Join the Discussion
Connect

# Maddening Methods:

## Fundamentals of Risk Assessment and Analysis

By Benjamin Tomhave – ISSA member, Northern Virginia, USA Chapter

**This article seeks to level-set on the contentious topic of risk assessment and analysis, providing information that is often left out in anti-risk-assessment arguments.**

### Abstract

Considerable confusion exists in the security industry around the effectiveness of risk assessment and analysis methodologies. Points of contention often focus on specific attributes of a given method, such as data quality, statistical analysis, or a qualitative versus quantitative approach. There are reasonable, viable answers to these points of contention that resolve most of these concerns.

There has been a lot of negative, cynical chatter lately about risk assessment and analysis. The average person does not understand it, and people who should understand it oftentimes throw up their hands in despair when citing examples such as the failures of Wall Street that led to the current economic mess or the aftermath of the Deepwater Horizon[1] incident. Unfortunately, all of this despair and cynicism seeks to throw out the baby with the bath water, as if to say that one bad apple spoils an entire orchard. This article seeks to level-set on the topic, providing information that is often left out in anti-risk-assessment arguments.

Beyond these fundamental concerns, it seems that some of the biggest challenges to risk management lie in a few key areas: accountability, consequences, and formalized assessment methods. The first two areas are easy to explain. If you are doing a good job assessing and managing risk, then you can hold people accountable for their decisions and actions. Accountability should link to consequences (positive or negative) and ultimately to success or failure. Unfortunately, this modern era seems to be one in which failure is feared, causing us to insulate ourselves, our families, our investments, and our world from negative consequences. Without negative consequences, what is the point of managing risk?

This last area is where much of the focus has turned as of late in the security industry. Formalized risk assessment methodologies are still generally immature, and they can frequently be problematic. However, many of the arguments made tend to be divisive at best and willfully ignorant at worse. We should be very concerned about this last type of argument, because it tends to lead to a path of FUD, fueled by pseudo-experts who stir the pot with confusion and veiled intentions.

## A bit of background

Before going into the various arguments about risk assessment and analysis, it is first important to know a little bit about the conversation, key players, and history.

### Key players

**Quals** – This group of people makes use of *qualitative* risk assessment practices. That is, rather than using numbers and calculations, they instead develop rubrics that are descriptive in nature. These descriptive scales may be mapped to numbers, but the relationship is subjective and arbitrary.

**Quants** – This group of people makes use of *quantitative* risk assessment practices. They rely heavily on statistical methods, seeking to put numbers, math, and science behind their reasoning. Their approaches are not widely understood because of the underlying complexity, leaving the door open to criticism.

**Risk cynics** – There is an increasingly vocal group of cynics that repeatedly make arguments about how risk assessment is a failed discipline, how it will never succeed, and pointing out what they see as fatal flaws in the various approaches. Their arguments tend to be monotonous and repetitive, and you will note that they generally just tear down without offering viable alternatives.

**Indies** – Finally, there is a group of independent thinkers who hold out hope for finding future solutions that will resolve the concerns of the other three groups. This group is actively working to help resolve concerns while maintaining a healthy degree of skepticism about the current and near-term states.

[1] http://en.wikipedia.org/wiki/Deepwater_Horizon.

### Common attacks

More details will be provided throughout the article, but following is a summary of some of the common attacks on risk assessment and analysis methodologies:

**Inadequate data** – One of the most common arguments is that there is not enough data from which to derive reasonable estimates for anything (loss, probability, frequencies, etc.). The actuarial tables leveraged by the insurance industry are frequently cited, with a quip that since no such thing exists for information risk management, nothing can be done.

**Faulty value/loss estimates** – Related to the first argument, this argument keys in on the estimates used to measure the impact of a loss event and suggests that it is not possible to reliably estimate the impact of an event like a breach. We might know how much it would cost to monitor, discover, and recover from a breach, but the enduring impact, such as to stock price or consumer confidence, is fuzzy at best.

**Faulty probability estimates** – Also related to the first argument, this argument looks specifically at the probability estimates typically used in risk calculations and says "there are too many unknowns – especially unknown unknowns – to make these estimates even remotely reasonable." With no basis for estimating probabilities, the probability models "fall apart." This line of argument tends to lead to the next quip.

**Unknown unknowns** – This argument states that we essentially do not know everything, and thus cannot know anything. Given that the world is infinite, there are threats and vulnerabilities that have not been envisioned, which means that the likelihood of occurrence cannot be estimated, let alone what the impact would be to the business. Fortunately, life is not constrained by a need for "perfect" information.

The main notion to bear in mind with risk management as a discipline is that it has been around for a long time and is, in fact, very mature. Information risk management is indeed a relatively new subset within the overall discipline, and it is suffering through growing pains as one might expect, but that does not nullify the entire discipline.

## Data quips

The most common starting point for criticism of information risk management is to target the data.[2] Common complaints are:

- There is not enough data
- The data is not reliable
- The data is not consistent

Each complaint is valid, at least to a point. However, like with most weaknesses in information security, there are ways to mitigate these concerns (not entirely, but to a level that makes it acceptable and useful).

### Not enough

The first quip, particularly from the risk cynics in the crowd affectionately referred to as "concretists," is that there simply is not enough data. If only we had reams upon reams of actuarial data like ye olde insurance firm, then perhaps we might be able to make use of it. This argument ignores current practices in statistics and probability, and really looks for an easy "out" from the discussion by assuming (incorrectly) that having a small data set limits the ability to use that data.

First and foremost, some data is better than no data. Second, even if that data is "salted" (i.e., not all good), it is still useful. Through the mathematics principles of Bayesian statistics, we can still run calculations and create distributions using small data sets. What we will find, particularly through use of tools like Monte Carlo simulations, is that less data may decrease our degree of confidence in the data (i.e., the scatter plot may be a bit scattered). However, as our data grows, scatter plot focus will improve, reflecting the consistency and effectiveness of our model and data.

Second, you have to start somewhere. The sad part of the "not enough" argument is the corresponding "do nothing" alternative proposed. Or, more correctly, rather than working to adopt a scientific, repeatable method, opponents seem to argue in favor of a less rigorous approach (such as advocated by the quals), which in reality leads to arbitrary assessments lacking grounding in a given organization's context. If you have good data, why not use it?

Last, there is a question of how long we have had to develop data. There are a number of reasons why there is limited data on breaches, impact, threat frequency, and the like. Chief among those reasons is that we are really only talking about 15 years of computing history. Consider for a moment how much the business world has changed since 1995. In 15 years we have gone from very limited, if any, Internet connectivity in the workplace to ubiquitous connectivity.

### Not reliable

Apparently what we need in this world is perfection. Is it not the ultimate goal? It is also unrealistic. There is no such thing as perfect data (or practices, or methodologies). Yet this does not stop people from criticizing the sources of data and how the data was collected or classified.

As already stated, some data is better than no data, even if we recognize its imperfect state. Also, as just noted above, there is at most 15 years of run time during which data could have been gathered. Realistically, it is even less than that – probably in the range of 10 years. Critics complain that the data we have collected is unreliable, often citing software engineering theory, but the simple fact is that we must do the best we can with what we have (all while developing better data).

The data is adequately reliable, assuming one knows how to break it out and use it. Look at the Attrition.org and DatalossDB.org archives. Adequate data is available for analysis, including perform trend analysis. There is no reason that this

[2] For an excellent discussion of risk data, please see Alex Hutton's post "Risk Appetite: Counting Risk Calories is All You Can Do" on the Verizon Business Security Blog at http://securityblog.verizonbusiness.com/2010/06/17/risk-appetite-counting-risk-calories-is-all-you-can-do.

data cannot also be leveraged for risk assessment and analysis.

The key, however, is making sure that confidence is factored into calculations and that single absolute numbers are shunned. If you know that the data is a bit unreliable, then you can account for that case explicitly. To that end, it is typical to work with ranges instead of single numbers (see Douglas Hubbard's *How To Measure Anything*[3]). The tighter the defined range, the higher the confidence in that range, which will then show through in calculations and visualizations. Notice that this technique allows the use of potentially unreliable data, rather than simply throwing everything away and working from a purely subjective perspective.

### Not consistent

One of the more valid concerns about data is its consistency. This context pertains to the consistency of data collection and classification. Data breach reporting is a perfect example where, in lieu of standard reporting requirements, the same types of data are not reported with each incident. This is a problem that organizations like DatalossDB.org encounter on a regular basis. Even if a standard data collection approach is used, such as Verizon Business's VerIS Framework,[4] consistency challenges may still be encountered when comparing one repository to another (e.g., comparing data between Verizon DBIR,[5] Veracode,[6] WASC,[7] and DatalossDB.org).

This challenge underscores the need for mandates and standardization around data breach reporting, in particular, but it also highlights the need to be cognizant of the data's source before acting on the data. It is desirable to gather as much data as possible, and for it to be as reliable as possible, which then means that extra effort must be expended in standardizing data sets to ensure consistency and to help weed out bad data. Little tweaks, such as ensuring consistent rounding of numbers, can go a long way toward helping ensure that data sets are more usable and useful.

In the end, of course, we come back to the same quip as above: some data is better than no data, even if our data confidence is only moderate. Once you have a start, you can then refine your models and data sets over time to ensure better quality data, and to improve your overall analysis.

### Cast a wider net

The last point to make with respect to data is to ensure that the right data sources are being sought. As useful as the various reports are in evaluating probabilities relative to threat frequency and vulnerability, most of these numbers are limited primarily by industry and organization size, and they do not address estimating loss magnitude.

However, there are ways to develop adequate loss magnitude estimates within a given context if the right sources are sought out. The sub-text is this: get outside of IT and information security; they are not experts in business analysis. Instead, it is of vital importance to seek out other subject-matter experts in the organization being assessed to properly gauge the values and assumptions that go into primary and secondary impacts. Business managers understand business costs far better than information security teams do, just as the legal team will understand legal liability far better than the average security analyst.

Finally, again, it is key not to try and work with absolute numbers, but to instead make use of ranges that meet a reasonable degree of confidence (typically targeting 90% confidence – see Hubbard[8] for more on this topic). It is also important to note that the goal of this exercise is not to predict the future, but rather to provide reasonable and adequate data points with which to make informed decisions. Risk assessment methodologies should be seen as decision-analysis tools, not as some panacea of future forecasting.

## Method quips

The next most common argument against information risk assessment is that the methodologies are wholly inadequate. Typically this criticism is leveled against the qualitative methods out there, though quantitative methods are also criticized. There can be an appropriate time and place for qualitative assessments, such as part of a larger evidence-based risk assessment approach, but it is important to carefully separate and weight the results accordingly.

One of the primary problems with risk assessment methods is the historical reliance on Annualized Loss Expectancy (ALE). More often than not we end up getting ourselves into trouble by pulling arbitrary numbers out of the air, when instead what is needed is complete transparency and illumination to see how a number was calculated or derived (as with cryptography research, insight into how numbers are manufactured is vital to proving the integrity and reliability of the overall system or methodology). The source of numbers is rather important, especially when performing a quantitative risk assessment.

Too much time is wasted on this attack against risk assessment. ALE is not an end-all-be-all kind of number, and is really heavily abused. Frankly, it is simply not correct to use it on its own and out of context. More importantly, single numbers should be eschewed in favor of ranges. At the same time, it is also valuable to adopt Jack Jones' preferred approach of breaking these estimated impacts into primary and secondary.[9] It turns out that real, direct costs for a given security

[3] Douglas W. Hubbard, *How To Measure Anything*, 2nd ed., (Hoboken: John Wiley & Sons, 2010).

[4] See http://securityblog.verizonbusiness.com/2010/02/19/veris-framework-2.

[5] See the *2009 Data Breach Investigations Report* at http://securityblog.verizonbusiness.com/2009/04/15/2009-dbir.

[6] See Veracode's *State of Software Security Report* at http://www.veracode.com/reports/index.html.

[7] See Web Application Security Consortium (WASC) Security Statistics at http://projects.webappsec.org/Web-Application-Security-Statistics.

[8] Ibid., Douglas W. Hubbard.

[9] For more on this topic, http://riskmanagementinsight.com/riskanalysis.

incident can be estimated fairly well. It is the indirect costs where a much broader scatter is seen, prompting us to compensate for that uncertainty accordingly.

A specific quip observed recently was that there is "no sound method of actually measuring loss magnitude."[10] This quip is essentially half right. Yes, ahead of time it is extremely difficult to precisely estimate the combined primary and secondary impact. However, getting back to semantics, there are a few key points:

- *How much precision do we need?* As already discussed, using ranges can help improve estimates, which can then be used to perform statistical analysis on multiple data sets to improve these estimates. However, do not forget that we are not talking about an exact science (if it were, then we would not be having these arguments, nor would we be relying on statistical models quite so much). What is needed is enough precision to make quality decisions, but without believing in some mystical, magical "perfect" result that will solve all problems.
- *It helps to split impact between primary and secondary.* Direct costs can be estimated fairly well. Hardware, software, and resource time are all generally known quantities. There are reasonable estimates on how long it takes to detect and correct major classes of issues. It is the secondary costs that introduce a higher degree of uncertainty. However, at the same time there is an ever-increasing data set, thanks to large publicized incidents, that facilitate making a reasonable guess at the secondary cost of an incident, even if it that estimate necessitates a wide range. It bears remembering that the goal is providing a good enough assessment result to facilitate making quality decisions.
- *Why the heavy focus on impact?* The focus on financial cost is natural to the business, but it also seems to have its roots in the long-since-debunked myth of Security ROI (or ROSI). Risk management decisions based on information risk assessment and analysis should not be oriented toward trying to estimate (or justify) a return, but rather on loss control/management and, more correctly, legal defensibility. Information security controls help defend against, and optimize recovery of, security incidents. Information risk management provides useful data points to see where to improve resource allocations to optimize defensibility and recoverability.

## Consistency: (not) the risk analysis panacea

Jack Jones of Risk Management Insight wrote a blog post in June 2010 titled, "Managing Inconsistency,"[11] in which he talks about the dream state of having consistency between assessments. In the dream state, two assessors will walk into an organization with their own tools and will produce results that have high degree of parity. That is, they will gather their own data, make their own calculations, and yet find that they get the same effective results. It's a nice idea, but is it really all that important? (answer: yes and no)

On the one hand, yes, we do need consistency. Without consistency we get into issues of integrity and bias (e.g., a recent study[12] showed that security managers tend to skew road maps to their own personal bailiwicks instead of doing what is objectively right for their respective organizations). So, yes, consistency is needed in order to help reign in some of the chaos that comes from implicit and explicit bias.

However, on the other hand, it is important to make sure that information risk management is not seen as panacea. Instead, information risk management is a tool in the overall toolbox that we need to use in information security, just like words are the tools we use to build and convey thoughts. The English language is very instructive on this point in that there are typically multiple ways to say something, getting the same meaning across, all while using different words or word-order (e.g., "My name is Ben" and "Ben is what I am called" are functionally equivalent, yet completely different sentences). What degree of parity is necessary?

The key point here, derived from the "Managing Inconsistency" post above, is that variance in risk assessment and analysis is manageable, and it is secondary to the overall outcome of the method and the ability of management to make meaningful use of the method's results. As with the data concerns, if an assessment is conducted while knowing that there is the potential (likelihood) for variance, then the variance can be compensated for programmatically. Over time, methodologies should become refined and better tuned to help reduce variance, but until then, compensating for it will have to suffice.

### Methods, methods everywhere

One last point to consider is that there are numerous methods, and they are not necessarily all the same or equal (e.g. FAIR,[13] TARA,[14] OCTAVE,[15] NIST RMF,[16] WFITW[17]). It is time to start moving more aggressively away from the "security is more art than science" mentality. WFITW is not a legally defensible strategy.

## Unknown unknowns

The final common argument – and by far the most inane – against risk assessment is that of death by unknown unknowns. The argument goes that, because we lack data, and because we do not know exhaustively and definitively what else is out there (in terms of threats, vulnerabilities, and attackers), then we simply cannot make any sort of reasonable estimate of anything. This argument is akin to the iceberg

[10] See the comments in the Securosis blog post "FireStarter: The Only Value/Loss Metric That Matters" available online at http://www.securosis.com/blog/firestarter-the-only-value-loss-metric-that-matters.

[11] See http://riskmanagementinsight.com/riskanalysis/?p=726.

[12] Reference unavailable.

[13] See http://fairwiki.riskmanagementinsight.com.

[14] See http://download.intel.com/it/pdf/Prioritizing_Info_Security_Risks_with_TARA.pdf.

[15] See http://www.cert.org/octave.

[16] See http://csrc.nist.gov/groups/SMA/fisma/framework.html.

[17] "Wet Finger In The Wind."

analogy, saying that we can plan all we want for the visible tip of the iceberg, but we will eventually be sunk by the other 7/8ths of the iceberg that is hidden under the surface.

There are some problems with this argument. First, we now know how to deal with icebergs. Sonar did not originally exist, but it does now, allowing us to better foresee the problems posed. Similarly, our data sources are improving continuously, allowing us to better foresee and estimate threats, vulnerabilities, and impacts. Second, our statistical analyses leverage ranges to better compensate for unknowns. While we need to care about unknowns, our ability to compensate for them allows us to work around any holes or "fuzziness" in our methods. Third, there is no requirement (or need) to exhaustively enumerate all threats in the universe. Instead, using an information-centric approach it is possible to optimize defensibility and recoverability, which is in-and-of-itself a sound strategy.

The bottom line is this: as with all the arguments, we know that we are not dealing with or striving for perfection, and can thus compensate accordingly. In all of these arguments are grains of truth, but none of them is insurmountable. Moreover, given an alternative of "doing nothing" or "working blindly," we ought to happily and enthusiastically adopt approaches with a known margin for error. The data exists, it is valid and useful, and it behooves us as an industry to use it accordingly.

## Cynicism

There is much that could be said about the cynics, but it can be boiled down to two quick points:

- Go read David Mortman's post "Decision Making Not Analysis Paralysis,"[18] in which he quotes Ben Horowitz saying, "…every decision that a CEO makes is based on incomplete information."
- If you are criticizing without contributing, then you are not really helping much.

We have come to a weird place in the evolution of the security industry. After a plateau of more than a couple years, we are now seeing a noticeable backlash against misunderstood areas like risk assessment and analysis. It is time to quit whining about the problem and start helping to solve it. Minimally, it seems recommended that many of us:

- Contribute financially or as a volunteer to the Open Security Foundation[19]
- Drive our respective organizations to opt into data breach reporting using a framework like the VerIS Framework

These two steps alone will help build the data set that is needed to improve risk assessment models and methodologies, while proactively undercutting the "do nothing" data arguments. Whether we understand it or appreciate it, it is time to bury the arguments against risk management and start pitching in to help solve deficiencies (whether real or perceived).

## Path to the future

Information risk management provides us with a viable future. It will, in fact, continue to be core to what we should be doing from an overall assurance perspective. That being said, there are a few clichés that we should keep in mind as we march on:

- Security is a journey, not a destination.
- Perfection is a myth that does not help us evolve the industry. Idealism, on the other hand, is very useful, so long as it is tempered by a touch of realism. Idealism is not the same as perfection.
- There are no silver bullets. Risk is no panacea.
- Risk management is not broken, but rather is evolving and improving over time. Risk assessment and analysis methodologies – especially in the qualitative space – are broken, to a degree, and must be fixed. Evidence-based risk management is helping relieve some these growing pains.
- Recent failings in risk management (e.g., Wall Street and the economy, BP and Deepwater Horizon) are reflective of the need to ensure that the risk vs. reward balance, complete with negative consequences, be allowed to function and flourish. If you remove negative consequences, then there is no reason to manage risk.
- This is not "Lord of the Rings."[20] There is not one single risk measurement to rule them all. Different valid approaches exist, just as there are different data sources with equally valuable, yet distinct, datasets.
- There is a time and place for constructive criticism. Outright, non-contributory cynicism does not qualify as constructive criticism.

If you have found this article to be interesting or useful, or if you have an interest in earnestly contributing to the development and evolution of information risk management, then I highly recommend joining "The Society of Information Risk Analysts" mailing list.[21]

## About the Author

*Ben Tomhave, CISSP, is a Senior Security Analyst with Gemini Security Solutions in the Mid-Atlantic region of the U.S., specializing in solutions architecture, security planning, program development and management, and other strategic security solutions. Ben holds a Master of Science in Information Security Management from The George Washington University. He may be reached at tomhave@secureconsulting.net.*

[18] See http://newschoolsecurity.com/2010/06/decision-making-not-analysis-paralysis.

[19] http://opensecurityfoundation.org.

[20] See http://en.wikipedia.org/wiki/The_Lord_of_the_Rings. Specifically, note the quote "One Ring to rule them all, One Ring to find them, / One Ring to bring them all and in the darkness bind them."

[21] See http://groups.google.com/group/InfoRiskSociety.