CHAPTER 2

The Misperception of Risk

> *The moment we want to believe something, we suddenly see all the arguments for it, and become blind to the arguments against it.*
>
> —George Bernard Shaw

One hundred years ago, the "unsinkable" *Titanic* foundered after striking an iceberg off the coast of Newfoundland. More than 1,500 people died in what became one of the deadliest maritime accidents ever. Several factors contributed to this massive death toll, but perhaps the most critical was that there simply weren't enough lifeboats. The ship carried 2,224 people, but fewer than half of them could squeeze into the boats.

As we know, passengers who didn't get a spot in one of those lifeboats quickly died in the freezing waters of the North Atlantic. What's less well known is that the *Titanic*'s supply of lifeboats was in full compliance with the British marine regulations in force at time. The law required the ship to carry 16 lifeboats; the *Titanic* actually had 20 lifeboats.

The ship's owners did a good job of providing enough boats to address the regulatory risk of noncompliance. Unfortunately, meeting regulatory requirements did little to prevent the tragic loss of life.

This is a case of *misperception of risk*. The owners focused on mitigating the regulatory risk, apparently blind to the much larger risk of disaster. A sad footnote: reports suggest the *Titanic* had enough capacity to easily add enough lifeboats for everyone on board, had the owners chosen to do so.

What does this example have to do with information security? We encounter misperceptions every day within the realm of enterprise risk and security. Furthermore, unless we mitigate these misperceptions, they can have disastrous consequences. As a result, I believe the misperception of risk is the most significant vulnerability facing enterprises today.

The Subjectivity of Risk Perception

As security professionals, we tend to think about objective ways to estimate risk—to assess the likelihood and extent of harm that can occur due to specific threats and vulnerabilities.

But in reality, the way people perceive risk has a strong subjective component. Economic and psychological factors greatly affect how each of us perceives the likelihood

and potential impact of harm from specific actions or situations. Within an organization, each individual's perception of risk varies depending on his or her job role, goals, background, and peer group. This means managers, security professionals, and end users all may have a different view of the risk associated with a specific technology or action.

Misperceiving risk has serious consequences because our actions are shaped by our perception of risk. An employee may think posting personal and work-related information on a social-media site is relatively harmless. However, hackers might use this publicly available information in phishing e-mails to gain access to enterprise systems via the employee's computer, ultimately resulting in detrimental security breaches.

End users are not the only members of the organization who can misperceive risk. Everyone is capable of misperceiving risk, including risk and security professionals. As I'll explain later in this chapter, misperceptions occur at the group level as well as the individual level. Members of a group may share the same bias in their perception of risk and benefit.

The decisions that result from these misperceptions can weaken the entire organization's security posture. If an organization underestimates a risk, it will under spend on controls to mitigate that risk, increasing the likelihood and potential impact of major problems such as data breaches. On the other hand, if the organization overestimates a risk, it will allocate a disproportionately large share of its security resources to the risk, leaving other parts of the risk landscape underprotected.

In this chapter, I'll discuss how and why different people within an organization misperceive risk—whether they are acting as information technology users, security professionals, or managerial decision makers. To explore these misperceptions, I've drawn on research across the broader field of risk psychology, notably *The Psychology of Risk*, a book by Professor Dame Glynis Breakwell, Vice Chancellor of the University of Bath (Cambridge University Press, 2007). I'll examine how these ideas about risk perception apply to information risk and security. I'll explain some of the consequences of those misperceptions, and I'll discuss some of the ways an organization can address them.

How Employees Misperceive Risk

Research shows that if we like an activity, we tend to judge its benefits to be high and its risk to be low (Slovic 2010). Conversely, if we dislike the activity, we judge it as low-benefit and high-risk. Because of this, the perception of risk by individuals and groups within an organization tends to be biased by their preferences, roles, and objectives. Everyone is trying to achieve their individual or group goals within the organization, so they tend to see activities and technologies that support those goals as beneficial, and therefore they tend to underestimate the risk.

So if employees like social media, their attraction to the technology skews their perception of benefit and risk. Because they judge the benefit to be high and the risk to be low, they feel comfortable posting information such as their job title, location, and even the projects they're working on. They may even allow sites to capture their location, using the global positioning system in their cell phone, and display the location in real time.

Unfortunately, these employees may not think about how a malicious individual could use the information. Today, as we've seen, an individual's use of technology can harm not only the individual, but the entire organization. Attackers exploit publicly

available personal information to craft spearphishing e-mails that are particularly convincing because they appear to demonstrate a relationship with the recipient, making the employee more likely to click on a link that downloads malware to the system. From there, the attack spreads to the rest of the corporate network. In addition, information posted by individuals is now routinely aggregated, analyzed to identify patterns, and sold, often to a company's competitors.

The risk and security team may also misperceive the risk of social media, but in the opposite direction—they overestimate the risk and underestimate the benefits. They may not like social media because it creates vulnerabilities, and their perception then drives them to focus on minimizing the risk by trying to block the use of the technology.

Other psychological factors also come into play in shaping end users' risk perception. People in general tend to believe they are personally less likely than others to experience negative events, and more likely to experience positive events, leading to a sense of personal invulnerability (Breakwell 2007). In addition, users also are more likely to behave in risky ways if their colleagues do so. "It's conformity—being seen to be doing what everybody else is doing," Breakwell says (pers. comm.). Many social media sites encourage this conformist tendency—if all your friends are using a social media site, you're likely to join the site too because it enables you to see what they are doing and share information with them more easily.

The likelihood that individuals will behave in ways risky to the organization also increases when their individual interests don't align with the company's. This divergence is most likely when employees are discontented, resentful, demoralized, or simply don't trust IT or the broader organization.

In economic theory, the problem resulting from this lack of alignment is known as a *moral hazard*: a situation in which someone behaves differently from the way they would if they were fully exposed to the risk. A useful moral hazard analogy is renting a car with full insurance coverage. People are likely to be less careful with the rental car than they would be with their own car if they're not responsible for the consequences. The attitude is "if it's not mine, it doesn't matter."

In the realm of enterprise IT, moral hazards may be a bigger concern than many appreciate. A Cisco survey (2011a) found that 61 percent of employees felt they were not responsible for protecting information and devices, believing instead that their IT groups or IT service providers were accountable. Ominously, 70 percent of these surveyed employees said they frequently ignored IT policies.

One indicator of the extent of moral hazard within an organization may be how employees treat company-provided laptops. Higher-than-average loss or damage rates might suggest employees don't care about the laptops and may be an indication they don't care about other corporate assets either. As I'll discuss in Chapter 5, I believe allowing reasonable personal use of laptops can help reduce the risk of moral hazard because it aligns personal interests with those of the organization.

More broadly, organizations can address the moral hazard issue by taking steps to align the goals and concerns of everyone involved: end users, information security professionals, and executives. This returns to the theme of the book—as information security professionals, our mission is to Protect to Enable. This mission aligns our security goals with those of the business. It helps maintain the perception of shared values. Research suggests that people with whom we share values are deemed more trustworthy (Breakwell 2007, 143). If employees trust us, they are more likely to believe our warnings and act on our recommendations.

One further point to remember is that everyone in the organization, regardless of the job role, is an end user. Therefore, we can all fall prey to the same tendencies. For example, we may be attracted to new consumer technologies and tend to ignore the risks.

How Security Professionals Misperceive Risk

While end users tend to underestimate the risks of a desirable activity or technology, security professionals sometimes display the opposite tendency. We focus obsessively on the information risk associated with a specific threat or vulnerability. In doing so, we completely miss bigger risks.

This phenomenon is known as *target fixation*—a term originally coined to describe a situation in which fighter-bomber pilots focus so intently on a target during a strafing or bombing run that they fail to notice the bigger risk to themselves and crash into the target as a result (Colgan 2010, 44). As information security professionals, we can develop a similar fixation. We focus so intently on one risk that our awareness of larger hazards is diminished. This target fixation can also occur in other groups with "control" functions within the organization, such as internal audit, legal compliance, and corporate risk management.

Here is an example from our own experience at Intel, which I'll discuss further in Chapter 9. Several years ago, we discovered that malware had been introduced onto our network from an employee's personal computer. We became so focused on this source of danger that we eliminated all personal devices from our network. We further fueled our target fixation by labeling these devices non-Intel managed systems (NIMS), a term that reflected the frustration over our lack of control. I vowed we would never again allow network access from devices that we didn't fully control.

However, by becoming fixated on a single threat, we may have created some larger risks and additional costs. For example, we needed to issue contract employees with corporate PCs, each of which allowed broader access to the Intel environment. If we had instead focused on how we could provide limited access to the environment from "untrusted" devices, we might have managed the risk with lower total cost and obtained a head start in developing a key aspect of our current security strategy, as I'll describe in Chapter 8.

As security professionals, we also may misperceive risk due to the tendency to "set and forget" security controls. This common security loophole is described in the sixth Irrefutable Law of Information Security in Chapter 1, which states that the efficacy of a control deteriorates with time. Once in place, controls tend to remain static, while the threats they are intended to mitigate continue to evolve and change, sometimes in very dynamic ways. Controls that are initially very effective can become inadequate over time. Ultimately, an adverse event may occur and may even have disastrous consequences.

Think about the history of major oil tanker spills. For years, regulations allowed tankers to be built with a single hull, instead of a double (inner and outer) hull to provide additional protection in the event of a leak. Meanwhile, tankers grew steadily larger because bigger ships could transport oil more efficiently than smaller ones. It wasn't until the *Exxon Valdez* ran aground, puncturing its hull and creating a giant oil leak that contaminated huge stretches of Alaska's coast, that authorities were spurred to create new regulations requiring double hulls in oil tankers (EPA 2011).

Within enterprise IT, a typical "set and forget" error is the failure to keep controls up-to-date, particularly if the controls are designed to mitigate a relatively low risk. A case in point: *distributed denial of service (DDoS)* threats were a big concern more than a decade ago, due to widely publicized attacks by worms such as Code Red, Nimda, and SQL Slammer. These attacks disabled corporate web sites or flooded internal networks by overloading them with requests. To mitigate the availability risk, many organizations invested in defenses against DDoS attacks.

Over time, however, DDoS attacks became less frequent, and organizations were assailed by newer threats. With limited resources, information security groups focused on mitigating these new threats rather than continuing to build defenses against DDoS attacks. At the same time, though, businesses were increasing their online presence. Web sites evolved from being used primarily for advertising and displaying static corporate information to managing business-critical data and applications. Some organizations began conducting all their business online. Even traditional brick-and-mortar businesses moved customer support, order management, and other critical business processes onto the Web. The larger online presence multiplied the potential impact of a successful attack. As a result, when DDoS attacks from a variety of groups resurfaced in the past few years, they created even greater disruption to business operations as well as damage to corporate brands.

Another example: over the past few years, many organizations have become much more diligent about scrubbing data from the hard drives of old computers before disposing of them or reselling them. But they failed to follow similar precautions for other business devices that have evolved to include hard drives.

Nearly every digital copier contains a drive storing an image of each document copied, scanned, or e-mailed by the machine. When CBS News reporters visited a company that specialized in reselling used copiers, they found businesses and agencies had discarded machines containing lists of wanted sex offenders, drug raid targets, pay stubs with Social Security numbers, and check images. One copier's hard drive even contained 300 pages of individual medical records, including a cancer diagnosis, which is a potential breach of federal privacy law (Keteyian 2010).

MISMATCHING CONTROLS TO THREATS

Businesses sometimes devote considerable time and resources to implement security controls that are completely irrelevant to the threats the companies are trying to mitigate. These mismatches reveal a lack of understanding of the security technology and the threat. The controls may further add to the risk by providing a false sense of security. In reality, deploying the wrong control is like carrying a lightning rod to protect oneself from getting wet in a storm.

Typical mismatches include:

- Using firewalls to prevent data theft from applications that are allowed to operate through the firewall
- Using standard antivirus tools that are effective only against previously identified threats, to protect against zero-day attacks

- Using controls at the operating-system level to detect application-layer attacks

This mismatch does not mean that these controls are worthless. It simply means that if our goal is to deal with a specific threat, we must understand both the attacks and the controls well enough to identify which controls are applicable, and where it is necessary to add other controls. For example, if a firewall cannot prevent attacks against an application, we might deploy an additional control behind the firewall.

How Decision Makers Misperceive Risk

A manager makes decisions based on information from technical specialists and other experts. Therefore, the decisions managers make are only as good as the information they receive. Decision makers can misperceive risk when their decisions are based on biased or incomplete information.

Bias can influence these decisions every day. If people are trying to sell a particular proposal or point of view to their manager, what are they likely to do? They tend to select data supporting their arguments and often ignore data contradicting those arguments.

The danger of misperception is particularly acute when decision makers rely on a narrow range of sources who all share similar viewpoints. Without obtaining a diversity of viewpoints, managers don't get a full picture of the risk. Like-minded individuals tend to agree with each other, as you might expect. When a group is composed solely of people with similar backgrounds and viewpoints, it may be particularly prone to group polarization (Breakwell 2007, 99) and the group's decision may be more extreme than the mean of their individual views. This problem may be especially acute when the people involved share the same mental model of the world, as is likely to be the case when the group consists only of specialists from the same organization.

An even broader concern is how a focus on business goals can drive people to make unethical decisions. When these decisions are made by managers at the organizational level rather than at the individual level, the impact is compounded by the potential for widespread disaster.

After the *Challenger* space shuttle exploded in 1986, extensive post-crash analysis revealed the tragedy was caused because an O-ring on one of the shuttle's booster rockets failed to seal due to the low ambient temperature at launch time.

However, it subsequently emerged that engineers had warned of the potential danger before the launch. Engineers from NASA contractor Morton Thiokol recommended the shuttle not be launched at low temperatures after analyzing data that indicated a link between low temperatures and O-ring problems. After NASA responded negatively to the engineers' recommendation, Morton Thiokol's general manager reportedly decided to treat the question of whether to launch was a "management decision." Against the objections of their own engineers, Morton Thiokol's managers then recommended NASA go ahead and launch, and NASA quickly accepted this recommendation (Bazerman and Tenbrunsel 2011, 13–16).

For Morton Thiokol's managers, the desire to meet the business goal of pleasing the company's customer, NASA, apparently caused the ethical dimensions of the problem to fade from consideration—with terrible consequences.

According to Tenbrunsel, this "ethical fading" is not uncommon. The way a decision is framed can limit our perspective. If the decision is framed purely in terms of meeting business goals, ethical considerations may fade from view. In fact, we may become blind to the fact that we are confronting an ethical problem at all (Joffe-Walt and Spiegel 2012).

Another infamous ethical lapse involved the Ford Pinto, whose gas tank exploded in a number of rear-end collisions, resulting in fatalities. As Bazerman and Tenbrunsel describe (2011, 69–71), Ford discovered the dangers in preproduction testing. However, facing intense business competition, the company decided to go ahead with manufacturing anyway. The decision was based on a cost-benefit analysis. Ford apparently considered the choice as a business decision rather than an ethical decision and determined it would be cheaper to pay off lawsuits than make the repair. The impact of dehumanizing this risk decision was disastrous.

In the past, many information technology risk decisions have often been considered only in terms of their potential business impact. As information technology is integrated into more and more products, decisions about information risk will increasingly affect the lives of millions of people, making it essential to consider the ethical as well as the business dimensions of information risks. It becomes even more important that we, as CISOs, keep ethical considerations to the forefront. What is the potential impact of a security breach when a car's sensors and control systems can be accessed via the Internet? Or when medical life-support equipment can be remotely controlled using wireless links?

How to Mitigate the Misperception of Risk

It should be apparent by now that the tendency to misperceive risk is universal. We need to find ways to help compensate for this misperception, given that it is our job to manage risk. As security professionals and managers, how can we mitigate the misperception of risk?

We can start by ensuring we include a diversity of viewpoints when making risk management decisions. Whenever possible, we should involve a broad cross-section of individuals representing groups across the organization. This diversity helps compensate for individual biases.

However, assembling the right mix of people is only the first step in building a more complete picture of risk. As information security professionals, we need to ensure that the discussion brings up new perspectives and views. We must ask penetrating questions designed to bring alternative viewpoints to the surface. We need to continually seek out the minority report, the view that is contrary to perceived wisdom. If the majority is telling me to turn right, are we missing something important that we'd find out by turning left?

This questioning counteracts the inevitable bias due to target fixation. We can also help counter target fixation by simply recognizing it exists, and then consciously trying to see the problem from someone else's viewpoint.

Uncovering New Perspectives During Risk Assessments

Risk assessment models can be valuable tools for helping to evaluate risks and to prioritize security resources. But all models have limitations. If we base our decisions solely on the results generated by a model, we may miss important risks.

At Intel, we typically use a risk assessment model based on a standard methodology. The model scores each risk using the formula:

```
Impact of Asset Loss × Probability of Threat × Vulnerability Exposure =
Total Risk Points
```

For each risk, we assign a rating to each of the three contributing factors in the formula. To illustrate, I'll use a scale of 1 to 5. A high-value asset, such as a microprocessor design, might warrant a rating of 5.

We then multiply the three ratings to obtain the total risk points. In this example, the maximum possible risk score is therefore 53, or 125.

A simple approach to risk management, using the output of the model, would be to divide the security budget among the highest-scoring risks.

The model is valuable because it provides a consistent method for helping compare and prioritize a broad spectrum of risks. However, allocating resources based only on the overall risk score can miss potentially disastrous "black swan" events that have very low probability but extremely high impact (Taleb 2007). Because the formula simply multiplies three ratings to obtain the overall score, black swans tend not to score as highly as lower-impact events with higher probability.

To counteract this problem, we can examine the information in the model in more detail, from different perspectives. We can create a list of the 20 most valuable assets and consider whether they need additional controls. In the same way, we can examine the top threats and vulnerability areas.

The point is that any model used to calculate risk should be used as a framework to drive a dialogue about all the variables and options, rather than as a tool that generates the answers to our problems. By discussing the issues from a variety of perspectives, we may identify important concerns we'd miss if we simply look at the overall risk scores.

Before I moved into the information security field, I worked in finance. In our finance group, we found the same principle held true when conducting ROI (return on investment) analysis. Our ROI model generated forecasts. However, it was by discussing the model's assumptions that we determined whether or not the model's predicted financial returns were reasonable.

Another method for prioritizing information systems risk management is to examine systems from the perspective of critical business processes and to consider the impact of a loss of confidentiality, integrity, or availability.

An application that prints shipping labels may initially appear to be low priority because it is small, inexpensive, and doesn't contain confidential data—it simply takes the information it needs from a customer information system on the network. However, if it's unavailable because the network is experiencing problems, the impact is huge because the company cannot ship products.

The potential impact to a business process of losing confidentiality, integrity, or availability may also vary depending on the stage of the business cycle. Consider a payroll system. Information confidentiality and integrity are always important; but availability is exceptionally critical on payday.

Communication Is Essential

Communication is an essential part of any strategy to mitigate the misperception of risk. To alter the way people behave, we need to change their perception of risk. To effect that change, we must communicate with them.

Changing perceptions is difficult. We may need to address long-held preconceptions about what is risky and what is not. Once people form an initial estimate of risk, they can be remarkably resistant to adjusting their perception, even when given new information (Breakwell 2007, 59).

In addition, each person may have a different perception of risk. To communicate effectively, we may need to understand an individual's viewpoint and then tailor our communication accordingly. Consider the example of taking laptops to countries with a high risk of information theft (see sidebar). People who are extremely concerned may need a patient, thorough explanation of the risks and benefits of taking their laptop versus leaving it in the office. A less fearful individual may just need a quick reassurance and a few basic facts.

Though changing risk perceptions can be challenging, we don't have any choice but to try. Employees will use social media whether we like it or not. When they do, they may not only put themselves at risk; they could be putting the company at risk too, if they are not careful.

Communication can reduce the issue of misperception due to asymmetry of information. This asymmetry is created when security professionals know about risks but don't share the information with end users within their organization. When two parties differ in their knowledge of a threat or vulnerability, their perception of risk is likely to differ also. In other words, it is difficult for users to care about a hazard if they don't even know it exists.

To succeed in changing users' perceptions, we must communicate in ways that engage them, using language they understand rather than technical jargon. At Intel, we have employed entertaining, interactive video tools to help engage users and teach them how to spot dangers such as phishing web sites. As I'll explain further in Chapter 5, we've found these methods have been highly effective in changing users' awareness and perceptions, and ultimately in shaping their behavior.

Patiently explaining to users the consequences of their actions can also help shape their perception of risk. In some countries, pirating software is so commonplace that it is almost an accepted part of the culture. This poses a problem for many multinational companies. Employees in these countries may not even believe copying software is wrong, let alone view it as an illegal act. It can be useful to describe the potential consequences of copyright infringement for the individual and for the organization. We can explain to employees that a decision to pirate software can expose the company to software license compliance risks. The consequences may be even more far-reaching if the copied software is then incorporated into the company's technology-based products or services. If a product is discovered to include stolen software, the company may be

unable to ship it to customers, which means a significant loss of revenue. Of course, employees may experience personal consequences too: if they copy software, they run a high risk of losing their jobs.

Organizations as a whole may also be blind to risks, or simply choose to ignore them. One way to overcome this misperception is to patiently build up a list of examples showing how other organizations ignored similar risks and experienced adverse consequences as a result, according to Breakwell, the University of Bath psychologist (pers. comm. 2012). The more examples in the list, the harder they are to ignore.

"Organizations stick their heads in the sand, ostrich-like," she says. "But if you have a database of examples illustrating where things have gone wrong elsewhere, it becomes harder and harder to find enough sand to stick your head in."

CHALLENGING PRECONCEPTIONS: TAKING LAPTOPS TO HIGH-RISK COUNTRIES

It may be necessary to challenge perceived wisdom in order to expose a clear picture of the real risks, and consequently make the right decision.

Some companies react to the higher rates of intellectual property theft in certain countries by barring employees from taking their corporate laptops on business trips to those countries. In some cases, the companies issue employees with a new "clean" system from which all corporate data has been purged.

The goal is to prevent situations in which information theft might occur, such as when an employee leaves a laptop containing corporate data unattended in a hotel room. A malicious individual could then get physical access to the system and copy the data or implant software that will surreptitiously steal information over time.

But does preventing employees from taking their familiar laptops really solve the problem? Let's suppose we issue employees with a new, data-free laptop. To do their jobs, they'll still need to use this system to log into their corporate e-mail and other applications—providing an opportunity for hackers to intercept the network traffic.

Furthermore, if attackers really want to target an individual, they have ways to do it without gaining physical access to the system. With a spearphishing attack, they can induce the individual to click on a malicious link that remotely downloads malware.

Preventing employees from taking their laptops and information also deprives the organization of the key business benefits of using a full-featured portable computing device; employees will likely be less productive as a result. So when assessing the risks of traveling with mobile devices, an organization needs to think through the tradeoff between risk and benefit, including the cost of providing what they believe to be a "clean" system and the impact on the user.

Building Credibility

Ultimately, our ability to influence people's risk perception depends on our credibility. We need to build trusted relationships with executives and specialists across the organization to ensure our security concerns are seriously considered rather than seen as fearmongering or target fixation.

Trust is hard to create and easy to destroy. If business groups think we are providing unreliable and exaggerated information, will they trust us to provide their security? If we create a security scare about a threat that turns out to be irrelevant or overblown, we may be seen as just another source of misperception.

As I'll describe in more detail in Chapter 9, we can establish credibility by demonstrating consistency, striving for objectivity, and showing that we can accurately predict the real security issues affecting the organization, and then communicate them in an effective and timely way. Credibility is also built on the competence that comes from understanding the business and technology as well as possessing core security skills. As the scope and importance of information security continue to expand, creating this credibility provides an opportunity to step into a more valuable, high-profile role within the organization.