Foreword

From Kelvin's "[W]hen you cannot express it in numbers, your knowledge is of a meagre and unsatisfactory kind" to Maxwell's "To measure is to know" to Galbraith's "Measurement motivates," there is little need to argue here on behalf of numbers. Doubtless you would not now be holding this book if you didn't have some faith in the proposition that security needs numbers.

But what kind of numbers? Ay, there's the rub. We need numbers that tell a story and, which is more, say something that allows us to steer for where we are going, not just log from whence we have come. We have to acknowledge the central creed of the statistician: all numbers have bias; the question is whether you can correct for it. As security practitioners we have to know our place: Security is a means, not an end. We have to share with preachers and teachers the understanding that the best must never be the enemy of the good. So let's begin by laying bare the bias of this essayist and this book alike.

In *The Book of Risk*, Borge reminds and instructs us, "The purpose of risk management is to improve the future, not to explain the past." The past is a beautiful thing, whether we read about it or stand reverentially in the graveyard, but it is not what we, or this book, can afford to fix our gaze upon. No, we have to manage risks, seeking that optimal set point between, on the one hand, accepting silly risks and, on the other, burning our entire fortune fleeing that which cannot be escaped. Borge goes further: "Risk management means taking deliberate action to shift the odds in your favor—increasing the odds of good outcomes and reducing the odds of bad outcomes." This is my job description, this is your job description, this is our job description: to shift the odds in our favor. Need we remind ourselves that our opponents understand this 'pert well?

Need we remind ourselves that our opponents pick the places where the odds are in their favor?

To change the odds we have to know what those odds are, and we have to be able to detect when the odds change under our influence. To do this, we need security metrics. The job description for those security metrics is this:

> *Security metrics are the servants of risk management, and risk management is about making decisions. Therefore, the only security metrics we are interested in are those that support decision making about risk for the purpose of managing that risk.*

If you want to argue with that, this book is not for you, and I know without asking that Andrew would agree. Neither he nor I want to waste time convincing folks who have no use for numbers that numbers are somehow cool; what Andrew does at length and I am doing in shorter form here is to say that there is no need for despair—quite the opposite: This is the that idea whose time has come. It is an idea whose time has come so much so that despair of life without it or indifference to life with it is not even a luxury that you cannot afford. So let me compliment you for your intuition that this idea is timely (and how), as well as that this is the author to read (double and how). You are right on both counts.

Understand that numbers are insufficient, but they are necessary. To be the raw material for decisions, they need to be well enough thought out that you know what they mean. We have to be careful with what we claim to be measuring, and we have to make sure that our readers have some understanding of what we are measuring on their behalf. Numbers can mislead if they are not understood. Numbers are much like security in that they are a means rather than an end. Numbers exhibit vulnerabilities like computer systems in that whether misuse is intentional or inadvertent matters little if misuse is at hand. Numbers, like the surgeon's scalpel, can harm or heal, but we need them. As Fred Mosteller put it, "It is easy to lie with statistics, but it is easier to lie without them."

In November 2003, the Computing Research Association convened an invitation-only workshop on what "Grand Challenges" in digital security the National Science Foundation should concentrate a decade of funding on. There were four:

- No further large-scale epidemics
- Effective tools with which to certify systems
- Achieving low/no skill as a requirement to be safe
- Quantitative information risk management on par with financial risk management

That last one is why we are here. We need to do for the security arena what the quants have done for the financial markets: We need to understand, quantify, measure, score, package, and trade digital security risks as effectively as all the other risks with which the financial services sector already deals.

Finance has a concept—namely, that of "value at risk" (VaR), which is a daily number summing up a bank's exposure to loss. VaR is not without its faults (and Andrew takes up this idea later), but the core idea is that there is a target for risk and a metric that says how far off the target you are. It was my privilege to sit through a full review of a leading bank's VaR calculation. At the end of the day, the top economist leans over the lectern and says to his peers, "Now, you may ask yourself why all this works. (Pregnant pause.) It works because there is *zero ambiguity* about which of you owns what risks." At that moment of epiphany, I realized that that is what separated his field from mine—in our field, there is nothing *but* ambiguity about who owns what risk. We will not achieve the meaningful equivalent of a VaR calculation, and we will not fulfill the NSF's Grand Challenge, unless and until we have a way to score the game.

For any game, without a way to score the play, you cannot improve your performance as a player. That is where we are today: no way to score the game and no way to improve our play. This is not just a failing; it is a risk in and of itself. If we cannot make headway on measuring, on scoring, on understanding our risks well enough to change the odds in our favor by backstopping decisions about risk, we will have created one of those vacuums that Nature abhors. If we cannot measure risk and predict its consequences, the public, acting through its legislatures, will simply assign all the risk to some unlucky player. If assigning the risk and liability does not kill the unlucky assignee, from that point forward new players will find a barrier to entry like no other. In plain terms, innovation in the digital sphere and in the Internet itself is what is at issue here. If we cannot find a way to measure the security problem, I am afraid our choices will become difficult.

Some would say that regulation, not measurement, is what is needed. Some would say that the answer is to somehow force more investment in security—that, if this is an arms race, the good guys can win, because at the end of the day, they can outspend the opposition. That may be true in the physical world, but it is a dangerous delusion to think it can work in the digital world. In the digital world, the defender's work factor is proportional to the sum of all the methods the attackers possess times the complexity of that which is to be defended. The attacker's work factor is the cost of creating new methods as fast as old ones must be retired while complexity ensures that the supply of new methods can never be exhausted.

This asymmetry does not allow an "Outspend Them" strategy. Writing on 23 October 2001, six weeks after 9/11, the then-Chief U.S. Economist for Morgan Stanley wrote in the *New York Times* (in paraphrase) that "The next ten years will be a referendum on whether we consume the entire productivity growth of the U.S. economy for increased security spending." Perhaps the U.S. or some other country could endure that, but if it does, security, whether that means us or armies, becomes a fundamental rate-limiting block to wealth creation. Being a block to wealth creation is a side effect that cannot be sustained by us or anybody else.

For these reasons, the canon of digital security is now that of cost effectiveness, and the analysis of cost effectiveness runs on the fuel of measurement: measurement of inputs and outputs, of states and rates, of befores and afters. There is no greater need in the digital security sphere than to have a set of definitions of what to measure and to measure against those definitions. This will produce no miracles; human nature ensures that: The insecurity a computing monoculture engenders is long since proven theoretically and long since demonstrated empirically, but those proofs have not caused much behavior change; the buying public still longs for a technical fix, and the sellers of technical fixes are glad they do. But we—all of us who practice in this space—must find a way to measure and thus manage the risks that are rising all 'round us. My friend, the Red Queen, told us to run faster and faster if we wanted to stay in the same place. Evidently we are not yet running fast enough, because we are not achieving stationarity.

It is my fervent hope that Andrew's book—this book—sets off a competitive frenzy to measure things, to fulfill that Grand Challenge of a quantitative information risk management that is as sophisticated as what our brothers and sisters in finance have constructed with the very computers that we are now trying to protect. A competition of that sort will be brutal to some, but I, for one, would rather compete in that sphere than endure competition based on the side effects of needless complexity—complexity crafted by knaves to make a trap for fools. I want to compete on my terms with those who would manage risk at the minimax of lowest marginal cost and highest achievable control, rather than compete with the legions of digital doom on their terms.

This is your chance, Dear Reader, to be present at the creation or at least to be a passive consumer. Andrew, by writing this book, has erased any chance you thought you had to say, "Can't be done" or even "No one is doing that." As soon as this ink dries, both of those excuses are terminated with extreme prejudice. We, you and I, cannot waste our time on wishful thinking. To go on from here we cannot use words; they do not say enough. If you want to be a hero, make *your* book as much a step function as this one is.

Daniel E. Geer, Jr., Sc.D.
Cambridge, Massachusetts
November 2006