't happen. We judge the effectiveness of a marketing plan, a CEO, a product innovation, or any other business measure by comparing how the business did after employing that measure with how we expected the business to do if it hadn't employed that measure. We assess *all* business measures by quantifying things that didn't happen.

The problem with most of the methods currently employed for estimating loss of value is not that they deal with hypothetical events. The problem is that they are measuring the wrong things in the wrong places. By stopping to look at some of the mistakes that people habitually make when they try to estimate the losses from cyber and physical attacks, we can put this subject on a sounder foundation. Four of these mistakes that people make are especially illuminating. Each mistake, once its fallacies are understood, leads directly to an indispensable component of the correct method.

### **Overcoming the Fixation on Assets**

One of the biggest and most widespread mistakes is to think that what people lose, when they suffer a loss of value, are "assets." This is the notion behind the widespread belief that security programs should be designed to protect a company's assets. Indeed, the idea that security professionals should be protecting assets seems so self-evident that most security planners fail to notice the extent to which this policy ignores most what matters to a business.

In fact, there is little correlation between the value of an asset and the extent to which a business would be hurt by an attack on that asset. Often a large business operation can be shut down for weeks or caused to have huge liabilities by an attack on a piece of equipment

that is in itself of little value. Imagine a physical attack that ruptures an inexpensive piece of piping or a cyber attack that causes a low-cost pressure gauge to give a false reading. If this attack causes a toxic gas to be released upwind of a large city, the result could be harm to thousands of people and bankruptcy for the company. Yet the asset involved might be of little value.

The way to avoid this kind of mistake is to begin by looking at activities a business carries out to create value, not its static assets. Closely related to these activities are the ways a business could be caused to destroy value if it operated defectively. Instead of imagining that value resides in static assets, it is vital to understand that value is something created by the operations that a business carries out to turn inputs into outputs. Except in certain parts of the financial industry, what matters most is the value being created, not the value being stored.

### **Overcoming the Fixation on Market Value**

Another of the biggest and most widespread mistakes is thinking that the value someone loses when they are deprived of something is the “market value” of that thing. This is another misconception that sounds so reasonable to people, they rarely stop to consider how rarely it is true.

Consider what the losses would be to you personally or to your company if you were deprived for a month of some basic product that you ordinarily take for granted, such as your telecommunications, your heating or air conditioning, your computers, your transportation, your medical care, or your food. Would the losses that you suffer be equal to the amount that you ordinarily pay each month for that type of product? Does the amount you pay each month for something like your telephone service represent even a good first approximation of what that service is worth to you? Of course not! Does the fact that diamonds are expensive mean that you would suffer a great loss if you were deprived of them? Unless your main activity is modeling jewelry, it’s not likely. The fact is that “market value” of something has almost no relation to the losses that we would suffer if we were deprived of that thing.

What, then, do we use as a measuring point if we can’t use the market value? The secret is to start with the “indifference points” that people use to decide whether a deal is worth doing.

For a customer, the indifference point is that customer’s Willingness-to-Pay. If the customer is offered a price that is lower than its Willingness-to-Pay, the customer will make the purchase. If the customer is offered a price that is higher, then the customer won’t make the purchase.

For a supplier, the indifference point is that supplier’s Opportunity Cost. If the supplier is offered a price that is higher than its Opportunity Cost, then the supplier will make the sale. If the supplier is offered a price that is lower, then the supplier won’t make the sale.

Both the customer's and the supplier's indifference points are determined by the best available alternative to the deal currently being offered. The customer's Willingness-to-Pay is determined by what that customer could do with the same resources if the customer didn't make that purchase. The supplier's Opportunity Cost is determined by what that supplier could do with the same resources if the supplier didn't make that sale. This means that both the Willingness-to-Pay and the Opportunity Cost are usually definite numbers. If there is a negotiation involved, the customer and the supplier will usually each go into the negotiation with a pretty definite idea of the maximum or minimum price that would be acceptable.

These indifference points mark the endpoints of a unified business activity that turns inputs into outputs. They determine whether that activity is worth carrying out. The fact that we can usually tell whether an activity is worth carrying out means we are already able to estimate, at least roughly, the amount of value created by a business activity.

This business activity is the core of economics. All of the other economic components are defined, at least implicitly, in relation to this business activity. The customers are whoever receives the outputs of that business activity and pays for those outputs. The suppliers are whoever supplies the inputs to that business activity and is paid for those inputs. Markets are the forums in which suppliers are matched with customers, and in which alternative suppliers or alternative customers can be substituted.

The precise value created by any business activity is the Willingness to Pay of the customers, minus the Opportunity Costs of the suppliers.

The actual price in any deal is the point at which the value created is divided between the supplier and the customer. Price isn't as fundamental as Willingness-to-Pay and Opportunity Cost. But it is still vitally important, because an individual company's share of the value being created is determined by the way prices divide up that value. The value captured by the supplier is the price, minus the Opportunity Cost. The value captured by the customer is the Willingness-to-Pay, minus the price.

Although companies will be most concerned with the value that they are capturing, they also need to have some idea of how much overall value they and their trading partners are creating by their deals together. Without this knowledge, companies will not be able to foresee they way prices could change as a result of security events. Hence, without the broader picture of value creation, companies will not be able to estimate their own potential losses.

### **Overcoming the Fixation on Productivity**

Yet another big and widespread mistake is to think that value creation and value destruction can be understood in terms of the productivity of individual components. Those taking this approach usually try to assign a productivity rate to a given type of equipment or system. Next they measure how long that equipment was shut down as a result of the

security event. Then they multiply the productivity rate times the length of time the equipment was shut down. If the people applying this method incorporate a lot of empirical data, they can make it sound as though they are doing something reasonable until someone asks a “dumb question,” such as, “how much value is that piece of equipment over there actually creating?”

Suppose the piece of equipment is a telephone. The amount of value that telephone is being used to create depends entirely on who is using it and what that person is using it for. A telephone that used by the company president to negotiate the company’s biggest deals has a very different level of productivity than a telephone in a storage room that people have forgotten is there. A telephone in the trading room of a financial services firm has a very different level of productivity than a phone in the office that sells money orders and travelers’ checks.

In addition, the amount of value created by a piece of equipment depends on its marginal productivity. If a telephone is the last one put in service or the first one taken out of service, its marginal productivity might be scarcely more than zero. If that telephone is unavailable, another can easily substitute. On the other hand, if the telephone is the last one the company has that is still working, its marginal value might be enormous. In fact, in certain circumstances, the very survival of the company could depend on it.

And that brings up another condition: the value created by a piece of equipment depends entirely on the market environment in which it is being operated. If there is no demand at that moment for whatever the company produces, then the value being created at that moment by every piece of equipment in the company might be almost nothing. A telephone can’t create any value if no one needs to call in or out. Alternatively, if the market conditions are such that every company in that market can sell everything it can produce, then the slightest reduction of capacity caused by any piece of equipment being unavailable might have an immediate impact on the company’s total profits.

These points might sound obvious, but people who try to calculate the value lost or created by adding up the productivity of individual pieces of equipment are ignoring every one of them.

The remedy is to start with the value created by the entire business operation, rather than trying to figure out the value created by adding up the parts. Then, after assessing the larger business operation, it is possible to determine how much that larger operation would be disrupted by damage to one of its contributing parts.

### **Overcoming the Neglect of Substitutes**

A final big and widespread mistake is to forget to pay adequate attention to what substitutes. This is the mistake that is being committed every time someone states what system or activity was lost without also stating what system or activity replaced it. Often a

piece of equipment will be put out of commission, but the people responsible for the operations will find an ingenious work-around that will allow them to continue the operations anyway. Even when the larger operation needs to be shut down for a period of time, the people and equipment will often do something else useful, such as catching up on maintenance, stock inventories, or clerical work. The substitute activities may create considerably less value than the normal activities that were interrupted. But the value they create is usually far from negligible, and it may be a large portion of the value that would have been created ordinarily.

Any estimate of loss that doesn't take account of the substitute activities is likely to be highly inaccurate. Something always substitutes, even though it may be very different from whatever it is replacing.

The remedy is to assume that the loss calculation will always be a relative one. The value lost as a result of an attack can only be calculated by taking the value created before the attack and subtracting the value created after. If this isn't a central part of the loss calculation, then that calculation must be regarded as incomplete.

Interestingly enough, it is often possible to estimate the drop in the value created without knowing the absolute quantity of value created before or after the attack. This is because there will often be an unknown portion of value created that remains essentially unaffected by the attack. In such cases, subtracting the value created after the attack will also subtract the portion of the value created that is unknown.

### **Taking Account of the Duration and Extent of the Effects**

The principles already described give us most of the guidance we need for calculating the value destroyed by an attack. We know how to estimate the value being created before the attack, and how to subtract the value created after the attack.

The total value that is lost, both at the level of the individual company and at the level of the market, depends not just on the degree to which the value ordinarily created by a business is being destroyed, but also on the duration of the destruction. This means that it is vital to identify all the effects that would have lasting consequences: irrevocably lost opportunities, possible damage to business relationships, lasting damage to production capabilities.

For an individual business, one of the effects of a security event that is especially likely to cause persisting losses is damage to business relationships. The strength of these relationships can be measured by the total switching costs that a trading partner would need to pay in order to replace that customer or supplier with another one. Switching costs for customers include things like researching a new product, establishing a new account, educating the new supplier about the company's special needs, training employees in the somewhat different procedures necessary to utilize the new product, adapting or replacing

other products and systems that need to be made compatible with the new product, and accepting the risks of a less familiar supplier. There is a similar, partially symmetrical list of switching costs for a supplier.

A security event causes lasting damages to relationships by forcing the company's trading partners to pay part or all of the switching costs that would allow them to move their business to a new customer or supplier. Even the threat of losing a customer or supplier will cause companies to spend part of the switching costs, so that they will be ready if their previous customer or supplier becomes unavailable to them. If a company needs to shut down its operations for a significant period of time, its trading partners may need to start doing business with an alternative customer or supplier. This may force the trading partners to pay out most of the total switching costs.

Once business relationships have been lost or damaged, the consequent losses can be very long term. Even a slightly damaged business relationship may force the company deemed unreliable to lower its prices. Meanwhile, a lost business relationship may cause the company to lose the total value it would have captured from doing business with that trading partner over the entire period of that relationship. This is a loss that can be estimated by assessing the probable life span of such relationships without any significant security events.

In addition to considering the duration of the destruction, it is also important to trace how far the damage extends. For the individual business, this means looking at potential liabilities. For the larger economy, this means looking at the knock-on effects. It is easy, in assessing losses, to consider too narrow a range of effects and to overlook consequences that are less direct.

### **Distinguishing between the Different Business Categories of Attacks**

This entire analysis of costs needs to be founded on a clear understanding of what security events can do to a business. Here, it is necessary to distinguish sharply between indiscriminant, natural events and highly discriminant, malicious events. Indiscriminant, natural events usually interrupt operations and require extra expenditures to put things back in order. Discriminant, malicious events can harm businesses in a wider variety of ways.

From a business standpoint, there are four things an attack can do:

- 1) The attack can interrupt the business operations.
- 2) The attack can cause the business operations to be carried out in a defective way.
- 3) The attack can discredit certain business operations, so that they are abandoned.
- 4) The attack can undermine the basis for the business operations so that they can no longer be carried out profitably.

In practice, each of these business four business effects of attacks needs to be analyzed somewhat differently, and each results in very different cost curves over time.

The classic information assurance categories (availability, confidentiality, integrity) are not useful categories for understanding the business effect of cyber attacks, because they only identify the mechanism, not the consequences. With a bit of imagination, it is possible to see how a breakdown in any of the information assurance categories could be used to produce any of the business consequences. Breakdowns in availability, for example, could be used to discredit a business operation. Breakdowns in confidentiality could be used to interrupt a business operation. There is no close relationship between the information assurance categories and their business effects.

### **Putting the Proper Risk Estimates Back into the ROI Calculation**

Once the quantification of consequences described here is inserted into the existing analyses of risks, it becomes possible to start calculating the ROI's for security correctly. In addition to providing the central term in the risk calculation, a clear and quantitative understanding of consequences allows the estimates of the threats and the vulnerabilities to be made much more precise. This can provide a sound economic basis for decisions and policies regarding security and privacy.

[A more complete explanation of this method and these concepts will be found in Scott Borg, *Cyber Attacks: A Handbook for Understanding the Economic and Strategic Risks* (forthcoming).]

Scott Borg is the Director and Chief Economist of the U.S. Cyber Consequences Unit (US-CCU), an independent, non-profit research institute that investigates the strategic and economic consequences of possible cyber attacks. He is responsible for many of the concepts that are currently being used to understand the implications of cyber attacks in business contexts. In collaboration with John Bumgarner, he is author of the US-CCU Cyber-Security Check List, the most comprehensive survey to date of cyber vulnerabilities. He regularly advises a number of different U.S. government departments and industry associations. Before being asked by government officials to tackle cyber-security issues, Scott Borg was one of the principal developers of Value Creation Analysis, a set of business strategy models for understanding how much value can be created by various types and components of value chains. He has been a guest lecturer at Harvard, Yale, Columbia, and other leading universities. He is currently a Senior Research Fellow in International Security Studies at the Fletcher School of Law and Diplomacy of Tufts University.