Chapter 7

Life in the New School

There are not always easy answers. Solutions that might apply in the general case are not necessarily applicable in every case. We have spoken in the previous chapters about how applying ideas from the New School will create positive effects. It is possible to achieve a future in which information security is not riven by opinion and demagoguery. Using objective data and embracing other sciences can trump both fear and fashion. We have also spoken throughout this book about how there is no *single* solution to the majority of the challenges that exist within information security. We have railed against the notion that any one product is a "silver bullet," or that there is some special framework or methodology that an organization can simply employ. The New School is neither a product nor a service. It is an approach to the world that embraces the scientific method, new sources of objective data, and new perspectives from diverse fields from which new theories and approaches flow.

The turn toward the New School will influence how we speak about problems. It will alter our immediate and considered reaction to news, shifting the emphasis from overdramatization to a more calm assessment of the facts. The New School will influence how individuals and organizations purchase security technologies by providing better answers to questions such as "how big is this problem?" The model of employing scare tactics to drive sales will decline, because anyone will be able to search for information about problems and determine how prevalent they really are. It will also be easier to find out what approaches to problems

actually work. These changes will influence which information security products and service offerings are brought to market and which become successful.

Even though life in the New School will be a much more productive and enjoyable experience than in the old, it will not be a utopia. In acknowledging this, we hope to differentiate ourselves from those ideologues and snake oil salesmen in the information security industry who refuse to see the weaknesses in their own argument, product, or service. This chapter describes some areas and challenges in the field of information security for which the New School either does not provide an immediate solution or does not suggest a solution at all. We'll talk about human nature, the limitations of breach data, externalities, risk compensation, language, and organizational issues such as skills shortages.

Lastly, we hope and expect that the ideas within this book will be challenged. We look forward to that debate. By engaging in a conversation about how to make better security decisions, the New School can only be strengthened.

People Are People

The New School won't be a utopia, because people are people, and so why should it be? People still smoke despite the long-term risks because of the immediate pleasure it provides them. People don't exercise, even when they know they are overweight and at risk of heart disease and diabetes. People often make such choices knowing the risks. In many places, cigarette packs are decorated with horrifying pictures of people with diseases caused by smoking, or they carry harsh warnings such as "Cigarettes cause strokes" or "Tobacco use can make you impotent." People still choose to smoke, but doctors no longer endorse cigarettes. The availability of objective data allows better choices to be made by those who seek out that data and act on it. We expect that most will do so, and

that better choices will be made in aggregate. It seems inevitable that some individuals and organizations will continue to make choices that are bad, confusing, or morally distasteful to some.

Poor decision-making, stubbornness, and apathy can be expected. They are part of the human condition. Some of the reasons people might not embrace the use of objective data to make better security decisions are more subtle. A number of psychological effects come into play in security decisions. In Chapter 6, we discussed how an unnecessarily risk-averse mind-set can be created by the fear of loss, even when the probability of loss is low. This aversion to risk can lead to inefficient and misdirected spending as a response, even when data suggests a better, more optimal course of action. In America, car accidents kill a far greater number of people than acts of terrorism, yet spending on antiterrorism measures dwarfs spending on road safety. The spectacular, horrific nature of terrorism creates a strong emotional response. Studies have shown that people find it difficult to think rationally about risks that carry heavy costs. As we noted in Chapter 2 when we discussed security surveys, the way in which questions are framed and situations are posed also has a strong effect on people's expressed opinions and their subsequent actions. Even as we move from fear to data, data can still be manipulated to create an emotional response, and people still tend to put their own interests first. Participants in the New School seek to overcome these psychological effects by using objective data and its analysis as the key factor in effective decision-making. They also recognize that people do not always behave rationally.

Chapters 5 and 6 highlighted the use of economic models. Those economic models are typically predicated upon certain assumptions. A common assumption is the notion of a "rational actor," which is a model of human behavior in which individuals are expected to behave in ways that maximize the utility of

their efforts. These types of simplifications allow economic models to be built at the cost of glossing over underlying complexities. Security models often break down when the user violates assumptions. A classic example that is known to all security practitioners is the class of attack known as the "buffer overflow." In this attack, the programmer writing a program assumes that the user will not enter input that exceeds a certain length. The attacker deliberately enters more than the expected amount of input, possibly changing how the program runs. We can perhaps expect that economic models are no different from security models, in that the assumptions they make can be their undoing. Research has shown that human beings deviate from conventionally defined economic rationality. In other words, people sometimes *choose* to behave in a risky manner.

When we discussed economics, we spoke of "agency problems." Although the economics of information security might enable us to propose strategies for making better security decisions, agency problems won't go away. Individuals and organizations will continue to be swayed by subjective or emotional factors. Some security executives promote themselves by pointing to their embrace of the latest security technologies. They might continue to pursue certain security projects that are seen as attractive and use them for self-promotion. Many security practitioners also want to obtain practical experience using the latest technologies because of their desire to increase their marketability. It would be naive to believe that the availability of economic models that can guide security decisions will have the universal effect of purging personality and organizational politics from decision-making. They will help, though.

Similarly, if a market continues to exist for security products that make the buyer feel better, but only superficially solve problems, companies will continue to provide products

and services to satisfy that market. A case in point is antibacterial soap, which seems to make people no healthier than standard soap and leads to more-resistant bacteria. Chapter 2 noted that many of the commercial information security products that exist today seek to correct only the symptoms of problems rather than the problems themselves. Companies build and sell these products because the effort involved in "solving" the more superficial aspects is less than would be required to address the problems at their root. Those who are responsible for making purchasing decisions within an organization might not realize the fundamental nature of their challenges, so they might unwittingly support that market. As more objective data becomes available, this behavior will be reduced, but some people might have the misfortune of remaining unaware of the benefits of using the data.

Some people choose to hoard data or form cliques to restrict its distribution, and this might be a difficult habit to break. Within the security field there is a great deal of secrecy around topics such as new vulnerabilities and attack techniques. Groups and cliques form to create or maintain power structures, often based on the secrecy of the knowledge they hold. As Lord Acton said, "Power tends to corrupt; absolute power corrupts absolutely."

Some cliques are commercial endeavors. A number of "information-sharing" organizations charge a membership fee. These fees are often structured such that only the wealthiest organizations can afford to be members. Perhaps this makes sense, to ensure that the organization will send appropriately senior representatives. Or the reason for the high fees might be elitism. In either case, the exclusivity of such organizations stands in opposition to the notion of free and open data sharing that is championed by adherents of the New School. A countervailing force to such cliques are truly open information-sharing initiatives. All else being equal, groups focusing

on open information sharing and analysis will create more value than those who invest in secrecy and exclusion. The open group will also reap the benefits of more diverse membership.

Breach Data Is Not Actuarial Data

Chapter 6 looked at traditional methods for attempting to calculate the value of security-spending decisions. We noted that these methods fail without actuarial data. Some of them require knowing the *likelihood* of a loss event such as a security incident. Without actuarial data, reliably determining that likelihood is very difficult. The breach data that is in the public domain today goes a significant way toward meeting that need. But to achieve an actuarial-like quality, it would need to be more detailed than it is today, and we would need more data to derive trends. This is not to say that breach data can never reach that required threshold, only that it does not do so today.

A complicating factor in creating actuarial data for information security that needs to be overcome is that the technology landscape is changing much more rapidly than, say, 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: floodplains are well mapped out, as are tornado-prone areas. Insurance companies only need to track a relatively small number of variables to generate actuarial data. The risks for other insurance markets, such as asbestos-related problems, can also be calculated. (Callous as it may seem, the upper boundary for asbestos insurance is the medical costs for everyone who has worked or will work in the asbestos industry.) In contrast to these topics, 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 systems decays over time.

If an insurance company offers homeowners' insurance only in Florida, it has a different exposure to hurricane-related risks than a company that offers homeowners' insurance only in California. There is a thriving business in which insurance companies sell part of the liability for their portfolios to re-insurers. Re-insurers then work hard to ensure that their portfolios are balanced across the existing pools of risk. The number of factors they must balance is rather small relative to those that might come into play in the electronic world. For example, criminals can execute electronic attacks at scale, with no consideration for issues such as geography. These factors make it difficult to create a model that would allow information security risks to be priced appropriately. (Whoever figures it out will make a mint.)

Breach data is a profoundly useful source of information, and the availability of objective data about successes and failures in information security will greatly improve our ability to make better security decisions. However, challenges still exist to obtaining actuarial data that would support insurance markets.

Powerful Externalities

Chapter 5 discussed externalities—costs that are not felt by the organization or individual who performs the action that creates them. One of the examples we used was of drivers of larger vehicles such as SUVs not experiencing more smog than drivers of smaller cars. Although SUV drivers create more smog, they do not feel the effects of poor air quality in proportion to the amount of air pollution they create. There are analogues to this type of situation in the security world. Failing to keep up to date with security patches might cause a computer to become compromised. The compromised computer might be

used to attack other computers. The impact of those attacks is an externality to the owner of the computer used to launch them. Assuming that the computer owner suffers no consequences for the attacker's actions, all the cost is borne by the party being attacked.

There are negative externalities, but also positive ones. Using the same example, an organization that invests in security measures creates *positive* externalities for the companies with which it interacts electronically. This is true because the organization's computers are less likely to become compromised or infected and therefore are less likely to spread a compromise or infection to the business partners. All sorts of social parallels exist, such as using immunization to reduce the spread of disease.

Within the study of the economics of information security, some strategies have begun to be identified to overcome these challenges. We've discussed some of them, such as product bundling and the use of incentives and penalties. Even so, the market today exhibits powerful externalities that work against security. An example is the level of security quality that software vendors, particularly start-ups, choose to implement in their products. The marketplace tends to reward companies that are the first to market in any new product space. This causes many companies to not worry about security properties for the first versions of their products. They might include some security features, such as encryption, but not worry about vulnerabilities in their code or design. The orientation of their programmers doesn't include worrying about security. The costs are borne by the customers of those products, due to their insecurity and the effort involved in improving their security later. Applying economic thinking to identify these challenges and propose solutions is extremely useful. However, the ingrained nature of many of the current problems suggests that the point at which positive effects will begin to be widely seen might be some time away.

The Human Computer Interface and Risk Compensation

Great user interface design is hard. There are more examples of bad interfaces than good ones. It's also much easier to point out that an interface is bad than to point out that it's good. (This has much in common with security.) For a long time, usability wasn't part of what sold software, never mind security software. Security was even supposed to be *hard* to use—to stop the user and force him or her to *think!* Interfaces are often designed to support particular tasks. If the designer's understanding of those tasks is wrong, the interface won't help the users get to their desired results. Thinking about usability requires thinking about the user, and that's a different orientation from thinking about the code. When a programmer thinks about the user, he often makes what he thinks are small, useful changes to the system, such as improving the Open and Save dialog boxes to "be more intuitive." In the programmer's view, these little visual tweaks are helpful. The customers may see lots of little tweaks from lots of places, and this can all add up to a cacophony. Icons aren't seen as the designer intended. The word the user searches for isn't the one the programmer thought the user would use. And all that is before security is considered.

People often think of security and usability as being at odds. Security certainly can prevent things from "just working." Unfortunately, if you can get to your files from anywhere without a password, so can anyone else. Security decisions often involve situations in which the computer can't make the decision by itself, so the user must get involved. Is the web site really the bank's, or is it a clever imitation? It might be easier if there was only one bank on the internet, but we bet they'd have lousy customer service.

Usability might seem to be a subjective set of questions, with fuzzy definitions and no way to resolve them. But usability professionals have rallied around testing their ideas as a

way out of the mess. Great usability experts believe in testing, refining, and testing again as a way to improve designs. (Security can learn a lot from usability.) Some of the challenges involved are that people are complex. They're very observant, and sometimes they pick up on subtle clues, which is great when it works but is hard to rely on. Returning to our accused Turkish hacker, only 0.3% of the potential victims fell prey to the false email. To put it another way, 99.7% of the people didn't fall for it. What happened to the others? Maybe they were on the phone as they were looking at their email. Maybe something distracted them, and they clicked the wrong link. Maybe we can use people's risk thermometers to keep them cautious, or maybe such attempts would blow up in our face. It's hard to test such things. Telling experimental subjects that a study is about security changes their behavior. Not telling them may violate ethics rules. Studying behavior that only three people in a thousand exhibit may require a very large sample. Making software usable is an empirical and iterative process, based on designing and running good tests.

Innovative work on these challenges is being shared at the Symposium on Usable Privacy and Security. Some of these papers have been assembled into an excellent introductory book, *Security and Usability*. This field is still in its infancy, and a great deal remains to be done. Even when we have done a generally good job of making security understandable and usable, attackers will still be motivated to find places where they can induce confusion and then guide users into making the wrong choices. This seems likely to be the case for as long as software is as complex as, say, a VCR. We know things can get better, but perfection isn't coming anytime soon.

Even when things have improved substantially at a technical level, the effects of risk compensation will keep things interesting. We've discussed how psychological factors can influence security decisions. The phenomenon of risk compensation (also known as risk homeostasis) can have a negative

effect on security measures. Because risk homeostasis can weaken security measures or render them ineffective, it would be useful to know how to construct systems so as to ward off, or at least mitigate, its effects. One approach we have described is designing security measures so that their operation is invisible to those who receive the protection. Good security design involves putting the "right" security on the default path. However, a great deal of today's software depends on end users making the correct security decisions. During the installation of a popular web-based collaboration product, the installation program presents a message that states "Always click 'Yes' or 'Always' when receiving messages from your web browser when using this program." This desensitizes the users to the decisions they are making.

As described in Chapter 6, we are skeptical of the traditional responses to this problem, such as the "security awareness training" of end users. First, there does not seem to be any good evidence that training users causes them to behave differently weeks or months after the training. Second, risk compensation allows us to suggest that users will believe the risk in any given situation will somehow be mitigated by other protective controls. Third, even if security awareness training were shown to be effective, the "weakest link" principle applies. In other words, an attacker needs to convince only *one* employee to make the wrong choice. People are people, and if end-user security awareness training would make a difference, everyone would exercise regularly, practice safe sex, and never smoke cigarettes.

We are stuck, then, between ideals of how to construct security systems to be secure and usable, and the reality that there are pervasive, deeply entrenched systems that are not implemented with those design goals in mind. Those platforms can improve, and they *will* improve as these concepts within economics and psychology are more widely understood, but

much work remains to be done. If computer security improves dramatically, people are likely to believe that computers are more secure. Believing that computers are more secure, people are likely to be more willing to accept instructions to perform risky actions. Risk compensation may well interact poorly with user interface improvements. We'd like to be clear: we are not arguing that people are stupid or careless. We are arguing that people are hard to change, and that we need to strive to help them in full recognition of these difficulties.

The Use and Abuse of Language

A great many of the words we use when discussing security, including trust, threat, risk, safety, privacy, and security, can have multiple meanings. Each is evocative and carries with it cultural baggage. We often find ourselves talking past each other because of the inexact nature of these terms. This is not an argument for prescriptivism in language. Languages are successful when and because they are vibrant means of communication. If we can think and speak clearly, we can do so in spite of imprecise terms. If we can't think clearly, having precisely defined terms won't help us.

Language can be abused, and it *is* abused. Chapter 2 discussed some of the sales tactics used within the commercial information security industry. Describing a product as "secure" reinforces the fallacy that security is somehow a binary value—that something can be either "secure" or not. That kind of black-and-white distinction works with, say, pregnancy, but not for security. Without active intervention, the security of a computer system degrades over time. This happens because new vulnerabilities emerge that can affect it, and because of a process akin to natural decay in which operational changes become security issues. Something that is "secure" can at the most only be said to be "secure right now." What is "secure"

today is unlikely to be "secure" tomorrow. Another example is referring to certain security architectures as having an "assured" security model. In fact, no security can unequivocally be "assured." In cryptography, a debate is raging over the use of the term "proven," for much the same reasons.

Some security practitioners understand that when they refer to something as "secure," they are implicitly including an unstated corollary of "...depending on this, that, and the other thing." Trying to define this, that, and the other thing—the external factors on which the security depends—is a game of infinite regression. The term "secure" might be seen as a simplification to cope with the situation's inherent complexity. This abstraction makes it easier for people to function practically in their jobs, but not everyone understands that subtlety. The preceding section discussed the challenge of making a system "secure and usable." We spent quite some time discussing a way to say this without using the word "secure." In the end, we decided to hope that you would see it as an example of a place where "secure" is easier to say, while glossing over underlying complexity.

Security companies often invent new terms for things. "Pharming" is a name for attacks against the Domain Name System. The meaning of "pharming" is not obvious. That makes it a poor name. The same criticism can be leveled against other terms within security, such as "pretexting." This was the technique used to illegally collect information about the Hewlett-Packard board of directors in 2006. Pretexting is actually "social engineering," which is just another word for lying.

Arguments about terminology have been unresolved for many years, and we will not solve them here. Attempts to create strictly defined vocabulary within information security are likely doomed to failure as long as English remains a living language.

Skills Shortages, Organizational Structure, and Collaboration

There's never enough time in the day. Much like our allocation of resources and dollars reflects our priorities, the allocation of time and training shows what skills an organization values.

There is much talk in the security industry about the need for more security practitioners. Given the state of information security in the world, the value of having more people to perpetuate the status quo might be questionable. Chapters 2 and 6 discussed how security needs merge into, and in many cases form a subset of, traditional activities within information technology. It seems that, rather than needing more people to think about security all the time, we need people with responsibilities in areas such as enterprise architecture, planning, and operations to think more effectively about security.

This change has two primary aspects. First, as long as secrecy and tribal knowledge are at the core of the information security field, practitioners in other fields will be unlikely to consider security. Tribal knowledge doesn't help us answer questions and justify our answers. Using objective data, security decisions can be more easily understood, and their benefits more easily expressed. Second, companies should consider how best to structure themselves to allow security thinking to spread throughout the organization.

Chapter 1 described how "phishing" attacks exploit the difficulty of distinguishing between real and fake email, and between real and fake web sites. We spoke about how some organizations send email that is difficult for their customers to authenticate. These organizations have no conscious desire to confuse their users and put them at risk, but within those organizations, the left hand is unaware of what the right hand is doing. The team responsible for communicating with customers does not collaborate with the information security

team. This is in fact another example of an externality. When an organization sends mixed messages about security, the cost is borne by the customers of that organization. Customers might rightly be confused by an organization that deploys new security technologies to increase the security of its web site but then sends email that can lead to phishing attacks. Such costs also reflect on the organizations at fault. They take the form of increased customer support costs, customers using channels such as the phone that are more expensive but more easily understood, and perhaps a few customers leaving. (Understanding and measuring the externalities imposed by companies on their customers and others might be a valuable research project for a consumer advocacy group.)

Put frankly, the availability of objective data and the embrace of other sciences does not solve problems of organizational structure or a lack of collaboration. In its favor, the New School is a set of ideas that can be understood and applied at all levels of the organization, regardless of technical aptitude.

In Conclusion

We have spoken throughout this book about challenges within the field of information security such as viruses, spam, and identity theft. We have also discussed structural issues within the security industry in which the traditional ways of approaching problems can actually cause them to be sustained. The New School reveals new approaches to many of these problems. As in other fields, however, some seem surmountable, whereas others appear more difficult or even intractable given the environment that created them.

We feel it is important to describe the challenges still ahead. To do otherwise would be to peddle the next panacea. Where economic models break down, or where decisions are

made ignoring data, those failures must be identified and used to craft better approaches. That approach *is* the New School: to identify causes of success or failure and use that information to improve.

A collective effort is still needed. We must continue to work. We can make things better.