A Doubt of the Benefit

Cost-benefit analysis in security is appealing as a standard approach, admirable for its simplicity, appreciated for its generality, but otherwise worthless. Every cost-benefit calculation requires a consistent scale, and the more people this affects, the less they're likely to agree on whatever rescaling this forces. Thus, questions such as "What is a human life worth?" or, in our case, "What is a secure machine worth?" yield indefensible answers, which serve as an awkward basis on which to begin formal analysis. For the record, we believe our lives to be more valuable than standard governmental estimates.

Cost-effectiveness analysis simply assumes that you'll spend the money, so it asks "How many lives can you save?" or, in our case, "How much breakage can you prevent?"

DANIEL E. GEER JR.
In-Q-Tel

DANIEL G. CONWAY
Augustana College

Take public health; to find undiagnosed cases of familial hypercholesterolaemia, you could (looking at the number of persons tested to find one case)

- Screen every 16-year-old: 1,365
- Add screening to all doctors' checkups: 938
- Screen all heart attack victims: 22
- Screen family members of known carriers: 2.6

In other words, it's 525 times more cost effective to know where to look. If every potential life saved had infinite value, we would screen everyone. We absorb some risk in the name of being cost-effective. The more expensive the diagnostic test, the more profitable this kind of thinking.

In security, we already think like this to some small extent. If we find a vulnerability in a library, we treat the use of that library like a gene, and we try to chase down the "family members" who share that library. We can go further; consider whether to put a patch out or not:

- Use application scanner to get some risk index r_i
- Apply the manufacturer's patch, and rescan to get r_j
- Determine the rollout cost c_r
- Cost per unit of risk reduction = $\frac{c_r}{(r_j - r_i)}$

You can set a cutoff for whether a fix is worth rolling out solo based on the cost-effectiveness of doing so. Microsoft may use such a scheme for whether to go out of Tuesday order. Neither you nor Microsoft has to be scientifically perfect in setting that cutoff, only consistent.

Again, some diagnostic tests are especially expensive—manual code review by practiced experts, say. The benefit of such work is undoubtedly grand, but how do you price it? You don't. You do a first pass with a cheap automated code analysis calibrated for low/no false negatives, then a second expensive pass with your experts (who generate low/no false positives).

Suppose you have 10^7 lines of code (LOC), the automated test costs 1¢/LOC, the expert code review costs $10/LOC, and one line in 10,000 (0.01%) has a security flaw. Suppose your automated test finds 99.99% of the flaws but has 10% false positives, and, for convenience, suppose that your experts call a truly safe LOC safe 99.99% of the time but have 10% false negatives.

auto pass 1	Pr(test+|true+) = 99.99% Pr(test+|true-) = 10.00%		
	true+	true−	
test+	1,000	999,900	1,000,900
test−	0	8,999,100	8,999,100
	1,000	9,999,000	10,000,000
human pass 2	Pr(test-|true+) = 10.00% Pr(test-|true-) = 99.99%		
	true+	true−	
test+	900	100	1,000
test−	100	999,800	999,900
	1,000	999,900	1,000,900
combined			
	true+	true−	
test+	900	100	1,000
test−	100	9,998,900	9,999,000
	1,000	9,999,000	10,000,000

The first test, if used alone, would leave you with nearly a million false positives—too many to fix; the second test, if used alone, would cost you $100,000,000—completely unaffordable; but used together and in that order, you find 90% of the flaws for US$11,233.34 apiece.

This is how cost-effectiveness works. Cost-benefit approaches such as Annualized Loss Expectation (ALE) will always prevaricate meaningful comparison in security because organizations will never converge on asset value. Nor should they. Cost-effectiveness is the apposite approach if we're going to advance today's "good measures" and ultimately leverage the established measurement giants of actuarial sciences.

Since last issue, the ØPI has risen $182.30 (0.2%) to $67,315.30; US credit cards have dropped to $1 bulk and have even seen $0.99. Fullz are up to $10 from $5 for now, and Windows bots are up as well, which might be due more to a consolidation of suppliers in the market than to a lack of supply. Our nascent Security Pressure Index (see last issue):

INDEX	PREVIOUS	CURRENT	TREND
Phishing	544	553	
Spam	286	299	
Workfactor	113	112	
Dataloss	93	83	
Composite Security Pressure	216	262	

Daniel E. Geer Jr. *is the chief information security officer for In-Q-Tel. He was formerly vice president and chief scientist at Verdasys, and is a past president of the Usenix Association. Contact him at dan@geer.org.*

Daniel G. Conway *is an associate professor of business administration at Augustana College. He previously served as associate professor at the University of Florida and Virginia Tech. Contact him at danielconway@augustana.edu.*

Interested in contacting *IEEE Security & Privacy* about what you've read? Please email letters to the editor to Kathy Clark-Fisher at kclark-fisher@computer.org.

For more information on these or any other computing topics, please visit the IEEE Computer Society's Digital Library at http://computer.org/publications/dlib, or our online portal, Computing Now, at http://computingnow.computer.org.