A close-up photograph of a bee on a green leaf. The bee is positioned on the left side of the frame, facing right. The leaf has prominent veins and some small dark spots. The background is slightly blurred, showing more of the leaf and some dark shapes.

THE EVOLUTION OF SECURITY

**What can nature tell us about
how best to manage our risks?**

DANIEL E. GEER, VERDASYS

Security people are never in charge unless an acute embarrassment has occurred. Otherwise, their advice is tempered by “economic reality,” which is to say that security is a means, not an end. This is as it should be. Since means are about tradeoffs, security is about tradeoffs, but you knew all that.

Our tradeoff decisions can be hard to make, and these hard-to-make decisions come in two varieties. One type occurs when the uncertainty of the alternatives is so great that they can’t be sorted in terms of probable effect. As such, other factors such as familiarity or convenience will drive the decision. This, too, is as it should be.

The other type of hard-to-make decision is when one must choose between a probable risk with tolerable cost and an improbable risk with intolerable cost. In metallurgical terms, this would be akin to “hardening” steel where *harden* can mean either *toughen* or *embrittle*. The tough steel will show every ding but will not fracture, whereas the brittle steel will show no dings but can be made to break. Perhaps this is best shown in a table where the implied risk cost is contrasted with the direct protection cost. Such a table is suitable for those who make dispassionate cost-based decisions.

Table 1 is OK as far as it goes: it implies that risk management is not all that hard when we enjoy easy precision about the costs of alternatives since the decisions are easy to make—if, but only if, you know the numeric dividing line between columns and between rows. Most days, we don’t.

Turn instead to figure 1, where the tradeoff is not a dividing line but a curve, and the subject matter is our subject matter: digital security.¹

Figure 1 is intended to illustrate exactly one idea: the total cost of a security regime is itself a tradeoff between the costs of protection (anticipation costs) and the

THE EVOLUTION OF SECURITY

Table 1 Categoric Cost-Based Tradeoffs

Cost of Protection	Downside Risk	
	Tolerable	Intolerable
High	bear the risk	capitalize
Low	whatever	do it yesterday

costs of nonprotection (failure costs). The expected cost for protection (anticipation) rises as the level of desired information assurance rises (the black line). Similarly, as the level of information assurance falls, the expected costs of cleanup (failure) rise (the red line). The curve of interest is the sum of the two (the green line). In the language of policy wonks and statisticians alike, what you want is a “minimax” solution—the *maximum* good for the *minimum* evil. Here the minimax point is where the sum of the cost curves for anticipation and failure reaches a bottom, as shown by the short vertical line.

Before we go any further, let’s review the ideas so far:

- Security is a set of tradeoffs.
- The existence of tradeoffs is why security = risk management.
- In the real world, tradeoffs are measured in cost.
- Cleanup and prevention are both necessary but neither is sufficient.

FAILURE MUST BE AN OPTION

The line of thought presented so far has interesting implications. Here’s one: The optimal number of security failures is greater than zero. If it is zero, then you are spending too much on protection. This is just as true for the number of medical side effects, the number of plane crashes, and the number of swindles. The (polite) term of art for spending too much is *cost ineffective* and, of course, what we want is to be *cost effective* (CE). Let’s be precise, however, about what that means.

Many readers will have heard the term *cost-benefit* (CB), which is a common way to describe whether some-

thing should be done, especially among policy types. The cost-benefit ratio is just what it sounds like:

$$CB_{ratio} = \frac{Cost_{new\ strategy}}{Benefit_{new\ strategy}}$$

This is favorable if $CB_{ratio} \leq 1.0$. Basically, this says that if the cost is less than the benefit, then take the benefit.

Cost-benefit analysis requires pricing the cost and the benefit on a common scale so that you can ask whether you would rather have the money (avoid the cost) or the benefit (incur the cost). This can be difficult as it leads to questions such as, “What is the dollar value of a human life?” “What timber-quality price will you pay to preserve wilderness?” “Is it worth paying for code compliance in affordable housing?” In contrast, the cost-effectiveness arithmetic looks like this:

$$CE_{ratio} = \frac{Cost_{new\ strategy} - Cost_{old\ strategy}}{Benefit_{new\ strategy} - Benefit_{old\ strategy}}$$

In other words, where cost-benefit asks whether you would rather have the money or the benefit, cost effectiveness assumes that you will, indeed, spend the money and thus your interest is in how much benefit you can get for your money, not whether you would rather keep your money in the first place. This means asking questions such as, “Would you save more lives by spending the \$10 billion on safer cars or on law enforcement?” “Would you get better availability by spending the \$1 million on 10 percent uptime or on instant recovery?” “Would your own pursuit of happiness lead you to spend \$100 on one fine dinner or on 20 lunches?”

CE is always tractable; CB is tractable only when the conversions of benefits to dollars are stable and noncontentious. To be blunt, CE is worth doing and CB is not. CE is decision support; CB is self-congratulation. If we are doing risk management rather than contemplating our navel or pandering to the electorate, then we must make decisions about allocating scarcity. We must remember that the purpose of risk management is to improve the future, not to explain the past.²

THE SPECIAL CASE

This brings us to the special case, the one that some of you were doubtless wondering about—namely, the issue of cost-effective risk management in and around software monocultures.

I've argued elsewhere that the only two classes of threats that matter at the national scale are:

- High-powered attacks on those few entities that for design reasons must be single points of failure—say, the literal root of DNS (domain name system).
- Cascade failure among the many.

The former is not a technical problem; it is a referendum on the willingness to spend the money required for defense in depth. The latter is where the action is, and where the monoculture question is centered.

The single most valid argument, as well as the single most head-scratching question, is whether one's security is advanced or retarded by having all computers just alike or all different. There are advantages to each. When they are all alike, the inherent ease of management is such that you might actually be able to manage them all, including risk management, using industrial-scale automation. On the other hand, when they are all different, there is no pathogen that can get them all. Perhaps table 2 will make this clear.

Let's go back to the cost-effectiveness point and think in terms of minimax solutions to tradeoff problems. We have risks, costs, and benefits from the all-alike alternative, and we have risks, costs, and benefits from the all-different alternative. Where's the tradeoff? What is cost effective? Is this a new problem never before seen? Is

there an answer? The answer is staring us in the face; the answer is in nature.

NATURAL LAW

Readers of *Queue* hardly need to be reminded that monoculture risk is real, that diversity can make coherent systems management challenging, or that risk management has to include tradeoffs around monoculture risk. There's nothing unique about digital security in that sense: farmers rotate their crops to do their kind of risk management. Big manufacturers second-source every critical part to do their kind.

Simulation studies done at George Mason University demonstrated that when about 40 percent of computers are alike, the risk of general collapse takes a leap upward.³ What a surprise! (Not.)

The science of disease prevention is often a search for what "model" can be used in developing this or that intervention—say, the bladder of a South American toad to study ion transport in the human kidney. Perhaps nature has a model or two for use in assessing monoculture risk.

The first observation is easy: Diversity accumulates over time if it is not edited by climate. In north temperate boreal forests, there might be 10 species of trees per acre. In Amazonia, there might be 200. The rain forests are the most ancient biomes we have, and they are not edited by the 110° F temperature swings that occur farther north, yet they have more predators per acre. In the north, the climate controls the vegetation. In Amazonia, the vegetation controls the climate (the outfall volume of the Amazon River divided over its watershed shows about the same rainfall input per unit of land area as is seen in Wyoming).⁴ There are lessons to be learned from the forest model, if only we try hard enough, but there is already a better model available.

Social insects are evolution's most fantastic success. In the perspective-of-scale department, there are approximately 2^{60} individual insects on this planet, of which 2^{50} are ants. Ants plus termites make up perhaps one-third of the biomass of all terrestrial animals.⁵ The economic benefit to agriculture from a single species of social insect, the honeybee (*Apis mellifera*), exceeds the economic losses to agriculture from all other insects combined.

Honeybees are a temperate species and they differ from other temperate wasps in an important way: their colonies persist in whole through the winter months, whereas the other wasps will, at the season's end, convert the entirety of their food stores to sexually mature adults who disperse to overwinter as they may. Think of the other wasps as annual plants and the dispersed adults as

Table 2 Monoculture vs. Diversity

	Nothing Happens	We Have Ignition
Monoculture	you win big	it's all over
Diversity	you wasted money	survive and gloat

THE EVOLUTION OF SECURITY

seeds. Think of a honeybee colony as a tree that retains the summer's food surplus for the next spring and thus lives through the winter in whole.

The downside to living through the winter is that diseases are not shed. A social insect colony is a veritable petri dish: moist, warm, open to the air, filled with great masses of soft-bodied young mixed with high-calorie, high-protein foodstuffs, and in continuous production. Disease pressure is thus, unsurprisingly, the greatest regulator of colony health, at least until an asteroid arrives. Honeybees are no exception: disease pressure is the most important contributor to colony death.

Honeybees, however, do something different—something interesting—compared with other wasps. According to the freshest research, most social insects have singly mated queens, but in honeybees each queen mates with numerous males to create a colony with a genetically diverse work force.⁶ The adaptive significance of polyandry by honeybee queens has been an evolutionary puzzle: mating with numerous males (on average, 12) requires that the virgin queen leave the hive and fly sufficient distance to find males, mate with them on the wing, and return home, thus exposing the hive itself to queen loss through predation of the queen, weather-related loss, and/or the possibility that the queen may be unable to find the home she has never before left. Where is the survival advantage in this? What eons-old risk-management tradeoff does this accomplish?

A singly mated queen produces worker daughters who are genetically identical. A multi-mated queen produces worker daugh-

ters who are composed of great blocks of identical sisters but where the blocks are half sisters to each other. In an elegant experiment, researchers showed that while both singly and multi-mated colonies had equal probabilities of infection, multi-mated colonies had much greater resistance to the effects of infection. In other words, multiple mating lowers the variance and raises the mean of the proportion of a colony's workers that survive, and it does so without lowering the rate of infection (susceptibility).

This beneficial effect of diversity is more pronounced when the disease is more virulent, and that is, if anything, yet more telling. If the pathogens of the Internet follow the evolutionary course of pathogens in biology, we should expect that selection pressure reduces the number of pathogens over time but that the surviving ones will be more virulent (i.e., the pathogens will spread faster and the genes [code snippets] for virulence will be shared).⁷ Not surprisingly, we have seen doubling times declining, as shown in figure 2.

PUTTING GOOD IDEAS TO USE

Social insects coordinate their life through a mix of genes and pheromones, the analogical equivalent of configuration files plus message passing. Disease pressure dominates colony health absent general breakdown of the environment, just as the background pressure of all the hundreds of automated attacks circulating on the Internet dominates absent the DNS completely failing. A monoculture is a *per se* genetic malformation wherever disease pressure is continuous and the survival of the fittest has

selected for nonzero genetic diversity, effectively valuing that diversity over the single-point-of-failure risk of queen loss during mating flights (whose sole purpose is to acquire that diversity).

The simulation studies of Gorman et al. have established an upper bound for how much computing monoculture you can tolerate without getting hosed, and the honeybee has selected for an upper bound on how much diversity to invest in. The NCMS (National Classification Management Society), an organization for industrial security professionals, has shown that all-protection is just as cost ineffective as all-cleanup, and logic tells us that the optimal number of failures is nonzero. Because the honeybee colony (enterprise) has permanently continuous disease pressure, unlike other wasps it invests in diversity of its endpoints, while ensuring that all the endpoints interoperate over the same (chemical) protocol. Because nature is parsimonious, the honeybee's degree of in-colony (inside the enterprise) diversity must be assumed to be an optimized minimum cost investment to maximally insure colony preservation against cascade failure.

Of course, this is speculative, but we may now be homing in on useful truths—compelling enough that the burden of proof shifts to those who would claim that nature is not a natural guide for computing. We computer types already copy nature by investing in centralized, feedback-based nervous control of the enterprise (sensor-fed operations centers)—control largely aimed at stasis. We already invest in primitive immune systems (intrusion prevention systems). We reproduce our computing tissue asexually by cloning some gold master somewhere, even though a pond full of identical blue-green algae can be thought of as success only when evolution is very young. We already have a pale kind of selective gene expression when installing enterprise-scale systems, though when we discover that there are 500 knobs (genes) to adjust, we tend to leave 90+ percent of them however they were set at the “factory” because, in truth, no one knows what happens if we reset all of them. We already run large data centers with board-level Linux machines that are simply thrown away when they look bad—repair is cost ineffective for the sysadmin staff (liver) to perform in a world where reliability is god.

It is time we learned something from the social insects because they are nature's success story *par excellence*. The purpose of today's essay has been to suggest that just as an evolutionary risk-management prerequisite to climb from single-cell organisms to multi-cell organisms is the self- vs. not-self-discrimination of an immune system, an evolutionary risk-management prerequisite to climb from

a multicell organism to a multi-individual colonial hive is a minimax tradeoff between the ease of control and the vulnerability resulting from unmitigated identity. When they, whatever “they” are, are all alike, their protective coating has to be flawless or they all die together. When they, whatever “they” are, are not all alike, they don't have to be perfect because they then have the law of large numbers on their side.⁸ Q

REFERENCES

1. Costs of Information Assurance. 2002. National Center for Manufacturing Sciences (August); <http://trust.ncms.org/pdf/CostInfoAssur-NCMS.pdf>.
2. Borge, D. 2001. *The Book of Risk*. John Wiley & Sons.
3. Gorman, S.P., Kulkarni, R., Schintler, L., Stough, R. 2004. Is Microsoft a threat to national security? The effect of technology monocultures on critical infrastructure. George Mason University, Infrastructure Mapping Project working paper; http://policy.gmu.edu/imp/research/Microsoft_Threat.pdf. (Full discussion of these results is outside the scope of this essay.)
4. Myneni, R.B., et al. 2007. Large seasonal swings in leaf area of Amazon rain forests. *Proceedings of the National Academy of Sciences* 104(12): 4820-4823.
5. Hölldobler, B., Wilson, E.O. 1990. *The Ants*. Cambridge, MA: Harvard University Press.
6. Seeley, T.D., and Tarpy, D.R. 2007. Queen promiscuity lowers disease within honeybee colonies. *Proceedings of the Royal Society of London* 274: 67-72.
7. Wassenaar, T.M., Blaser, M.J. 2002. Contagion on the Internet. *Journal of Emerging Infectious Diseases* 8(3).
8. Jones, J.C., Myerscough, M.R., Graham, S., Oldroyd, B.P. 2004. Honeybee nest thermoregulation: Diversity promotes stability. *Science* 305(5682): 402-404.

LOVE IT, HATE IT? LET US KNOW

feedback@acmqueue.com or www.acmqueue.com/forums

DANIEL GEER is chief scientist and vice president of Verdasys Inc. Early in his career, he oversaw development of the X Window System and Kerberos at MIT. Other highlights of his career include the first information security consulting firm on Wall Street, the first academic conference on e-commerce, and the seminal 1998 speech, “Risk Management is Where the Money Is.” Geer is a past president of Usenix and principal author of and spokesman for “Cyberinsecurity: The Cost of Monopoly,” a noted report published in 2003 by the Computer and Communications Industry Association. He is also co-founder of SecurityMetrics.org.

© 2007 ACM 1542-7730/07/0400 \$5.00