ICS Update

Daniel E. Geer Jr. | In-Q-Tel
Mukul Pareek | Consultant

> It is easy to lie with statistics, but it is easier to lie without them. —Frederick Mosteller

The natal announcement for the Index of Cyber Security (ICS) first appeared in these pages one year ago. As we promised at the outset, its first birthday marked the time for a review.

The ICS is composed from a survey of expert sentiment—that is to say, it asks a set of respondents what they think. Sentiment-based indices have a long history and wide acceptance; two (US) examples are the Consumer Confidence Index and the Purchasing Managers Index.

Generally speaking, sentiment-based indices are vulnerable to misinformed respondents. This is conquered either by large-scale sample randomization (Consumer Confidence) or by careful selection of respondents (Purchasing Managers). The ICS goes with the latter: it gathers a composite of cybersecurity expert opinions that aren't generalizable to any description of the public at large.

As everyone here knows, definitions of terms in the security space are imprecise. To conquer the problem of subtle differences in definitions of, say, malware, the ICS asks each respondent each month the same set of questions, in the following form:

> Compared to the previous month, the unmitigated threat to you from malware (any/all types) is falling fast, falling, static, rising, or rising fast.

This kind of scoring is called a Likert scale; it doesn't require precise calibration of every respondent's definition of malware, just that each respondent has a stable definition that isn't from outer space.

Out of respect for respondents' time, we ask them to answer the same 20 questions each month. The ICS's structure is such that we can add, subtract, or swap a question without derailing the survey's continuity. The math we use for this is precisely the same used for a broad spectrum of financial indices.

Respondents' answers aren't traceable to them, even by us. In other words, kidnapping either of us wouldn't get you any info.

We've taken several steps to make the ICS a reliable gauge of collective expert opinion. The feedback we've received corroborates the ICS's value as a component of risk management. As always, data sharing is the only way to determine if the risk pressure you see is unique to you or a general phenomenon.

To be clear, then, the ICS is an index of risk: if respondents perceive that risk is rising, then the ICS rises. As distributed (cybersecurityindex.org), a year's worth of the ICS looks like Figure 1.

Put in words, our respondents believe that risk in the aggregate is rising steadily (the blue line is the ICS itself), although the month-to-month change (the red line) is not so steady.

Because each question reflects a component of cybersecurity risk, it's quite possible for one of the individual components to be important to the change in the main ICS one month and to be irrelevant to it in some other month.

Variation between component contributions to the overall ICS is, in fact, what we've found. If, for each month, we rank order an individual component's contribution to the overall ICS, we get the rat's nest in Figure 2.

This demonstrates that overall risk might be rising relatively steadily, but that individual components vary in terms of their contribution to the overall rate of rise of the composite ICS itself.

Looking at the components in a different way, consider the cumulative contribution of a particular component to the overall cumulative change in the composite ICS as in Figure 3.

Now we see a couple of interesting things, such as how regulatory pressure in the Americas is inching upward in its contribution to the overall risk as captured by the ICS, that counterparty and hacktivist risks are currently running neck and neck for the biggest contributors, and that failure of defense measures has steadily come down to where it has joined

1540-7993/12/$31.00 © 2012 IEEE Copublished by the IEEE Computer and Reliability Societies May/June 2012

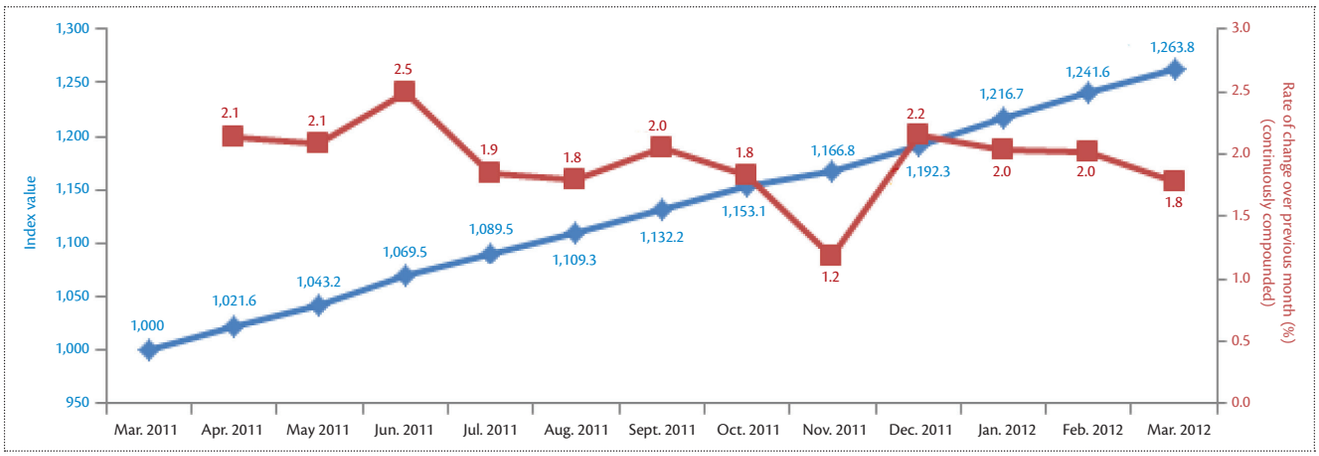

Figure 1. Year one of the Index of Cyber Security.

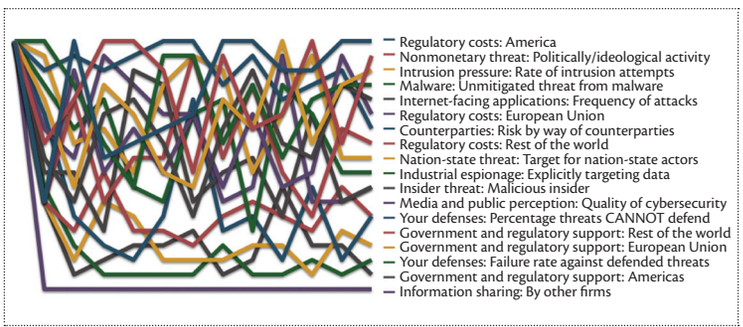

Figure 2. Rank order of month-by-month component contribution to the ICS.

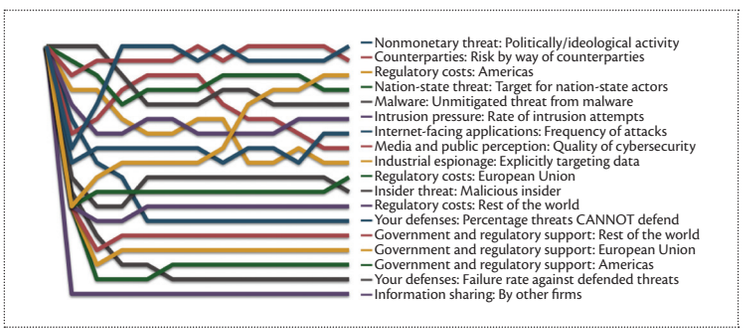

Figure 3. Rank order of cumulative component contribution to the ICS.

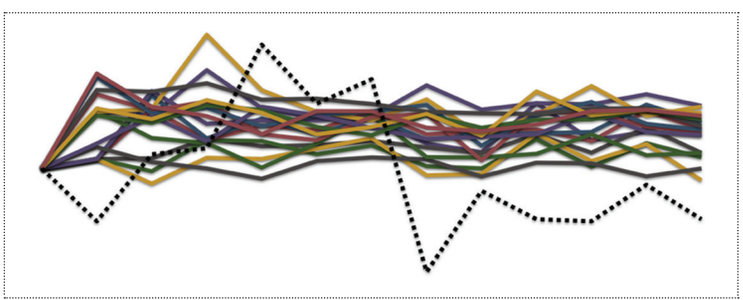

Figure 4. Superimposing VIX on month-by-month component fluctuation.

actionable information sharing by other firms besides your own as those areas with little contribution to the aggregate risk at all.

As we said at the outset, our plan had been to touch nothing for the first year. We're at that one-year point, and we're looking at changes to a few questions, but we'll do those incrementally as the months progress. At the cybersecurityindex.org website, you can see how we do the calculations for month-over-month symmetry and for question swaps. Again, you'll find nothing methodologically interesting here—everything is conventional financial markets math. Boring is good.

Like any index, the ICS has limited predictive value. It's a measure of what industry experts feel the risk to be at a certain point in time. Just as the S&P500 today isn't a pointer to what its value will be tomorrow, the ICS can't be used to extrapolate what the future will bring. Such extrapolation, if performed, would be based on the shaky assumption that the future will continue to be like the past.

We've often observed a direct correlation between news stories and what our respondents reported as the threats with the greatest increases in risk. But which way does the causation flow? Were our respondents' reactions affected by what they were reading in the newspapers, or were the news stories reflecting what our respondents were telling the press? Suppose we superimpose the Fear Index (VIX, the Chicago Board Options Exchange's Market Volatility Index) on the month-to-month changes in ICS subindices; might we see something worth analysis, as in Figure 4? (We don't think so, but you get the idea.)

What we most want to do, however, is where you come in: expand the base of respondents. One thing we've learned is that people have good intentions, but

consistent participation has many other things working against it. An order of magnitude more participants might let us break out the ICS by, say, industrial sector or geographic continent. We repeatedly get requests for various breakouts, but we've so far resisted. Certainly one requirement if we're to reconsider breakouts is that of getting sample sizes big enough that we can still guarantee nontraceability of any individual respondent's answers. Certainly another requirement is to only do breakouts where the implied security metric is actually of decision support value to the general community that consumes the ICS number or numbers; we aren't in this as stamp collectors but as decision support engineers.

We're asking some readers of this column to volunteer to be respondents. To be clear, we aren't looking for just anyone—we need people with direct, hands-on operational responsibility for cybersecurity, perhaps especially the person where cybersecurity data feeds of various sorts from elsewhere in the corporation converge. We want the opinions of people who are data driven but who have to have opinions because the data they get isn't intrinsically informative enough for decision-making as received. We aren't looking for researchers, policy people, executive management, or general counsels. We're looking for CISOs, people whose system administration responsibilities include security administration, and any person who, for whatever idiosyncratic reason, has a consistently current view of frontline operational reality.

If this is you, please be in touch. We can't pay you, but we can thank you with the coin that we have. Respondents get a monthly report that's more detailed than the report we put up on our public website. Among other things, you get the spread of answers question by question and analysis of the subindices (this column is a taste). We'd like to think that respondents would volunteer to be respondents simply because it's a good thing to do for our profession, and the detailed report we give back would be received as a professional courtesy, but if you want to think of this as bartering your data for our data, then feel free to do so.

Our hope is for the ICS to have a permanent value to the cybersecurity profession. Because there are only two of us, an errant bus could indeed wipe out the entire team. As such, we are evaluating where and how a permanent home might be found. We expect tradable instruments pegged to the ICS to appear, just as with Consumer Confidence, Purchasing Managers, and other indices, but that's a discussion for another day. ■

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.

Mukul Pareek is a New York–based risk professional working in the financial services sector. He has previously worked in audit, consulting, and industry. Contact him at mp@pareek.org.