

# To Be FAIR About It

A perspective on risk and risk management

## Table of Contents

Comparing Your Security Budget, or, The Lemming Approach to Risk Management	4
Risk Decision-Making: Whose Call Is It?	7
Considerations on Risk Modeling	14
“Vulnerability Events”	16
More Thoughts on Vulnerability	18
Measuring Vulnerability	19
Critical Thinking	24
Communicating About Risk -- Part I	25
Appropriate Funding	27
Communicating About Risk -- Part II	28
Physicians and Medics	31
Compliance is Critical	33
Aggregate Analysis	36
Some Thoughts on “Physics Envy”	40
Lipstick on Pigs	43
Lipstick on Pigs - Part II	45
Models Matter	48
Managing Inconsistency	51
Usefulness?	52
What’s a “risk” anyway?	54
How Much Risk...	55
Getting Loss Right	57
More Than Just Numbers	60
It’s still a choice	62
CVSS Review	63

To be FAIR about it

This document is an collection of my blog posts over the last couple of years ([www.riskanalys.is](http://www.riskanalys.is)). I've included the URL to each post so that you can also view the comments, which is often where some of the most interesting dialog takes place.

The views and concepts expressed are based on a combination of my professional experience and Factor Analysis of Information Risk (FAIR), an analytic framework I've been developing and applying since 2001.

# Comparing Your Security Budget, or, The Lemming Approach to Risk Management

(<http://riskmanagementinsight.com/riskanalysis/?p=221>)

Just the other day I was asked again what percentage of my employer's IT budget went toward security. My answer (as it's always been) was, "Why should I care?" As usual, the reaction I received ran along the lines of, "Well if you don't know, how can you determine whether your organization is spending enough on security?"

In exchanges like this, I'm often asked to explain myself. What self-respecting CISO doesn't benchmark him/herself against their peers? Don't get me wrong, I completely understand the desire to check yourself against your peers, and in some circumstances it's worthwhile. But I don't believe there's much value in budget comparisons for our profession today, and those comparisons may actually work against me as I try to help my employer manage its information-related risk.

## To what benefit?

What practical benefit is there to comparing my spend against the industry? If my numbers are lower than average, am I going to be able to use that to garner more support? Not in my experience. If I haven't effectively made my case already for the various security initiatives on my radar, the simple fact that my employer isn't spending an average amount isn't likely to pull a lot of weight.

On the other hand, if our numbers are about average, then I may very well be at a disadvantage in requesting additional funding for things that really do need attention. Likewise, if our numbers are high, then there's a very good chance I'll need to tighten the belt. Now, if the industry numbers were truly meaningful (more on this in a minute) then positive or negative budgetary adjustments on my part might be appropriate. But the numbers aren't meaningful and so comparison stands a better chance of hindering my ability to be effective than it does of aiding me.

Many of the security people I talk to will argue that their organization doesn't spend enough on security. Consequently, if a significant number of companies are "under spending" (according to their CISO), then setting an industry baseline based on averages derived from "under-spending companies" further erodes the usefulness of the metric.

Does leadership care about how much it's spending on security? Sure it does, but only within the context of whether it's the "right" amount (as differs from "the same amount as everyone else"). More on this later...

## What's included in the numbers?

As I've engaged in surveys and discussions with peers regarding security spend, I've seen a high degree of variability between organizations and what they consider "security spend". The simple fact is that organizational structures vary widely (and tend to change often in many companies) and, as with any comparative metric, if we can't normalize the data then the conclusions and resulting decisions are likely to be flawed.

To be FAIR about it


## Where are we in the curve?

By this I mean the “maturity curve”. In other words, is our security program just starting out, is it well established, or is it somewhere in-between? Keep in mind that the amount and nature of spending on security varies throughout the life-cycle of a security program. Therefore, it isn’t useful to compare organizations that are at different points on the curve. Sure, an argument can be made that by averaging we compensate for these differences, but it still leaves me unable to make a meaningful comparison regarding what my organization spends given its point on the curve.

## It’s not just how much we spend, it’s HOW WELL we spend it

One of my objectives as a CISO is to provide some competitive advantage to my employer by trying to achieve equivalent (or better) risk management at less cost than our competition. Now, I don’t know specifically what the competition is spending (but I do know the supposed “average”!), nor do I necessarily know what they’re spending it on (although I can guess with some degree of confidence because of the focus on “best practices” that seems common -- more on this in a minute). But I do know that if my target is simply to spend the same amount as everyone else, then I’m not focused on the right thing and I’m not being a responsible steward of my budget.

## Lemmings



If I use as a target the “average” security spend in the industry, then I am, by implication, assuming that the average company is doing a good job in how it manages information risk.

This is really a topic for another blog post, so I won't dive deeply here. Briefly, I believe our industry is still far too dependent on the shamanistic principles of:

- **FUD** -- scare the non-believers into following our advice. "The thunder-gods will get you." "The hackers will get you." -- not much difference there.
- **Best practices** -- "The tribe down the river does it this way, grandpa did it this way, so we have to do it this way." Some best practices are badly dated, others reek of vendor agenda, so there's no guarantee that best practice is the right solution for our particular risk issues and corporate risk tolerance. Perhaps worse, blind adherence to best practices violates our responsibility as stewards of our budget to look for cost-effective solutions.
- **Gut instinct of the practitioner** -- Don't get me wrong, many security practitioners have developed outstanding instincts. Furthermore, good instincts are a critical component of dealing effectively with almost any aspect of life. The problem is that without applying a dose of critical thinking and analysis to the complex problems we face, we're -- a) too vulnerable to personal bias, industry myth, and dogma, and b) unable to effectively defend our conclusions and recommendations to our stakeholders.

## Tolerance

Perhaps the most significant concern I have about budget benchmarking is that it implies there's some universally accepted "appropriate amount" of spend. Hogwash. Think about it this way -- how much automobile insurance do you carry? Is it above or below average? Would you change it if you knew it was above or below average? Some people might, but I select my coverage based on my income, savings, expenses, risk, and risk tolerance. This coverage is the "right" amount for me given these variables.

The fact is, every organization has different resources, expenses, risk levels, and risk tolerances from every other, and it's a fallacy to believe one-size-fits-all. The good news is that our organization's leaders know what their resources and overall expenses are, and they have an innate sense of what their risk tolerance is (because they're making risk decisions every day). I believe the challenge has been that we haven't been doing a great job of providing leadership with useful risk information. Until we can do that, the question of how much we're spending on security seems almost moot.

## If not benchmarking, then what?

The bottom line is that the "right amount" of security spend is unique to each organization. Furthermore, executive management's opinion is the only one that ought to matter regarding what that amount is. They are the ones who have a clear understanding of the company's condition, objectives, resources, competing risk issues, and risk tolerance. It's their job to manage the overall business risk portfolio. Our job is to help them make well-informed decisions regarding our piece of that puzzle by providing a clear, unbiased, and useful picture of their information-related risk and risk mitigation options. Until/unless we do that, then any argument regarding appropriate security spend isn't terribly useful.


## Risk Decision-Making: Whose Call Is It?

(<http://riskmanagementinsight.com/riskanalysis/?p=228>)

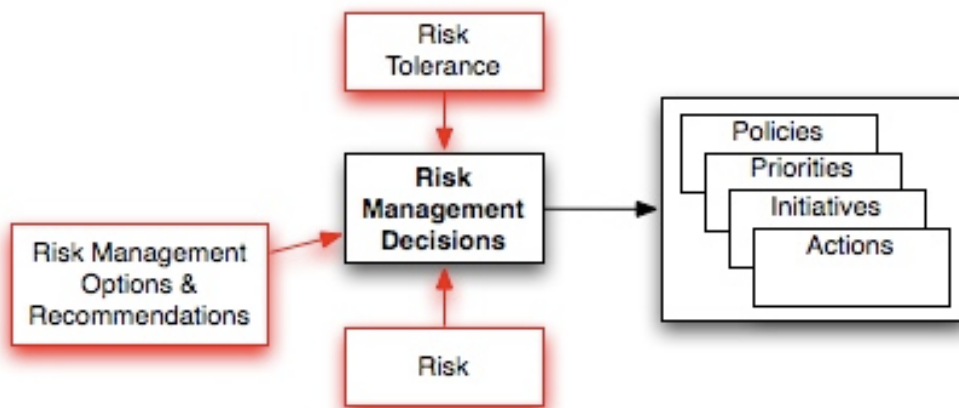
I still occasionally run into a debate with colleagues over whether security should be making the major information risk decisions for an organization, or whether it's business management's responsibility. Rather than just spew my opinion, let me try to build an illustration of how I view the problem.

Picture this...

1. Risk decisions are the things that drive policies, priorities, initiatives, and actions (this falls under the category of "duh").

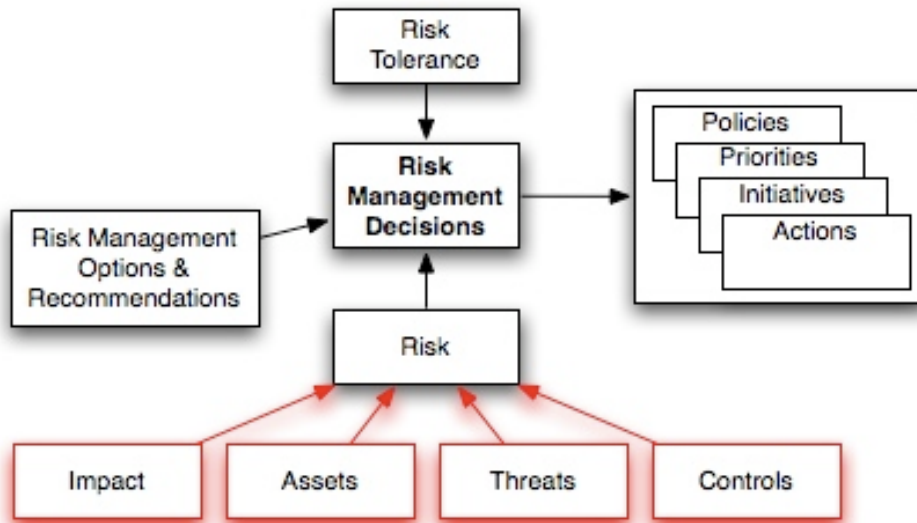


2. Well-informed risk decisions are dependent upon knowing the risk associated with the decisions, as well as the best risk management options. Risk tolerance also is an inevitable factor (we'll discuss the question of whose risk tolerance further on).

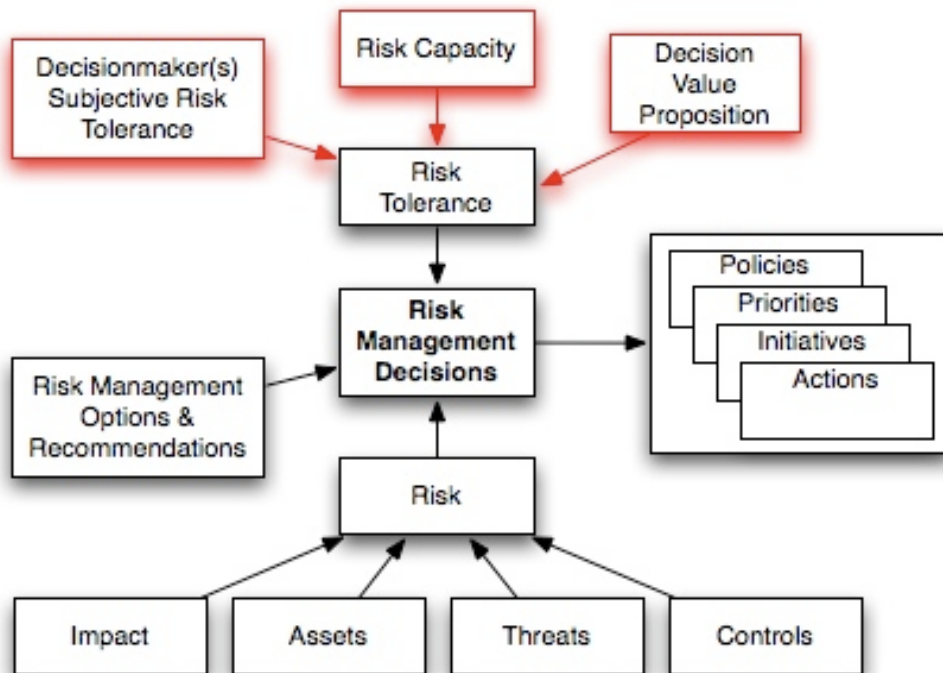


3. Understanding risk, of course, requires that we understand the factors that drive impact (stakeholders, laws, contracts, competitive landscape, etc.), the assets associated with impact, threats against those assets, and controls that are in place to manage risk. Absent any of these inputs, our understanding of risk can be seriously deficient and the resulting decisions flawed.

To be FAIR about it



4. So far – no surprises. At this point, however, things begin to get a bit more interesting... Specifically, risk tolerance is derived from three inputs; risk capacity, the decision’s value proposition (the potential upside associated with the risk scenario), and the decision-maker’s subjective risk tolerance (more on this further on).

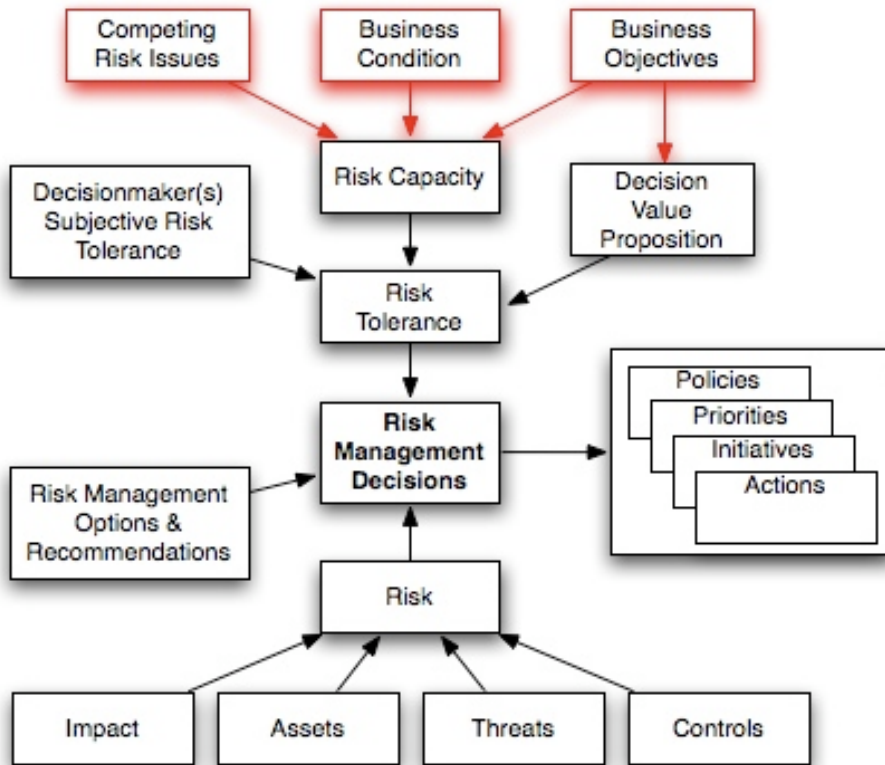


5. Risk capacity also has three inputs; the organization’s current condition relative to its objectives, as well as the portfolio of competing risk issues. It’s important to recognize, too, that these factors will often vary across the different types of loss (e.g., productivity, competitive advantage, resources, reputation, etc.). For example, an organization that has a significant stockpile of resources will have more capacity for resource loss than will an organization that operates on a shoestring. Likewise, an organization that is

To be FAIR about it

trying to build market share will have less capacity for reputation damage than will one that already leads the competition and/or that has a very loyal customer base.

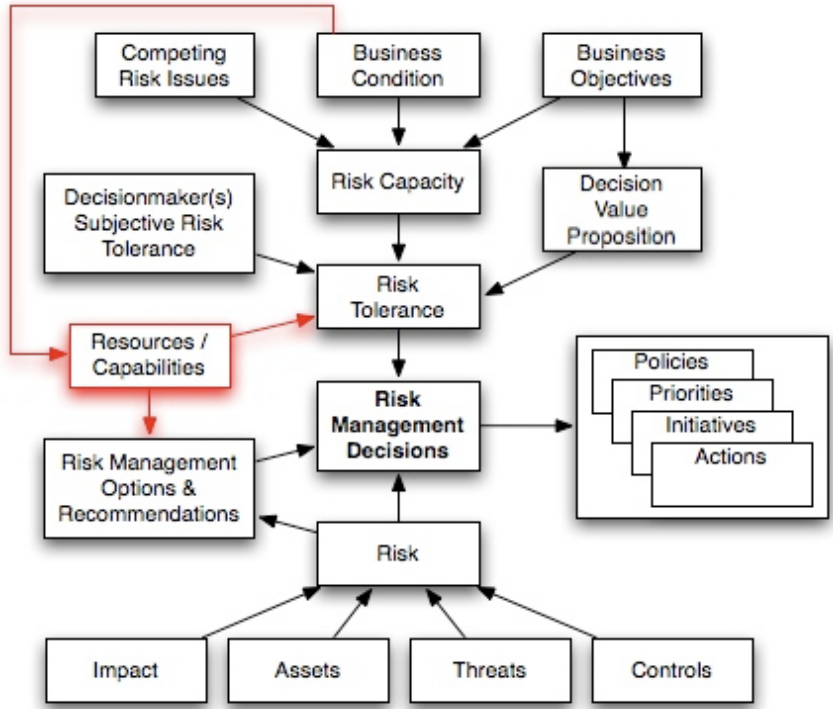
The point is, tolerances will vary not only between organizations but also between types of loss within an organization.



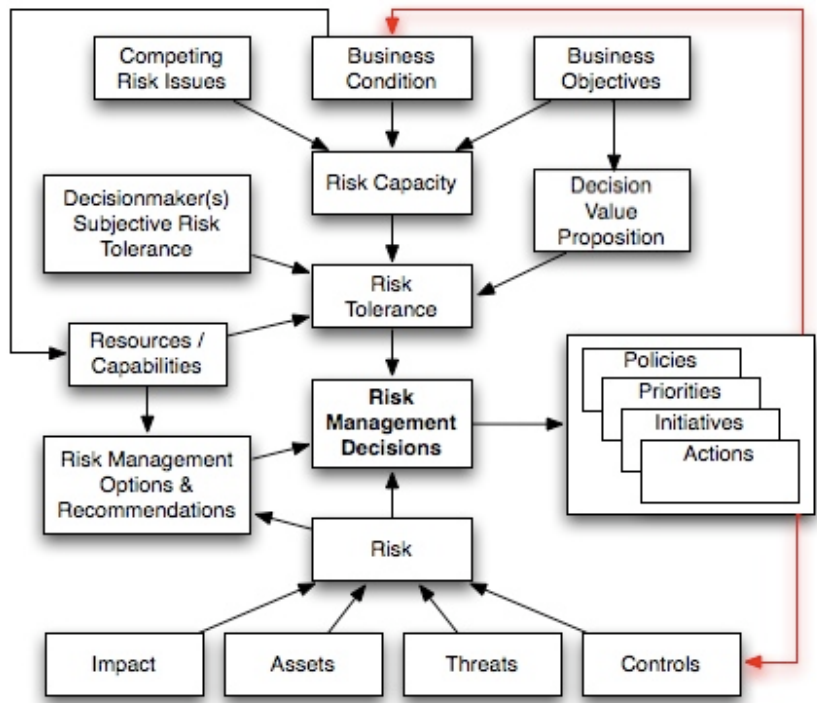
With regard to competing risk issues, it's important to keep in mind that information-related risk is only one of many risk domains management has to deal with (e.g., market, insurance, investment, etc.). Combine this with complex organizational conditions and objectives, as well as limited resources, and it becomes clear how important (and difficult) it is to strike the right balance in applying risk management resources.

6. Speaking of resources...available resources and capabilities help to drive which risk management options are feasible. These resources, of course, are dependent on the organization's condition. Note, too, that resources and capabilities can affect risk tolerance, as an organization with fewer resources for mitigating risk may be forced to accept more risk if, for example, a decision's value proposition is particularly compelling.

To be FAIR about it



7. And finally, the policies, priorities, initiatives, and actions that result from risk decisions will have an effect on risk and the organization's condition (for good or ill). At the very least, expenditures made to manage information risk are no longer available to use on competing risk issues and opportunities.



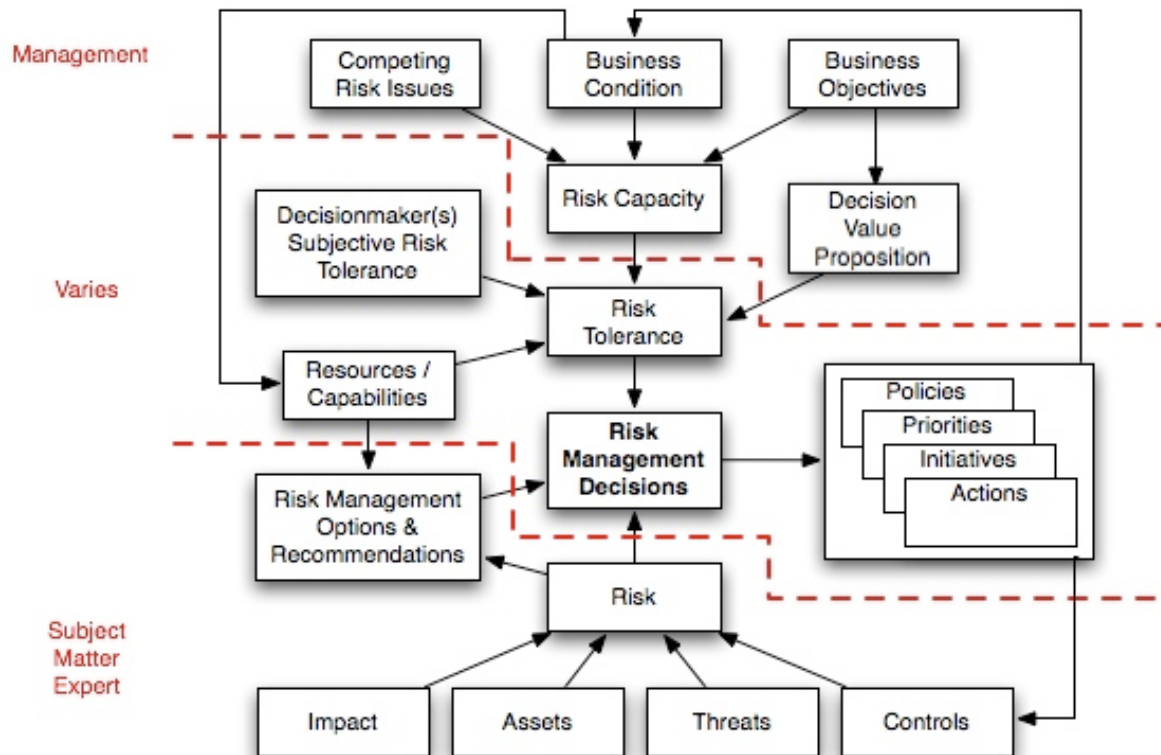
To be FAIR about it

Okay, if by now you haven't fallen asleep or decided to spend your time elsewhere, I'll tie all this back to the original question of who should be making the decisions regarding information risk...

## Carving it up

Using this illustration of the risk decision elements we can draw lines that carve the landscape into three parts –

- Those elements that would appear to belong to business management,
- Those elements that would appear to belong to the subject matter experts (in this case, us), and
- Those elements in the middle that, well, could go either way



Note that the decision itself falls into the “could go either way” domain, which means I can't give you a definitive, “This is how it should be” answer. What isn't surprising is that who makes the risk decisions will vary from organization to organization. What's unfortunate is that in many companies security leadership believes they are (or should be) empowered to make the major decisions while business leadership believes otherwise. Speaking from painful personal experience, this disconnect can cause significant trouble.

## Size matters

Of course what I mean is that the size (significance) of the risk decision also determines who can/should/will make the decision. Business management isn't usually going to be involved in day-to-day operational risk decisions. Furthermore, security management can't personally be involved in each discreet risk decision that takes place throughout the organization (e.g., Clerk: “Hmmm. Should I shred

To be FAIR about it

this document, or just chuck it in the trash?”). These day-to-day and discreet risk decisions are where good policies, procedures, and risk awareness education come in.

At the end of the day, decision significance is a continuum rather than a binary or clearly differentiated scale. Consequently, some decisions fall into a grey area regarding who should make what call. For these issues, the question of who should make the decision will vary from organization to organization. You can, however, work with management to come up with some ground rules, for example; policies, policy exceptions, strategic initiatives, and significant expenditures fall into business management’s court, and security deals with the rest.

## Look again

With regard to discreet risk decisions, take a close look at the risk decision diagram. You’ll see that the diagram applies quite well whether we’re talking about major strategic decisions or the discreet risk decisions being made by employees countless times each day. The only difference is that, in the absence of a clear understanding of organizational risk tolerance, employees WILL substitute their own views of organizational risk tolerance (or leave it out of the equation altogether). In any event, employees often will be placed in the unfortunate position of having to reconcile organizational risk tolerance with their own conditions/objectives/competing risk issues, etc. (e.g., the question of choosing compliance with security policy over meeting the deadline their bonus is resting on...). This highlights the need to be aware of, and manage, issues related to competing individual and organizational priorities.

Something else to think about is that policies and processes will never cover all of the potential risk decisions our employees face. As a result, it’s critical that education and awareness efforts go beyond regurgitating policy, and include information that helps employees understand risk and the organization’s risk tolerance so that they can make good judgment calls. This better understanding also helps them tolerate those policies they otherwise chafe at.

## Things to consider

The simple fact is, security leadership will never know as much about the business-related elements at the top of the illustration, and business management will never know as much about the risk elements at the bottom. Consequently, if security is empowered to make the major decisions, then they need to spend the time and effort to learn as much as they can about the business-related elements. On the other hand, if business leadership is making the major risk decisions, then security must provide clear, unbiased, and useful information so that the decisions are well informed.

(For those who are curious, I strongly prefer that business management make the major risk decisions where I work. I’m far more comfortable in my ability to provide them with good risk information and mitigation options than I am in my ability to sufficiently learn and understand the complex business landscape. Besides, when they’re the ones who have made the decisions, pushback and arguments are largely eliminated. I’ve also found that you have far more influence as a trusted advisor than as a combatant.)

A decision-maker will to some degree ALWAYS apply his or her own personal risk tolerance to a decision. Consequently, if security leadership has been empowered to make major risk decisions, they should try very hard to be as aware as possible of business management’s risk tolerances. If security leadership isn’t careful on this, then they will, invariably, run into issues where business management doesn’t support security’s decisions. And if the misalignment is bad enough (and I’ve both witnessed this and come close to having it happen to me – long ago) then it can become a “terminal” condition. At the very least it makes the waters far choppier than necessary.

To be FAIR about it

I make it a point to review the risk decision question (and now the diagram) with business management whenever I take a new job or have a new business leader join the organization I work for, even if I'm pretty confident about where they stand. When I've had these conversations it's always generated very productive dialog and has strengthened the relationship.

Note: This posting will soon be reproduced as a white paper and/or PowerPoint on the RMI website.

## Considerations on Risk Modeling

(<http://riskmanagementinsight.com/riskanalysis/?p=315>)

As Alex discussed a couple of weeks ago, Mike Rothman posted an [article](#) discussing concerns he has with risk management models. In his article, Mike reminds us that risk management is not a silver bullet, that we should only do as much risk modeling as is necessary in order to achieve our goals (I assume he means the organization's goals), and that calculating risk to the Nth degree doesn't keep attackers at bay. I couldn't agree more.

That said... there are a number of things Mike has in his article that I would like to offer a different perspective on. Specifically:

In his third paragraph, Mike states that "*we (the industry) are all about mitigating risk.*" Actually, it might be more accurate to state that we, the industry, should be all about *managing risk*. The difference is that managing risk (or managing anything) is about achieving a desired result. In our case, it's about managing how often bad things happen and how bad they are when they do happen, to a degree that's acceptable to management. Mitigation, by definition, assumes reduction and seems to ignore the notion of the balance point where management is comfortable with their risk position and/or maybe even decides there's a need to increase risk in order to take advantage of opportunities.

Mike also implies that it takes "considerable resources" to build a risk model. Certainly it CAN take considerable resources, but that doesn't mean it HAS to. To Mike's earlier point about not calculating risk to the Nth degree, pragmatic risk analysts seek to find the right level of abstraction/complexity in their analyses to fit the need and available resources.

I'm a bit unclear on Mike's analogy to building a hotel in a swamp... He is correct that all risk-modeling requires assumptions and estimates. Yup. No question of that. So instead he's suggesting that we should, ummm, NOT model our risk scenarios and instead just shoot from the hip? That seems to necessarily entail even greater assumptions and fuzzier estimates that are at least as vulnerable to error (and arguably even more vulnerable). This seems even swamplier, so to speak, than risk modeling.

Mike is right about the lack of precision in risk quantification. In fact, if your estimates of future events are dead-on, you were lucky. This is true in any analysis of future events. Mike states that actuaries "know" we'll have 3.7 car accidents in our lifetime. Of course, they don't really KNOW this because they can't predict the future, even with their terabytes of empirical data. If we, as individuals, actually experience 3.7 car accidents in the course of our life (assuming we could have a .7 accident), it's coincidence. Almost everyone will experience more or fewer than 3.7 accidents. What he really means is that, on average, the population will have somewhere in the neighborhood of 3.7 car accidents. There is no means of predicting precisely what the experience of a single individual will be.

The point is, precision in risk analysis is a pipe dream. But it isn't precision we should be seeking, it's accuracy. (This was pointed out to me by two senior vice presidents of the actuarial departments in a large insurance company when I first started working on FAIR.) You can be precise and yet terribly inaccurate -- e.g., an estimate that John Doe will make \$4,325,211 dollars in 2007 would have been a precise figure but inaccurate (he only made roughly \$500k). If, on the other hand, we had estimated that John would make between \$400k and \$550k, our estimate would not have been precise but it would have been accurate. The degree of required precision for any estimate is dependent upon the situation. I don't know any business people who expect precision in our risk estimates. They expect ballpark accuracy.

To be FAIR about it

I'd also differ with Mike's assertion that "*it's all about measuring the relative risk*" and that "*What security managers should strive to do is get a relative idea of the risk to each major business system...*"

Yes, relative risk is important because you want to focus on the higher risk issues, but this disregards the question of whether the amount of current risk is acceptable. Focusing too strongly on relative risk seems to take us down the old rabbit trail of continually winnowing down the risk issues until they just disappear altogether -- which, of course, is also a pipe dream.

Mike's statement that "*any security person worth his or her salt should already know what major systems are most important and what the potential risks to those systems are*" is true. That's not the point though, or at least it isn't enough. We can know (or think we know) what the important systems are, and what the relevant threats and vulnerabilities are, but what matters is whether the combination of threats, vulnerabilities, and value/liability at risk is acceptable and, if not, how far off it is (i.e., the degree of unacceptability). And at the end of the day, acceptability is not the security professional's call, it's the organization's leaders' call.

We, as security professionals, are still on the hook to get our jobs done. Mike's absolutely right about that. This, however, begs the question of what our job is. If we think it's to "secure things", then I believe we're on the wrong track. Securing things is simply a means to a more important end. Loss happens. It will always happen. And because loss happens, we have to recognize that, ultimately, we're being paid to help ensure that the losses our employers experience are at an acceptable level. This is called risk management. And risk analysis/modeling -- done well -- provides the means to manage risk more effectively on a more consistent basis.

## “Vulnerability Events”

(<http://riskmanagementinsight.com/riskanalysis/?p=241>)

When a new vulnerability is discovered in (for example) an operating system, does that mean the system was vulnerable all along? As I see it, the answer is "No".

The rationale behind this answer is based on the fact that weakness (a.k.a. vulnerability) is a relative term. Logically, a relative term requires at least two components – one relative to another. Oh, it's true that the "flawed" condition within the operating system existed all along, but in order for that condition to actually BE vulnerable, the capability to exploit the condition had to exist. And within the context of a human threat community, capability requires two things: knowledge and resources. Consequently, until the condition was known to be exploitable, it couldn't be leveraged and wasn't a vulnerability.

So, if a vulnerable condition occurs when available force becomes greater than the ability to resist that force, then vulnerability can come about in one or more of three ways:

- Resistance strength is diminished in some manner (e.g., cutting part-way through a rope)
- Available force increases so that it exceeds existing levels of resistance (e.g., more weight is added to the end of the rope)
- An asset is newly exposed to threat elements, either because the threat elements are new to its landscape or it enters a threat landscape it didn't exist in before (more on this in a second)

Regardless of the cause, whenever available force becomes greater than the ability to resist, you have what can be referred to as a “vulnerability event” – i.e., vulnerability now exists where it didn't before.

In our operating system scenario, nothing changed about the operating system itself. What changed was threat capability, which increased as soon as the threat community became aware of the condition's exploitability. At that instant, the knowledge component of the threat community's capability changed, and their resources likely changed soon after, when exploit code was developed.

### Vulnerability, not loss

Here's another example prompted by an excellent question posed by Stacy on the ["layer8.itsecuritygeek blog"](#) -- essentially, how should we classify "near miss" events where, for example, someone sends sensitive information unencrypted over the Internet? Is that a “loss event”? By my reckoning, the answer is no – unless and until actual loss to the organization materializes. Instead, it's another example of a vulnerability event – i.e., vulnerability to loss now exists where it didn't before (ref. #3 above).

### Why “vulnerability events” matter

If history provides any clues to the future, some folks are going to question why I feel the need to define yet another term. It's a fair question (pun intended).

If you're familiar with FAIR you already know that we define two other event types – Threat Events and Loss Events. Threat events occur when a threat agent acts against an asset. Loss events occur when loss

To be FAIR about it

results from a Threat Event (i.e., as happens when force exceeds resistance). The reason it's important that we make distinctions between event types is three-fold:

- It helps us to better understand our problem space, which is always a good thing,
- It allows us to communicate more consistently and effectively, and
- It enables us to identify and make meaningful use of metrics

This last point is especially important as we try to make better use of metrics.

## More Thoughts on Vulnerability

(<http://riskmanagementinsight.com/riskanalysis/?p=347>)

Take a look at the following list and ask yourself which of the following would be labeled “vulnerable”:

- An eight-character password made up of alpha and numeric characters
- A six-character password made up solely of alphabetic characters
- A four-character PIN made up solely of numbers
- A fourteen-character password made up of alpha, numeric, and special characters

Actually, there are a couple of rational answers -- 1) “it depends”, and 2) “all of them, to some degree”. As I think about it, maybe these are both the same answer stated from slightly different perspectives.

### It Depends

The “*it depends*” answer comes from the fact that we haven’t identified the threat agent we’re up against. If we’re talking about a threat agent who isn’t particularly skilled, isn’t leveraging a toolset that makes up for their lack of skill, and/or doesn’t have much time in which to carry out their attack, then even the four-character numeric PIN might be more than they’re capable of defeating. On the other hand, if the threat agent is highly skilled, has powerful tools, and has lots of time, then even the fourteen-character password can be defeated. This, it seems, also supports the “*all of them*” answer. The point is, everything is potentially vulnerable under the right (or wrong) circumstances.

Unfortunately, we tend to use the term vulnerability as if it’s a binary condition. Something is vulnerable or it’s not. But whether we realize it or not, what we’re really doing when we say that something is or isn’t vulnerable, is making unstated assumptions and generalizations about threat capability relative to the control in question.

Of course, some folks insist that we have to rate controls against the “most capable” threat agent. A couple of problems with that include:

- Who’s to say what the most capable threat agent is capable of?
- If we’re judging against the most capable threat agent, then everything is theoretically vulnerable (given enough skill, resources, and motivation)

The fact is, when someone calls something vulnerable (or not vulnerable) they’re consciously or subconsciously quantifying the threat capability as well as the control condition, comparing the two, and then making a judgment about the degree of vulnerability. Or, I suppose, they may just be blindly following someone else’s proclamation that “this is vulnerable” and “that isn’t”.

So, if we’re performing subconscious quantification and comparison when we rate the vulnerability of something, is there any reason we can’t/shouldn’t be more [conscious about it](#)? What’s the downside? And is there any reason to believe conscious analysis would be less accurate than the subconscious one? Think about it. Subconscious assessment is at least as exposed (and arguably much more exposed) to errors of omission, errors in estimation, and personal bias/gaming, which means conscious analysis can be no worse and has the opportunity to be much better.

# Measuring Vulnerability

(<http://riskmanagementinsight.com/riskanalysis/?p=348>)

Apologies in advance, for the length of this post...

## In a perfect world...

...we'd know which specific threat agent was going to act against us and know the capability of that threat agent in absolute terms (e.g., pounds per square inch), as well as know (through testing) what our resistance capabilities are in those same absolute terms. If we had this information AND assuming this information was precisely correct all of the time, vulnerability becomes a clear and simple binary consideration -- we will be or we won't be.

## Stating the obvious (anyway)

Losses occur when threat events take place that we're vulnerable to. This is true whether we're talking about weather events, human error, or malicious acts. Obviously, we don't experience loss with every threat event, which means we're only vulnerable sometimes -- i.e., less than 100% of the time. This means there is some probability associated with whether we'll be vulnerable to any given threat event. The process of measuring vulnerability is intended to help us understand what that probability is likely to be.

## Simplest approach

Perhaps the simplest approach is to identify the threat community you're analyzing risk against and simply estimate your ability to resist the capabilities of that threat community. For example, we might estimate that our web application is capable of resisting all but the top 2% of the cyber-criminal threat community -- i.e., two out of a hundred hackers have the skill and resources to defeat the application's security.

This works as a quick-and-dirty solution, and in many cases is good enough. Read on if you're interested in a somewhat more involved approach.

## Uncertainty

Unfortunately, in the real world we usually don't know:

- Which threat agent is going to act next,
- What their capabilities are, or
- What our resistance capability is going to be

Making matters even more challenging:

- We don't have an absolute measurement scale for some threat categories (e.g., human capability)
- Our measurements are imprecise (e.g., we can't measure force or resistance perfectly)
- One or more of the values being measured may vary over time (e.g., hurricane wind speed varies throughout the lifetime of the storm, and strength can change throughout the lifetime of a control )


To be FAIR about it

- One or more of the values being measured may vary across a population (e.g., not all hurricanes have the same wind speed)

### When absolute scales apply

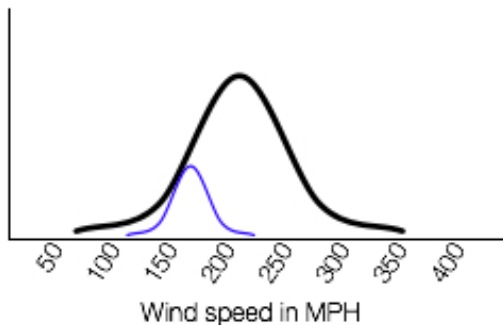
(Warning: This is an illustration and not an engineering exercise, for those who might want to argue details.)

Some types of threat categories can be measured using absolute scales (e.g., wind speed in miles per hour), which makes things a bit more straightforward. For example, thru testing we could estimate that a structure should be capable of resisting wind forces between 150 and 200 MPH.



By using a distribution to describe this measurement, we account for the fact that under some circumstances wind speeds of less than 150 MPH might compromise the structure, while in some circumstances the structure may be able to withstand speeds greater than 200 MPH.

If we wanted to measure the structure's vulnerability to a specific type of storm (e.g., a tornado) we could plot a similar distribution for tornado wind speeds (black curve below). This distribution reflects the fact that wind speeds vary from tornado to tornado, ranging from under 100 MPH to over 300 MPH, with most falling in the 200 MPH range. (Keep in mind this is just an illustration and isn't intended to reflect actual tornado data.)




In order to determine the probability of being vulnerable, we'd use a Monte Carlo function to:

- Take a random value from the tornado distribution and from the structural resistance distribution
- Compare the values -- i.e., for this iteration, determine whether wind speed was greater than resistance
- If wind speed was greater, increment a counter that tracks the number of vulnerable instances
- Repeat a thousand iterations (or ten thousand, a million, etc.),
- After completing all of the iterations, the vulnerability counter divided by the number of iterations provides the probability of this structure being vulnerable to tornado winds

### When an absolute scale doesn't exist (the human threat community)


Human threat capability can be boiled down to skills and resources. Because skills and resources vary from individual to individual, we can characterize threat community capability as a distribution. At one end of the distribution are those threat agents who have the least capability, while at the other end are those who are the most capable. As seems to be the case for most things in nature (e.g., weather events), the distribution is probably pretty close to being bell-shaped (i.e., the majority of threat agents fall somewhere below those who are most capable and above those who are least capable).



A “100% secure” control (if such a thing existed) could be illustrated as existing outside of the threat community capability distribution. It would be 0% vulnerable.



More realistically, we can in most cases expect that some portion of the threat population would have the skill and resources to compromise a control (shown below).




Now, because of the uncertainties regarding threat capabilities and control strength, it would be more accurate to describe control strength as a distribution as well. For example, we expect the control is at least resistant to 90% of the general threat population, and may be resistant to as much as 99%+ of the population.

To be FAIR about it



This is fine as far as it goes, but it doesn't get us the answer we're looking for in most circumstances. Most of the time it isn't enough to know our vulnerability to the general threat population. In most analyses, we want to know what our vulnerability is to a particular threat community (e.g., cyber criminals, nation-state intel units, etc.). In that case, we'd have to plot the capability of the threat community in question (red distribution).




With that plotted, we can run our Monte Carlo function again, generating a probable vulnerability by taking random samples from the control distribution and the distribution of the specific threat community in question.

The key to measuring vulnerability in the absence of an absolute scale is to use the general threat population capability as the comparative baseline for both control strength and the capability of the threat community in question.

### Considerations

Of course, because some malicious threat communities tend to share knowledge and tools, there can be an equalizing effect, which potentially narrows the width of the threat capability curve (shown below) but likely wouldn't change its fundamental bell-shape. The good news is that this narrowing effect wouldn't alter how we measure. The bad news is that it does affect vulnerability, which we know intuitively anyway.



Another consideration is the fact that the capability of the malicious population evolves over time -- i.e., the curve shifts to the right along the continuum. For example, at one time in the past DES was

To be FAIR about it

considered invulnerable to brute force cracking. It isn't any longer. In other words, we could say that the control stayed in place along the continuum, but the capability curve shifted to the right. This highlights the fact that it's important to keep abreast of how threat capability evolves, so that you can evolve your defenses as well. Also, this is good fodder for the importance of defense-in-depth.

### Concerns

An obvious concern is the inexact nature of these estimates and the potential for the analyst to estimate badly for various reasons. We've covered this issue previously in other postings, so I won't go into it in depth now. Suffice it to say that yes, this is an inexact measurement fraught with all of the goblins that any measurement approach is subject to. That said, keep in mind a few things:

- The ability to estimate effectively can be significantly improved using [calibration techniques](#)
- There's no such thing as a perfectly exact measurement, whether you're using a laser or the width of your thumb to do the measuring. Therefore, the purpose of measurement is to reduce uncertainty, not eliminate it
- You can apply confidence levels to your estimates, both to describe the probability of actual values being outside of the estimated minimum and maximum, and to shape the peakedness/flatness of the curve

### Monte Carlo analysis is designed to help account for the uncertainty in measurements

You should never convey to management that these numbers are exact. In my experience management won't have any problem with this, as the numbers they're given from other business disciplines have precision challenges of their own.

Bottom line --- If you're trying to quantify risk, then you have to quantify vulnerability. This is one logical means of doing so. What's more, it seems to accurately reflect how we subconsciously evaluate and quantify vulnerability anyway, only it brings the analysis to the surface. And by bringing it to the surface, it allows us to better understand and analyze risk scenarios.

If there's interest, I can provide a couple of examples in a future post. Also, if there's interest, I can include an example where the threat event is due to error rather than malicious intent.

## Critical Thinking

(<http://riskmanagementinsight.com/riskanalysis/?p=349>)

Another perspective on risk management that I've found useful is to recognize that risk issues are "open-ended" in nature rather than "well-structured". Well-structured problems can be reasoned to a single correct answer – e.g.,  $3+3=6$ , or "Will I overdraw my bank account if I write this check?" Open-ended problems, on the other hand, are those that can't be reasoned to a single, undisputed correct answer. Examples of open-ended problems include:

- What's the right solution for peace in the Middle East?
- What's the best financial investment or insurance plan?
- Should I step on the accelerator or the brake at this yellow traffic signal?

Most of the information security/risk problems we face are open-ended – in other words, there are very few clear, undisputed correct answers. Examples of open-ended questions we're forced to deal with include:

- What's the best solution for this risk issue?
- Is this amount of risk acceptable?
- What is the highest priority of our many security issues?

Because these issues defy simple, indisputable answers, and because each of our circumstances will vary, we're forced to apply [critical thinking skills](#). This, of course, flies in the face of prescriptive standards and "best practices" that try to portray the risk landscape as black and white (well-structured) when it's clearly shades of grey (open-ended). To be fair, non-prescriptive standards and "best practices" play an important role as directional references -- compasses so-to-speak. But even a really good compass can't always account for the unique circumstances we encounter.

As I see it, any grade school graduate can recite a standard or compare a checklist against what they see in front of them. Whether we realize it or not and whether we like it or not, we have to prioritize, make decisions, and defend/explain our rationale within a complex open-ended environment. Sometimes a specific best practice or standard will be the most cost-effective solution for a given circumstance; sometimes it won't. The important thing is being able to recognize the difference. That's where critical thinking comes in, and that's where we provide real value as professionals.

An interesting one-page matrix (in a Word document) that categorizes thought process maturity can be found [here](#). It's worth a read.

## Communicating About Risk -- Part I

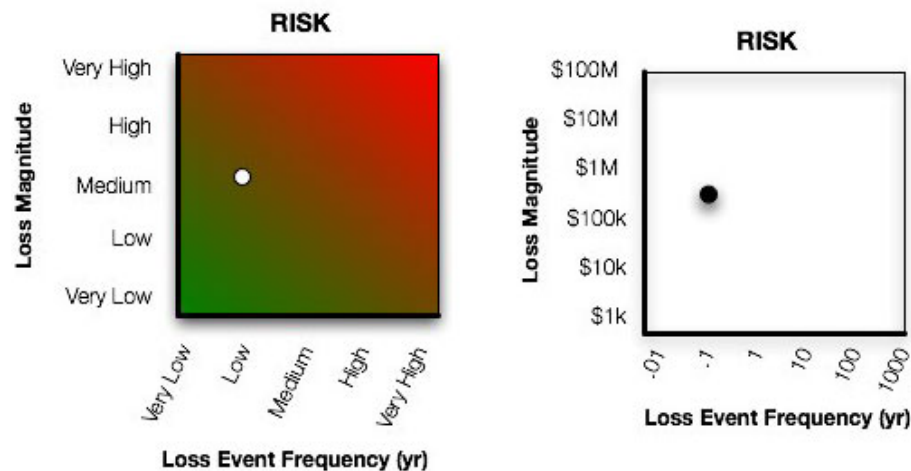
(<http://riskmanagementinsight.com/riskanalysis/?p=351>)

In his comments a couple of weeks ago, Walter brought up an important point. Paraphrased, he pointed out that misrepresenting the precision of an analysis is a bad thing. He also pointed out that this isn't so much a problem with the analysis model (although it's more likely to occur with a quantitative model), but rather tends to be a problem with how an analyst communicates results to management.

With that in mind, I thought I'd write a couple of posts about communicating risk. In this week's post, I'll talk about "risk qualifiers" that can be critical in helping management understand the true nature of some risk scenarios.

"I can live with this..."

Let's say that you've done an analysis and the results look something like what's shown in the charts below (I've included both a qualitative and a quantitative version):



At first glance, a decision maker might think "This doesn't look so bad. I can live with this level of risk." But that's not necessarily the whole story...

### Unstable conditions

An unstable risk condition exists when the following characteristics co-exist:

- Threat event frequency is low
- Vulnerability is high
- Probable loss magnitude is significant

When these conditions exist, the low loss event frequency is driven solely by the low threat event frequency. In other words, we're not actively managing loss event frequency; we're just trusting to luck. If threat event frequency changes (or an event occurs at all), then significant impact will likely occur. An example might be an internal application that handles a significant volume of sensitive consumer records, but that has little or no authentication or authorization control in place.

To be FAIR about it

Now, if all we provided management was a qualitative “Medium/Low” risk statement or a quantitative statement that “probable loss event frequency is roughly once every ten years with a probable loss magnitude of \$500k”, then we haven’t really allowed management to make an informed decision.

This additional information about the unstable nature of the risk condition is critical for a couple of reasons: 1) it allows management to decide whether they want to gamble, and 2) instability can reflect poorly from a due diligence perspective.

## Fragile conditions

A fragile condition exists when the following characteristics co-exist:

- Threat event frequency is high
- Vulnerability is low, but dependent on a single effective control
- Probable loss magnitude is significant

At a glance, this will look similar to an unstable condition. In this case however, a single control is all that prevents a high loss event frequency. An example might be a single layer Internet architecture, where the volume of threat events is high but the firewall is generally quite effective.

## Differentiation

One big advantage these qualifiers provide is to be able to differentiate between risk conditions that, from a risk chart perspective, look the same. This differentiation allows us to prioritize better, which leads to more cost-effective risk management.

Another advantage is that it provides nomenclature for expressing what our intuition has probably already recognized. In other words, the experienced information security professional would intuitively recognize the difference between an unstable or fragile condition and one that isn’t (but that may look the same on a chart). In my experience, what we tend to do in those instances is label the condition “high risk”. The problem with this is that it lumps these scenarios in with those where loss event frequency and loss magnitude are high, which erodes management’s ability to prioritize effectively.

At the end of the day, effectively managing any complex set of issues requires an ability to differentiate. These qualifiers have proven to be extremely useful in that regard.

## Appropriate Funding

(<http://riskmanagementinsight.com/riskanalysis/?p=352>)

Because many organizations are beginning to wrestle the funding beast at this time of year, I thought I'd focus this week's post on the question of "appropriate funding". It only tangentially touches on the question of communicating about risk, but I'll return to part two of that series next week.

One of the arguments I've heard folks use to dismiss the notion of a risk-based approach to security is that it's been tried and failed. The argument goes on to claim that it isn't possible to get appropriate funding for security because management just doesn't "get it". And, while I agree that many (most?) past attempts at risk-based security have struggled, I'd submit that it was because the methods used didn't address risk effectively. They often focused solely on worst-case outcomes (which is the Chicken Little problem), didn't apply any rigor in determining risk, simply focused on vulnerability (but called it "risk"), or treated the problem as a possibility issue versus a probability issue.

Of course, the argument about funding begs the question of what constitutes "appropriate funding". It's naive (or arrogant) to believe that I -- as an information security professional -- am in a position to understand the incredible mix of business issues that determine the right risk-balance for an organization. Running a business requires weighing the various risk-domains management faces (investment, insurance, product, market, security, etc.) as well as complex value propositions in light of the organization's objectives and limited resources. And, while it's imperative that information security professionals seek to understand the business side of the equation, we are never going to have the same breadth and depth of vision into the organization's unique mix of business issues that executive management has. Combine that with the fact that it isn't our risk tolerance that matters, and it should be crystal clear that complaints of being "underfunded" have to be cast in the light of "Compared to what?". Compared to what **we** think it ought to be? Compared to some industry baseline of [questionable applicability to our organization?](#)

Of course, I struggled to get management support for years. I tried leveraging fear, uncertainty, and doubt. I also tried the old "You have to do it because it's best practice" card. And although both of these can work for awhile, at the end of the day, management's perspective will likely be that you're paranoid and you lack perspective about the nature of running a business. I've come to the conclusion that if I believe I'm underfunded, then it's likely:

- I haven't done a good job of communicating risk to the business,
- I don't sufficiently understand the risk tolerance of the organization's leadership, and/or
- I don't understand the mix of competing risk issues, resource limitations, or business objectives.

It's my responsibility to see that I'm not underfunded by providing high quality (unbiased) risk information to management. If I do that, then I can expect to receive an appropriate level of funding given the other business considerations management faces and their risk tolerance. The funding may be less than I'd like given my risk tolerance, but that's a personal problem.


Frankly, since taking a risk-based approach to my job, I've had very little difficulty getting management support for the stuff that matters most.

## Communicating About Risk -- Part II

(<http://riskmanagementinsight.com/riskanalysis/?p=354>)

### The trouble with likelihood

It's common to see charts similar to the one below used to communicate risk. On one axis we have Impact, and on the other we have Likelihood. We'll save a discussion regarding Impact for another post, but in this post I'd like to point out a couple of subtle but important limitations with the term "likelihood".



Likelihood connotes the probability of an event occurring. In fact, you may see explicit probability ranges assigned to each qualitative label (e.g., "Very High = 90% to 100% probable"). And, while this seems to be on the right track, there are two problems with it:

- It often doesn't include a timeframe reference. In other words, does the likelihood statement refer to the probability of the event occurring this week, this year, in this lifetime?
- It doesn't provide the means to differentiate between something that may happen once vs. something that may happen multiple times. For example, a statement; "*The likelihood of a virus infection is Very High*" doesn't differentiate whether the event is likely to happen once or many times.

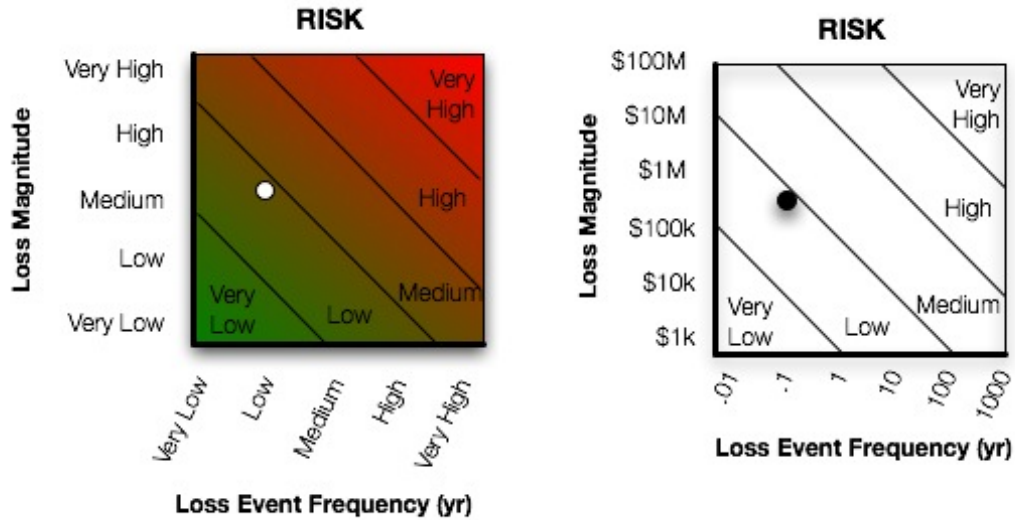
These two limitations become critical when we're trying to quantify and/or compare risk issues.

Using frequency, we can account for events that occur many times within the defined timeframe as well as those that occur fewer than once in the timeframe (e.g., .01 times per year, or once in one hundred years). Of course, this raises the question of how we determine frequency, particularly for infrequent events. In the interest of keeping this post to a reasonable length, I'll cover that another time (soon).

### Drawing lines

You may have seen charts like the ones below, with lines drawn to differentiate High from Medium, etc.

To be FAIR about it



(NOTE: Magnitude scales will vary based on the risk capacity/tolerance of the organization)

These can be useful, but a few challenges I've encountered with this approach include:

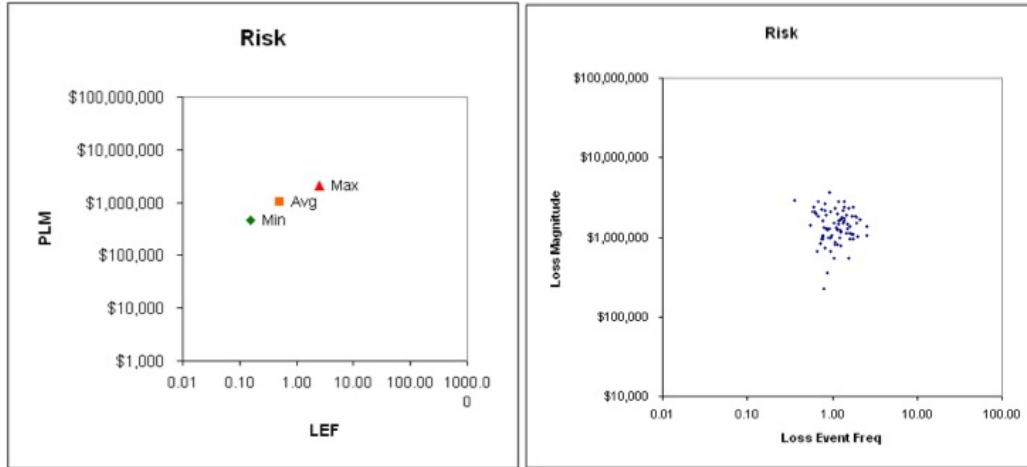
- If the risk point falls barely on one side of the line or the other, do the lines really serve a useful purpose, at least from the perspective of being able to assign a qualitative value?
- Who drew the lines? At one place I've worked, I couldn't get management to provide guidance on where to draw the lines so I took a stab at drawing them based on what I thought management's risk tolerance was given their earlier decisions. This seemed to work okay, as I didn't experience much push-back from management, but you need to constantly look for evidence that the lines need to be changed.

Particularly in larger companies with multiple affiliates or subsidiaries, line placement will vary because each part of the enterprise will have its own risk tolerance. A "critical" loss at the subsidiary level might not equate to a rounding error at the enterprise level. I've dealt with this by plotting results on two charts; one scaled to the enterprise risk tolerance, and another drawn to the subsidiary's tolerance.

Of course, the fact that the point isn't really a point at all, but the intersection of two ranges or distributions further affects the utility of lines.

I've found two ways of charting risk that seem to be well received by management (below).

To be FAIR about it



(NOTE: These charts were created using Monte Carlo analyses within FAIR-based applications)

My preference is the scatter plot, which does a nice job of visualizing the uncertainty that is a part of any risk analysis. A couple of things to note:

- No lines have been drawn to label the result "High", "Medium", etc.
- I haven't used a green-to-red background on the charts.

I will use those illustrative tools if requested by management, but I tend not to use them otherwise. Besides the challenges I noted above regarding lines, my rationale is that lines and colors tend to bias interpretation of the results. In other words, if someone sees a risk point plotted in a red background or in the "High" section of the chart, they equate those results as "unacceptable". The fact is, the acceptability of a risk condition is often dependent on the value proposition of the situation, the cost to mitigate risk, etc. I've found management is intelligent enough to know that the upper-right part of the chart means more risk than the lower-left.

## Physicians and Medics

(<http://riskmanagementinsight.com/riskanalysis/?p=360>)

My thanks to Mike Rothman who last week gave me credit for “[fighting the good fight](#)”. I’d like to think he’s right -- it has been a bit of a struggle over the years, I’d like to think I’m winning (or at least managing a draw) as I continue the struggle, and I’d like to think it’s worthwhile. Mike does seem to continue to question the pragmatism of my approach though, which is what this post is about.

Don’t get me wrong. I greatly admire the work Mike does and wish he and his book had been around when I started out as a CISO. Would have saved me significant pain and suffering. On the other hand, if I’d had Mike’s P-CSO I might have become complacent and ended up believing that’s all there was to being a CISO. Not that I think Mike is advocating complacency -- he’s not. I also don’t think he discounts risk analysis concepts. He’s simply focused on helping that component of our profession who’s just getting started or who faces other practical constraints in dealing with our very complex problem space. His is a necessary and highly valuable contribution, and he provides it in an entertaining way that’s too rare.

Let me set this discussion in a medical analogy context. If I was in the middle of nowhere or didn’t have the resources for a physician, then a medic who’s skilled in lifesaving basics would do just fine. However, if the situation called for a deeper understanding of the complex, sometime subtle health considerations, then I’d prefer a physician. Someone who **didn’t** say; “*Boy, this anatomy and physiology stuff is complicated. I’m just going to stick with ‘[The hip bone is connected to the back bone...](#)’*” My physician may, of course, choose to follow a pragmatic, commonly-used course of treatment, but they’d be able to do so with a deeper understanding of the problem space, greater (but not perfect) certainty that the course of treatment would work, and a better ability to explain to me, the patient, why I had to swallow this bitter pill, undergo the knife, or have this long tube snaked into one of my orifices.

Yes, I realize that physicians sometimes get it wrong, sometimes get wrapped up in fancy and even unnecessary procedures, and can drive up costs. That’s just as true as what can happen at the other end of the spectrum -- the shaman who operates entirely by superstition, faith, FUD, and intuition. The point is, there’s absolutely a need for both medics and physicians (and levels in between). We, as professionals, can choose where we want to be within that continuum. With this in mind, a few things to consider are:

- In the heat of battle, when resources are limited, or when it just makes sense, physicians always have the option of behaving as medics and sticking with the bare essentials (the reverse isn’t true). In fact, the best physicians I’ve encountered are pragmatic in their approach but have the deeper knowledge to leverage when need arises
- Medics might effectively deal with 80+% of our problems, but that remaining ~20% can be critical
- A person can start out as a medic and then become a physician later, as need and resources dictate
- Physicians tend to be paid more

Bottom line -- knowledge and understanding are never a bad thing, but it requires extra effort to acquire them. And, as Mike points out, the simple approach is often good enough and may be all we can hope for given our individual circumstances. For myself though, I prefer a deeper understanding of our complex problem space. I want to be able to answer the hard questions about why and how. But that’s just me.

To be FAIR about it

BTW - I was amused at Mike's characterization of risk analysis as Black Magic, as this phrase would also have been used in the past to describe medical and scientific concepts/practices we take for granted today.

# Compliance is Critical

(<http://riskmanagementinsight.com/riskanalysis/?p=369>)

Compliance has been getting a bad rap lately, and I'm here to set the record straight... compliance is CRITICAL.

Now, those of you who know me are probably picking your jaws up off the floor and asking whether I've suffered a stroke, have started drinking heavily, or have a gun pressed to my temple by a regulator or someone from the PCI lobby. Nope. I still have my full mental faculties (such as they are), and I make the statement without duress -- however...

## There's compliance, and then there's compliance

As usual, our profession tends to not be specific in our use of terms, which sets us up for confusion, inconsistency, and a host of other problems. When I say "compliance is critical", I don't mean compliance with some external standard like PCI, ISO, or some hypothetical "best practice". I mean compliance with an organization's own policies and standards. Compliance with external standards has its place too (unfortunately), but we'll pick that up in another post.

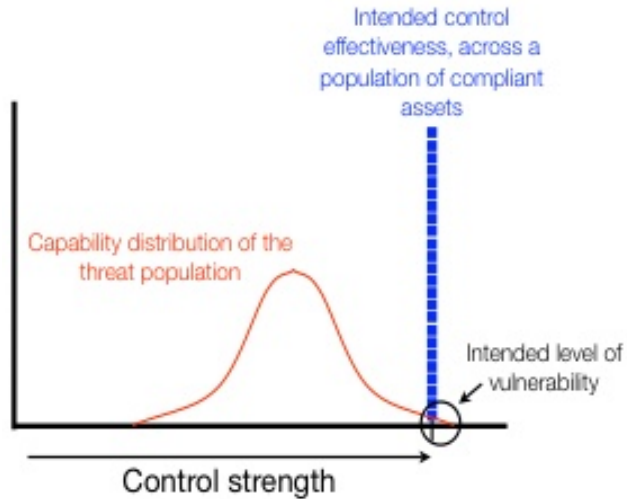
## Think about it...

In most cases, if an organization was completely, 100% compliant with its own policies and standards, it would almost certainly have a much lower level of risk exposure than most other organizations. In fact, in many cases a 100% compliant organization would be too secure to operate effectively. In other words, the more significant problem isn't typically a matter of how strong a policy is, it's the variance from intended/desired state that's described by policy.

## In a perfect world...

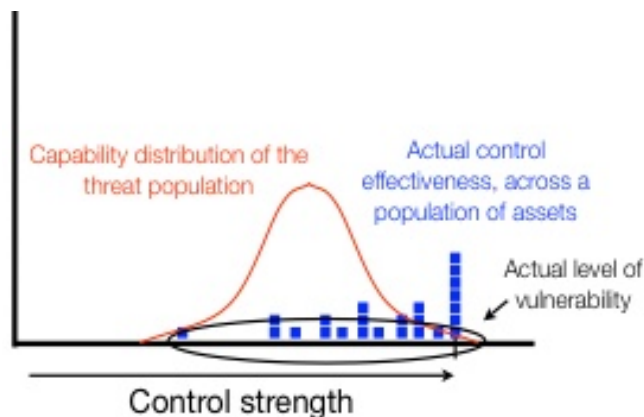
The illustration below is intended to represent a "perfect world" condition, where all of the assets/systems/whatever are compliant with an organization's policies/standards. It also reflects the fact that there is no perfect security, and that the organization has wisely established its policies/standards with an acceptance of some degree of vulnerability (and thus, risk).

To be FAIR about it



### The real world tends to be much different

The illustration below represents a more likely condition, where controls applied to a population of assets/ etc. tend to vary from what policy calls for. It also reflects the effect that has on vulnerability, which in turn affects risk.



### But we knew this already, right?

Yes, it's true that 99.9% of us already know that variability exists and that it's bad from a risk perspective -- so what's my point? My point is that variance is one of the most important risk-related metrics we have available to us. Here's why...

As we see from the illustration above, variance from policy can be a strong indicator of an organization's risk exposure. At the same time, it's also a marvelous indicator of an organization's ability to manage risk (i.e., decision making capabilities and/or the ability to execute against decisions). A little root cause analysis of a highly variant asset population can provide critical insights into what's not working, which can lead to far more cost-effective risk management measures.

One example of where this could be applied is in the evaluation of a third party's risk posture. Rather than send a 60 page questionnaire, why not evaluate the organization's compliance with its own policies

To be FAIR about it

across a cross-section of its information risk landscape. I submit that it would provide a more accurate and useful picture of risk exposure and risk management capabilities than the typical questionnaire, at less cost/effort to both parties.

## Aggregate Analysis

(<http://riskmanagementinsight.com/riskanalysis/?p=601>)

One of the questions I commonly encounter is "How do you take something like FAIR and apply it to a big problem, like measuring the aggregate risk within an entire organization?" In order to keep this post from becoming too long, I'll focus on the concepts of one approach in this post, and show a simple example in a following post.

### Measuring the surface area of Long Island

Imagine that you've been given the task of determining the surface area of Long Island. How are you going to go about it? One way would be to use a satellite photo like the one below and make a rough estimate based on the level of available detail and an appropriate scale.



With that information you'd probably report that the surface area is between X and Y square miles, with Z % of certainty. If this level of precision isn't good enough, you might want to carve the area into chunks, zoom in, measure each chunk, and then add them all up.



If even more precision is needed, you can carve the landscape into finer chunks...

To be FAIR about it



You could also take surveying equipment and go about the process of making measurements on the ground.

Extrapolating these approaches, it's conceivable (given sufficient time and tools) that you could try to measure the sub-atomic particles that make up the atoms, molecules, and objects that constitute the island's surface area. Let me know how that goes BTW...

Regardless of your level of abstraction and your approach, two things are clear:

There is a point of diminishing returns when it comes to measurement precision, and that point is found at the balance between how the measurement is going to be used and the cost/effort in making the measurement

There is always some degree of uncertainty and imprecision. Even at a sub-atomic level of measurement, the effects of erosion and other dynamic processes guarantee that by the time you've finished measuring, the subject being measured will have changed

## Measuring aggregate risk

In tackling the problem of aggregate risk, we likewise can carve up the landscape into chunks that are measurable and meaningful given our time and resources. One example might be to define a set of object groups such as:

- Network devices
- Network transmission media
- Servers
- Personal systems
- Applications
- Printed media
- etc...

To be FAIR about it

Harkening back to our notion of abstraction, you could carve the landscape even finer. Perhaps, for example, defining subsets of servers based on operating system, function, location, line of business, etc. Bottom line -- you'll want to define a level of abstraction that provides useful information given the available time and resources.

In order to perform a risk analysis, we also have to define our threat landscape. This, too, needs to be defined at an appropriate level of abstraction. For example, a high level breakdown might look like:

- External hackers
- Disgruntled employees
- Contractors
- etc...

You can also define threats that aren't malicious, such as....

- Acts of nature
- Weather
- Geological disturbances
- Extraterrestrial (No - not ET. I mean things like solar flares, etc.)
- Animals (non-humans)
- Errors & failures
- Employees
- Service providers
- Suppliers

... and/or become even more granular in your definitions.

With your landscape defined, you're in a position to begin a series of high-level FAIR analyses.

### Things to keep in mind:

It doesn't matter what level of "elevation" you're operating from, all measurements are estimates, which means there's always some degree of uncertainty and error. Therefore, the question isn't whether uncertainty and imprecision exists in a measurement, it's a question of whether the degree of certainty and precision is acceptable.

The acceptability of a measurement isn't for the measurer to decide. It's up to the people who are making decisions based on the information. What's important is that they understand the degree of (un)certainly and (im)precision in the measurements being provided.

Precision is often (always?) a function of the size, complexity, how dynamic the problem is, the quality of available tools and methods, and the time spent measuring. In other words, the more precision desired, the more time/money that's going to have to be spent.

Accuracy and precision aren't the same thing. I can be precise and not be accurate. For example an estimate that my 2009 income is going to be \$2,352,004.32 is highly precise. Unfortunately, it's not even close to being accurate (unless a miracle occurs). Conversely, an estimate that my income will be between \$50k and \$500k isn't very precise, but it's likely to be accurate. In an ideal world we could be accurate and highly precise. In the real world, at least when measuring risk, you want to be accurate and have an acceptable degree of precision.

To be FAIR about it

When performing an aggregate risk analysis, start at a relatively high level of abstraction and let the results from that analysis guide you regarding where to dive deeper. This helps to ensure that you find a feasible level of precision while managing your time and effort effectively.

Were we supposed to measure surface area at high tide or low tide? The point is, it is crucial to be very clear about what's in and out of scope within the analysis, as unstated and misaligned assumptions can result in measurements that don't meet management's objectives.

## Some Thoughts on "Physics Envy"

(<http://riskmanagementinsight.com/riskanalysis/?p=636>)

I just can't resist the need/desire to comment on someone else's misleading post on the subject of risk analysis.

This time it's [Richard Bejtlick, who appears to be on the warpath again](#). And, while Richard is extremely intelligent and well established as an expert in many security-related matters, I'd argue that he's not an expert in risk. It's clear that he reads on the topic, but he appears to interpret it through an infosec lens which, I believe, tends to be badly distorted by some unfortunate/inaccurate biases in the industry. That said, let's examine some of his concerns and see what we can learn...

### "False precision"

On this point, Richard and I agree -- the notion of precision in risk analysis (whether infosec-related or some other form of risk) is absurd. The future is uncertain, and risk analysis is fundamentally a discussion of the future. A precise statement (i.e., prediction) of exactly when something will happen, or how often, or to what effect, just isn't feasible in a complex problem space like risk. Where Richard and I appear to differ, however, is in our understanding of what risk analysis is and isn't.

Risk analysis isn't (or shouldn't be) put forth as a prediction of the future, but rather as a statement of probabilities given what's known or believed. Much like a statement that there's a 1/36th probability of rolling snake-eyes given that a pair of dice has six sides each, and that the dice have independent probabilities. Nobody in their right mind would believe they could predict on which roll the dice will come up snake-eyes, but it's still very useful as a decision-maker to know what the probabilities are.

Analysis results also should reflect the degree of (un)certainty involved, so that people making decisions based on the analysis have realistic expectations. This is where the use of distributions and ranges become very useful in portraying uncertainty and imprecision. However, if I read Richard's post correctly, he considers all risk analyses to be useless because they can't predict the future (i.e., aren't precise).

This "*if it ain't precise, it ain't useful*" position is one I run into frequently in the infosec community, presumably because the profession is made up of so many people with engineering backgrounds who are used to measuring things relatively precisely. Or, maybe, Richard's just pointing out that many risk analyses are flawed because they don't do a good job of conveying the degree of imprecision involved. He's really not clear on that, and seems to paint the entire issue with a single broad brush stroke.

### "Overweighting things that can be counted"

Here again, I tend to agree with Richard about the fundamental problem. There is an unfortunate tendency to look around us for things that can be easily counted, and then assume that they comprise the whole picture. This may not be as significant a problem if you're dealing with something that has a large volume of relatively clean data to work from, but is a huge problem when good data are sparse. For example, if I want to construct models of human life expectancy in order to profit as a life insurance provider, I can probably find enough good data to do so. Unfortunately, in the infosec realm, good data are harder to come by. As a result, models derived from available hard infosec data are much less likely to be complete/accurate.

To be FAIR about it

What I've described above, however, is an inductive approach to modeling -- i.e., evaluate data to derive a model. The other approach to modeling is deductive. I'll spare you a long-winded comparison and let you research these further if you're interested but, simply stated, a deductive approach constructs a model based on logical (and believed to be) true relationships between premises. For example, a model that states "*Loss events are predicated on threat events and vulnerability to those threat events*" is logical and "true". We didn't need data to construct that model -- it just makes sense logically.

Does that mean that deductive models are always accurate? Heck no. They're subject to potential problems too, but if they're well thought-out they're less likely to have the gaps an inductive model built on sparse data is likely to have. Deductive models also act as a guide to knowing what data we need in order to perform analyses.

Keep in mind though, that "[\*All models are wrong \(i.e., imprecise\), some are useful\*](#)". As I've said before, the world is far too complex to model exactly. Nonetheless, we should be looking for accuracy and a useful degree of precision in our models, which is entirely feasible. [Also, all models should be flexible enough to be adjusted as data and experience improves.](#)

### "Man with a spreadsheet syndrome"

Richard's basic concern is valid -- spreadsheets often connote a sense of validity that isn't always warranted. It isn't logical, however, to conclude that because some spreadsheets (or other quantitative tools) are flawed, all must be. Maybe that isn't what he meant, but Richard tends to use a broad brush, so it's hard to tell sometimes what he's really saying.

At the end of the day, the validity of a spreadsheet boils down to model accuracy and the quality of data. Since we've already covered models, I guess it's time to cover data...

### "A lot of guessing"

As Alex states, Richard (and others) toss the term "*guessing*" around like it's an insult. If what Richard means by "guessing" is "*estimates made in the absence of perfectly complete and precise data*", then welcome, Richard, to reality. All measurements in the real world are guesses to some degree. I assume, however, that what Richard is concerned about is whether the estimates (guesses) are accurate. The answer to that, of course, is "it depends".

If someone asks me what the wingspan of a 747 airliner is, and I answer "*Ummmm, I dunno. A hundred feet?*", then maybe we have a problem. For one thing, I've given a relatively precise answer (100 ft), but that answer may not be accurate. If, however, I answer "*Well, the wingspan is almost certain to be less than the length of a football field (300 ft) but greater than the length of my driveway (80 ft)*" then I've made an estimate that isn't precise but is much more likely to be accurate. With a little work, I can probably narrow the range significantly (i.e., get better precision) and still be accurate (especially if I have access to a subject matter expert). The question of whether it's precise enough (i.e., is useful) is a matter of what I need to use the information for.

Business decisions of almost any sort are based on imperfect and imprecise estimates of what might happen. That's reality. As long as decision-makers are aware of and okay with the imprecise nature of the information they're operating from, then it's not a problem.

What people seem to forget is that whether we perform formal analysis on a problem or not, a decision is still going to be made. The question then becomes, is the decision-maker going to be using conclusions drawn from:

To be FAIR about it

- Someone's unstructured, undocumented, mental model and the "guesses" they apply to it, or
- A structured model that has been documented, examined, and evolved through use, and "guesses" that have been given due consideration

Either way, some model will be applied and guesses/estimates will be used. The point of analysis is to give decision-makers better information than they would have had in the absence of analysis.

Bottom line -- it seems like Richard has accurately recognized the existence of some of the fundamental challenges in risk analysis, but it feels like he's drawn some extreme conclusions about their significance and the ability to effectively deal with them. It could be, of course, that the problem is simply the manner in which he described his conclusions. Perhaps he'll respond and clear it all up.

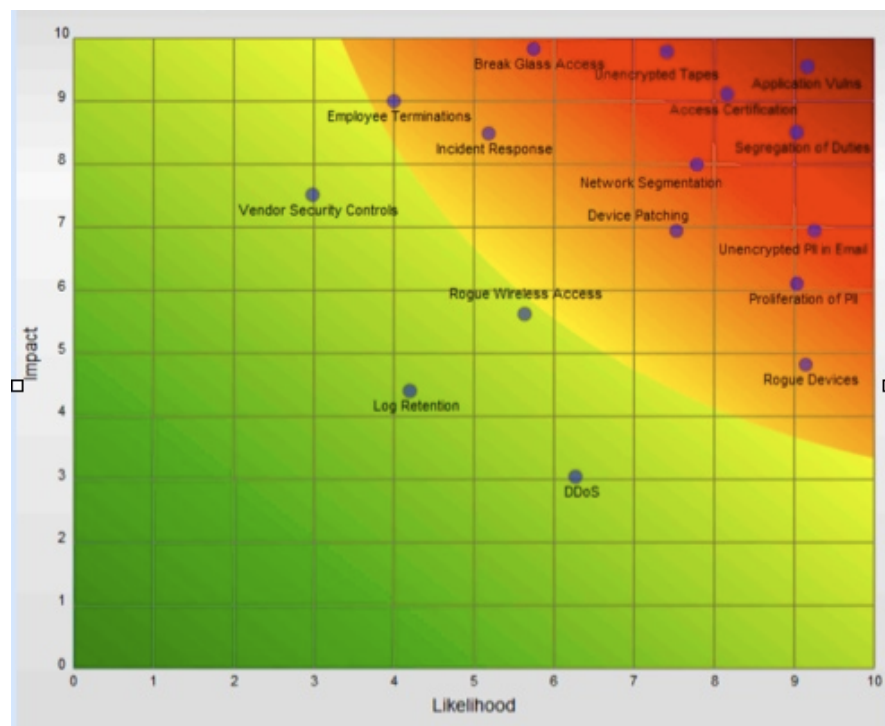
## Lipstick on Pigs

(<http://riskmanagementinsight.com/riskanalysis/?p=654>)

There are a lot of really slick looking “risk” assessment solutions on the market today, with lots of eye-candy and great marketing verbiage to go with them. Look under the makeup though and, well, you're going to see pork....

Now, before I start to point out the problems with those products, let me first say that some of them do provide value. Not the nearly the value they're claiming, but value nonetheless. At the very least, they can provide an interface and structure for their users to identify “stuff” they view as problematic, a way to apply some sort of significance rating for that stuff, and generate reports about the stuff. So in that sense perhaps they make us more efficient at what we already try to do. They also help us look good in front of the auditors and other stakeholders who aren't too interested, really, in digging in to see whether there's any meat on those bones.

What these products don't do, at least none of them that I've seen, is actually measure risk in any sort of defensible way. In fact, many of them don't even really seem know what risk is. For example, let's take a screenshot from one such product (which shall remain nameless).



### Which of these things are not like the others...?

On this screen are shown a number of “risks” and where they (supposedly) stand from a significance perspective. My first observation is that they're measuring apples, oranges, peas, and carrots here. Some of the “risks” are really control deficiencies (e.g., Unencrypted PII in email). Others are threat events (e.g., DDoS), and yet others appear to be security program elements (e.g., Incident Response), policies (e.g., Log Retention), or problematic characteristics of the asset landscape (e.g., Proliferation of PII). Clearly, each of these issues contribute, to some degree, to how much risk an organization has, but they're

To be FAIR about it

fundamentally not the same thing and therefore can't be measured and compared to one another -- at least not the way they're trying to measure them in this tool. Furthermore, they aren't independent of one another. For example, the "Proliferation of PII" (assets) is potentially a bigger problem within an organization that has "Rogue Devices" (control deficiency) and that does a poor job of dealing with "Employee Terminations" (threat landscape management).

### What are they measuring, anyway?

Their measurement scale for risk is, appropriately, Impact and Likelihood. That should be good news. Unfortunately, their application of it is pretty much useless. For example, looking at the *Incident Response* issue, what is "Likelihood" supposed to represent? The likelihood of a breach due to poor incident response? The likelihood that incident response is not going to be effective if an incident does occur?

What about Impact? How do they come up with and "8+" Impact score for *Incident Response*? What kind of incident are they referring to? What if the incident is minor -- i.e., someone changed the cafeteria menu?

### What's it mean?

While I'm on the subject of measurement, what does an Impact score represent here? Is it a most likely outcome, average outcome, or a worst case outcome? And what does a "9" represent? How much worse is it than a "7", or a "2"? Unless these values map to something meaningful to the business, it's just noise.

Similarly, does the Likelihood scale represent the likelihood of something occurring this year, this week, in our lifetime? What timeframe is this being measured against? And where on the scale do we account for things that might happen many times within whatever timeframe we're referencing?

Without a lot more clarity regarding what these scales represent, the results are ambiguous to say the least.

### Where the problem really lies is...

...with how our profession approaches "risk". A user of this tool who doesn't try to measure apples vs. oranges vs. peas, etc., but who instead only enters risk issues that can be measured in terms of frequency and magnitude will be FAR better off. They'll still be stuck with poorly defined and marginally useful ordinal scales, but at least they'd be measuring stuff that makes sense.

### At the end of the day...

All these tools are, are a way for an infosec professional to identify things that bother him/her about their current risk landscape and apply a rating that represents how much sleep they lose at night worrying about them. There's nothing wrong with that. Unfortunately, the tools tend to be designed poorly, we use them poorly, and then we try to pass the results off as something they're not. Oink.

## Lipstick on Pigs - Part II

(<http://riskmanagementinsight.com/riskanalysis/?p=680>)

In a conversation with Jared, he shared that the RC roadmap includes plans to improve its built-in model. And clearly, new product development always requires tough choices, trade-offs, and continual improvement, so I understand and empathize. He also reiterated that RC's first focus was on portfolio management -- i.e., helping security practitioners prioritize and communicate, which is a critical need in our profession. My reply was that if prioritization is going to be based on risk, then the method/model used to evaluate risk is foundational to the product's value. He agreed, but seems to have a much higher level of confidence than I do in how our profession approaches risk.

At Jared's suggestion, I logged into the RC demo and dug into the information in their help file that serves as guidance for its users. However, rather than post specific observations about RC, I thought it would be more helpful if I simply provided a brief "*Thinking Person's Guide to Risk Assessment Tool Selection*". Okay, maybe "*Things to watch out for*" is a better description. Regardless, I hope you find it useful.

### "First, do no harm" (Auguste Francois Chomel)

The phrase above was borrowed from Douglas Hubbard's book "*The Failure of Risk Management*". (Buy it. Read it.) One chapter in the book is entitled "*Worse Than Useless*", and in there he describes "structured" scoring methods that can, in fact, lead to worse decisions than if no scoring method was used at all. To limit the length of this post, I'll refer you to Douglas' book rather than repeat it here. Suffice it to say, amongst other things he describes the same concerns I've already posted about ordinal scales and scoring.

### There's likelihood and then there's "likelihood"

Many information security risk assessment tools view "Likelihood" as a measure of how likely it is that an attack will be successful. This is VERY different than a measure of how likely it is that an attack **will occur** and be successful. Without including the likelihood of occurrence we could rate the "Likelihood" of my being attacked by a polar bear on the streets of Dayton Ohio as "high" because I have no effective defenses from such an event. Bottom line -- understanding likelihood of success is not very useful if I don't also understand the likelihood of occurrence.

Of course, the first argument that someone's likely to raise is, "*But we don't know how often some of these events occur!*" I'll talk more about this in a future post, but the short answer is:

Baloney. I sit down regularly with clients who need to evaluate the risk associated with "rare" events or events where no direct evidence exists to draw from, and we're able to arrive at frequency ranges that make sense and can be defended. The key here is the term "ranges". We may not have the information we need to state **exactly** how frequently events might occur, but we absolutely have the means to generate frequency as a range. Again -- read Douglas Hubbard's work.

### Ambiguity and overlap

Besides the problems inherent in ordinal scales and scoring, another very significant problem is the lack of clarity and specificity in the elements being measured. Unfortunately, many of the models I see in use are very poorly defined, with lots of ambiguity and overlap/redundancy between variables. The result is

To be FAIR about it

that things are accounted for and measured multiple times. Combine this with the ordinal scale problems, and the results are not defensible under any sort of scrutiny.

## CMM limitations

Some models use a CMM scale to rate the effectiveness of controls. And although CMM is useful for rating process maturity, it's not intended for nor effective at rating technical controls.

## Compensate not, lest ye go awry

Many models have only one place to rate controls, and those control ratings tend to be applied solely to the Likelihood component of risk. (This was a problem in the first version of FAIR). Typically, what happens then is that users throw compensating controls in that bucket too, even though some compensating controls (e.g., recovery capabilities) affect Impact rather than Likelihood. As a result, the effect of these controls are accounted for in the wrong part of the equation.

## Chicken Little

The Impact ratings in most assessment models focus on what “can” result -- i.e., what’s “possible”. And, being the paranoid lot that most of us are, we turn this into an estimate of what a worst-case outcome might look like. I don’t know what your experience has been, but out of all of the incidents I’ve been witness to and victim of over the years, not one has approached a worst-case result despite the fact that some of them had significant potential for really nasty outcomes. In fact, as I’ve discussed this with colleagues in the past it’s become clear that worst-case outcomes are extremely unusual. By characterizing risk events purely in terms of worst-case outcomes we provide an exaggerated view of risk, which management recognizes intuitively and writes off as “Chicken Little syndrome”.

The simple fact is that outcomes from incidents can range from inconsequential to catastrophic. And although we can’t predict precisely which will occur from any future event, there are factors that we can use to help us understand and communicate the range of possible outcomes from worst-case to best-case and even what’s most likely. If we want to communicate useful and believable risk information to management, we need to be able to deal with loss magnitudes other than just the worst-case outcome.

## To summarize...

There are other issues I could raise, but here’s the short list:

- Be very skeptical of methods that use addition, subtraction, multiplication, or division with ordinal scales. If you do choose to use them, recognize that at the end of the day you’re not going to be able to defend the results as truly quantitative, and you may have a very difficult time defending their legitimacy.
- Make certain that Likelihood includes a frequency component or, better yet, that Frequency is used instead of Likelihood. Regardless, without some reference to the frequency/probability of occurrence the information’s usefulness is significantly reduced.
- Elements being measured, particularly if math is involved, must be as clearly defined as possible so that redundancies and overlaps can be avoided. This also helps to prevent having the wrong element in the wrong part of the equation.
- “Quality” scales like CMM should only be used to evaluate the things they’re intended for

To be FAIR about it

- If the tool only allows the user to describe one level of Impact (e.g., “High”), there’s a significant likelihood that users will choose a worst-case outcome. This almost invariably inflates the risk rating well beyond the actual level of risk, which increases the probability that management won’t take the results seriously.

Bottom line -- if we want our risk analyses to be taken seriously, it’s critical that we challenge the assumptions and models (including FAIR) underlying our tools. Unfortunately, much of what I encounter in our industry’s risk assessment tool kit are examples of faux sophistication and poor definition. Is it any wonder then, that many within our profession struggle to accept risk analysis as a viable approach?

### So how long DOES it take the Sun to orbit the Earth?

I do need to reiterate my concern about a “model-less” analytic tool. As Jared clearly states, RC is not constrained to any one model for measuring risk. It’s intended to be an efficiency tool that allows the use of any risk assessment model. From a marketing perspective that may be pure genius, I don’t know. I suppose it could translate into a larger potential market because they wouldn’t be locking out those clients who are strict adherents of one model or another. And certainly, if a user is leveraging a reasonably accurate model, then the tool’s effect should be very positive. Unfortunately, as I described in the first part of this post, much of what our profession uses to model risk is junk. In that case, a model-neutral tool is a bit like saying to an astronomer, *“Hey, if you want to analyze the solar system by modeling the planets and Sun orbiting the Earth -- go for it. And while you’re at it, if you’d rather measure gravitational pull in bushels rather than units of acceleration, that’s cool too. We’ll still allow you -- in fact we’ll help you -- to present the results as valid astronomy.”*

The fact is, models matter. A lot. In my next post I’m going to talk about the role models play and I’ll also draw a distinction between the different types of models I see our profession using.

## Models Matter

(<http://riskmanagementinsight.com/riskanalysis/?p=696>)

Let's say you're approaching an intersection and the traffic signal turns yellow. What do you do - slow down and stop, or hit the accelerator? The answer for most people, ultimately, is; "*It depends*". How fast am I going? How far am I from the intersection? Is there someone close behind me? Is there a police cruiser at the intersection? What is the road condition? Am I in a hurry to get somewhere? These are just a few of the considerations that may flash through our minds in an instant - at least some of them subconsciously. We then make a decision and act on that decision.

In that instant when we see the signal change color we take in and process a remarkable amount of data and analyze the scenario using whatever mental model we'd developed through experience and education. We then instantly apply the results of that analysis against our own tolerance for the different forms of risk that are in play (health/safety vs. legal, etc.). If our data and model are reasonably complete and accurate, then we probably survive these events with an acceptable frequency and magnitude of loss.

That "mental model" is a construct that represents our understanding of how the different elements in the decision play together. If our mental model is missing key elements -- e.g., the effect of icy road conditions on our ability to stop -- then our decision is much more likely to have an undesirable outcome. The same is true if our model contains erroneous or inaccurate structural elements or relationships -- e.g., a belief that icy roads will improve our ability to stop.

When we're faced with a decision regarding information security we will also apply a model. The model might be an informal mental model or something more formal like FAIR, TARA, CVSS or a host of other candidates. So the question we want to ask ourselves is; "*How accurately does the model we're using represent the problem we're trying to understand?*"

### Where models fit

It's implied above, but the strategic role risk models play in an organization's ability to be successful is outlined below:

- Effective management decisions are predicated on...
- Effective comparisons between the issues/options that are in play, which are predicated on...
- The ability to measure the issues/options in a meaningful way, which is predicated on...
- An accurate model (understanding) of the problem and its elements

The above is true regardless of whether you're using an informal mental model or something formally defined. Consequently, you can't expect to consistently and effectively manage a complex problem space if the underlying model for measurement and comparison is badly broken.

### An example

My last two posts already described some of the obvious and important ways in which risk models commonly used in infosec are broken. But how badly broken are they? For the sake of brevity, this blog post will not include a blow-by-blow analysis -- I'll save that for another time. I will cite an example, from when I was a CISO, of how a flawed model almost had a significant impact on my employer...

We'd brought in a big-4 consulting firm to perform an attack and penetration exercise against us. At one point in the exercise they came to the table claiming that they'd identified a number of "high risk" issues that needed to be addressed immediately. I took one look at those issues, applied a quick mental sniff-test, and told them they were wrong. I didn't believe any of those issues represented a level of risk that warranted a high impact (to the business) response. They agreed to sit down and review the issues with me so that they could show me the error of my ways. However, after we broke down each of the issues in detail using FAIR, they conceded that none of the issues warranted an immediate, high-impact response.

Their original analysis (measurement) of the issues was shown to be based on an inaccurate model of the problem (risk). This flawed measurement led to an inaccurate comparison of these issues versus the other issues and priorities the business faced, which would have had a significant negative effect on the business if we hadn't recognized the flawed analysis. (The flaws in their model involved how they treated threat event frequency and loss magnitude.)

Now, please don't interpret this as an indictment of big-4 firms. They have a lot of very bright people who do marvelous work. And besides, I'm pretty confident the scenario would have played out similarly with almost any firm because the big-4 firm was using a very common assessment method. The point is, if your model is broken badly enough, the results can significantly affect your organization.

### There are models and then there are "models"

As I see it, there are three types of "models" being used in our industry.

- Checklist "models" (e.g., ISO, PCI, etc.)
- Maturity models (e.g., SEI's CMM for software development)
- Analytic models (e.g., FAIR, TARA, CVSS, etc.)

It's a matter of opinion (/religious debate) whether checklists qualify as models. It doesn't matter to me what they're called as long as we understand what they are. Checklists are simply a set of security and/or risk management elements somebody believes is relevant and important. Presumably, if you follow the checklist you're better off than if you don't follow the checklist, and for most of the checklists in our industry, I'd agree (up to a point). So if you're looking for a quick and dirty "*are we generally doing the kinds of things we should be doing*" litmus test, then checklists are fine. They're also fine for comparing one organization against another ([which has some potential pitfalls](#)), and showing progress against, for example, last year's checklist results. The downside to checklists is that they tend to be one-size-fits-all, they don't help us prioritize or compare our options, and they don't help us understand why the different elements are important or how important they are.

Maturity models tend to focus on measuring process effectiveness and process improvement on a relative basis -- two very worthwhile objectives. What they don't do is explain the practical effect of process improvements -- the "why" or "how much". Consequently, similar to checklists, maturity models don't help us prioritize or compare options.

Analytic models (some might call them scientific models) attempt to describe how things work. If they're designed well and used well, they enable the practical and useful measurement of complex systems (systems in the scientific sense vs. the IT sense), explanation of cause and effect, and sophisticated what-if analyses. In other words -- if you want answers to questions like:

- "How much risk do we have?"
- "How much less/more risk will we have if...?"
- "Which of these issues is most significant and by how much?"
- "Which of these mitigation options is likely to be most cost-effective?"

To be FAIR about it

... then analytic models are the way to go. If, however, they're designed badly or used poorly, then they can very easily lead to inaccurate conclusions and poor decisions.

Bottom line -- all three approaches have their benefits and limitations. For most organizations, some combination of them will be the best bet for effective risk management. Speaking of which...

## Does organizational maturity play a role?

Steve Dotson raised a very good question in his comment to [Lipstick - Part 2](#). Paraphrasing -- he asked whether more mature organizations are better able to answer/analyze risk-related questions. The short answer is "probably".

More mature organizations generally have a more complete picture of their risk landscape -- i.e., they likely have better visibility into what their assets are, where their assets are, the control conditions surrounding their assets, the threat landscape, and the loss implications from events. This information should enable them to provide more precise data for analyses with better confidence than less mature organizations. That said, less mature organizations can still get accurate and useful results from analyses -- just often with less precision. And, having gone through an analysis, the less mature organization can acquire a very clear idea of where their information gaps are, the significance of those gaps, and what can be done to fill those gaps.

Organizational maturity also can play a role in how much emphasis an organization will likely place on checklists vs. maturity models vs. analytic models. An organization that's very immature may place primary emphasis on checklists, just to get things moving in the right general direction. They may only use maturity models to gauge where a few key processes are today and set preliminary goals for improvement. They also may initially limit the use of analytic models to key, high-impact decisions.

More mature organizations often are looking to become more cost-effective and/or need to make business cases for continued improvement. Good risk analyses can make that possible. Speaking from personal experience, once you get your security/risk organization to a point where management no longer views it as the brightest/hottest fire burning in their landscape, it can become very difficult to get their attention (unless that attention comes in the form of cutbacks...). Of course, even some "mature" organizations don't have their security/risk ducks in a row and struggle to get management to care. For these organizations, being able to explain "why" and "how much" through good risk analyses can be very important. Conversely, taking lame risk analyses results to management can erode credibility and make future dialog even tougher.

## Managing Inconsistency

(<http://riskmanagementinsight.com/riskanalysis/?p=726>)

In the LinkedIn discussion mentioned earlier, some very legitimate concerns were raised regarding the inconsistency (variance) that can exist between risk analyses performed by different individuals. Because not everyone is watching that discussion (and probably many who were got tired of it and moved on) I thought I'd post my thoughts on the consistency problem here.

For the sake of clarity, "consistency" as discussed within this post equates to "the likelihood that two independent analyses of a specific risk scenario will result in similar outcomes". In other words, analyst A's results will look very much like analyst B's.

In order to frame the problem, we should ask ourselves where inconsistency tends to come from. In my experience, there are four key sources of inconsistency within risk analyses:

- The scenario is scoped differently -- i.e., analyst A is operating from a different set of assumptions than analyst B (e.g., is including different threats, different assets, etc.). This, BTW, is a huge contributor to variance -- in many cases it's the single most significant contributor.
- The analysts are operating from different analytic models -- i.e., one analyst is using a model consisting of variables X, Y, and Z, while the other is using a model consisting of variables X, Y, and W. The models also may have different underlying formulas. The opportunity for inconsistency is especially problematic when analysts are using their own "mental models" for analysis, versus a structured model that can be explicitly referenced.
- The analysts may have different experience levels and data sources -- thus analyst A may estimate variable C to be between "5 and 10", and analyst B's estimate for the same variable may be between "40 and 100".
- Some people are lousy at estimating.

The first and second sources of inconsistency can be dramatically improved by ensuring that the analysts are singing from the same sheet of music -- i.e., using the same model/method for analysis.

The third source of inconsistency can be significantly reduced (but not eliminated) by getting the right subject matter experts involved in the analysis. For example, as a security/risk geek I shouldn't be estimating reputation damage. That's the domain of business personnel. It also helps to have more than one person involved in the analysis to increase the experience and perspective the estimates are based on.

The fourth source of inconsistency can be significantly reduced (but not eliminated) through calibration training similar to what Douglas Hubbard presents in his book "How to Measure Anything". You'd be surprised at how much improvement can be realized.

Bottom line -- inconsistency in analyses is manageable to where the degree of variance is not significant relative to the decisions being made and the inherent uncertainty in the data.

## Usefulness?

(<http://riskmanagementinsight.com/riskanalysis/?p=752>)

In the LinkedIn discussion regarding risk analysis, the question came up of whether risk analysis is even useful and, if so, how. Before diving in -- it has been pointed out that context is critical in any discussion, and particularly so when you're talking about something like "usefulness". "Useful for what?", is a question that's begged. "Useful to who?" is another. So, for the sake of clarity in this post, the "who" is business management and the "what" is better informed decisions.

My experience has been that organizations rarely have the resources necessary to address all of the business opportunities, operational costs, and risk issues (of many sorts) that they face. Consequently, they're forced to choose what to do now, what to do later, and what not to do at all. These choices invariably require some form of comparison between the issues, and comparisons are invariably based on some form of measurement (whether formal or informal/intuitive).

Between the three categories of issues (opportunities, costs, and risk), operational costs are probably the easiest to evaluate and forecast, although there's plenty of opportunity for operational factors to change in ways that weren't anticipated. Forecasts regarding opportunities and risk generally are much more speculative, and it's rare to see realistic business opportunity analyses that aren't expressed as ranges and/or distributions, with some form of confidence statement. After all, market demand, the competition, regulations, politics, and the organization itself (amongst other things) may change or behave in ways that weren't foreseen. In other words, despite best efforts, what was projected may turn out to be inaccurate.

Nonetheless, despite this inherent uncertainty, I suspect most business decision-makers would agree that they'd prefer to base their big opportunity decisions on some form of structured analysis that can help them understand what's known, what's less well known, and what's highly speculative. They would probably look pretty skeptically on conclusions and recommendations regarding a large, complex business opportunity that didn't have a structured analysis behind it, if for no other reason than without the analysis they wouldn't be able to challenge the assumptions and data the conclusions and recommendations were based on.

In my experience, business decision-makers likewise appreciate the information a well-structured risk analysis can provide. It allows them to understand how we came up with our conclusions and recommendations, and allows them to challenge our results if our assumptions (particularly regarding loss magnitude) seem off-base. If the analysis presents the information in more meaningful business-like terms (e.g., annualized loss exposure, or whatever), then so much the better.

### Example

A project where I worked had been identified as having a significant amount of risk (a new application required users, a lot of them, to have admin access on their workstations/laptops -- sound familiar?). Unfortunately, the project was the "pet" of one of our senior executives; a man who wasn't a fan of infosec to begin with. Worse, the recommendations we were making were going to delay the project and increase its costs.

As I anticipated, the executive was ummm... not happy with our conclusions and recommendations. I believe his words were, "You guys are always crying about high risk, why should I pay any attention now." At that point I explained that we had recently begun taking a more structured approach to analyzing risk, and I asked if I could describe how we came to our conclusions on his project. After about a five minute explanation and whiteboard demonstration of the analysis his response was, "Well, it's hard

To be FAIR about it

to argue with that. Let's look at your recommendations again." In the end, he agreed to all of our recommendations. More importantly, he became an advocate for us because he had an entirely different perspective on how we approached our work.

Useful? Yeah, I'd say risk analysis can be useful.

## What's a "risk" anyway?

(<http://riskmanagementinsight.com/riskanalysis/?p=765>)

Although there are a number of definitions for "risk" out there, most of us seem to gravitate around a definition that relates to the likelihood (or frequency) and consequences (or magnitude) of loss. So with that in mind I'm going to ask a question about something that's bugged me for a long time -- What is "a risk"? Likewise, what are "risks" (the plural of "a risk")?

If you survey a set of people who deal with risk or security professionally (inside or outside of infosec) and ask them to list key "risks" within their scope of responsibilities, you tend to get an interesting set of answers. For example, the list you get from an infosec professional might look something like:

- Insiders
- Lack of user awareness
- Data leakage
- Non-compliance
- Reputation
- Web applications

### Why it matters

Clearly, all of these can be issues worthy of concern for an organization, so what's the problem? Well, maybe nothing. If all you're looking for is a list of issues that contribute to the amount of risk an organization has, then a list like this is probably fine. A problem arises though, when you try to measure, compare, and/or prioritize these, for (at least) two reasons:

- They aren't the same kind of thing -- e.g., Insiders are a threat community, lack of user awareness is a control deficiency, data leakage is a type of loss event, non-compliance is a condition, reputation (damage) is an outcome, and web applications are a type of asset. A very apples vs. oranges problem.
- They aren't distinct or solitary in their contribution to risk. In other words, two or more of them can be combined in different ways to describe different risk scenarios with different probabilities and consequences. As a result, any individual measurement of significance in terms of risk is invalid.

The definition for risk mentioned above implies a measurement of some sort -- i.e., a pair of values (some version of likelihood and consequence) -- yet we use the terms "a risk" and "risks" in a way that implies reference to one or more objects or "things" rather than a value.

Unfortunately, I see a lot of instances where people have tried to characterize "risks" in terms of likelihood and consequence, and it's never pretty. The results are very difficult to defend logically, which I suspect contributes to people's notion that dealing with risk is hard. My experience has been that once you get clarity around risk terminology the kind of confusion that comes from "risks" goes away and the problem becomes a lot easier to wrap your head around.

## How Much Risk...

(<http://riskmanagementinsight.com/riskanalysis/?p=774>)

As professionals, two of our objectives include helping management understand:

- Which issues they should be concerned about, and
- The importance of each issue (for prioritization's sake)

As a profession, an argument could be made that we're pretty good at the first objective -- identifying problems. Even the lamest amongst us can grab a checklist of controls and an inventory of assets and threats, and pretty quickly identify a laundry list of concerns within most organizations. Giving us the benefit of the doubt, we're also reasonably good at intuitively identifying, roughly, which of the concerns in the list we believe are most problematic.

Okay, so let's say we've boiled down our laundry list to a set of top-ten "risks". And logically, I suspect that anything making the top-ten would be labeled a "high risk" risk. If this list is like most I've seen, it contains things like:

- Inappropriate access privileges
- Absence of patching
- Disgruntled insiders
- Wireless access points
- Data leakage
- Lack of user awareness
- Reputation damage
- Regulatory compliance
- Etc...

So what, if anything, is wrong with this list of "risks"? If you're just looking for an inventory of things that may significantly contribute to how often bad things are likely to happen and how bad they're likely to be when they do happen, then the list is probably fine. However, by calling each of the things in this list "a risk", several problems occur...

### Which of these things is not like the other?

By calling each of these things "a risk", a lot of people seem to lose sight of the fact that many of the issues in the list are fundamentally different. In fact, the list above contains six very different categories of issues:

- Control deficiencies (inappropriate access privileges, absence of patching, lack of user awareness)
- Threats (disgruntled insiders)
- Assets (wireless access points)
- Scenarios/events (data leakage)
- Objectives/requirements (regulatory compliance)
- Outcomes (reputation damage)

To be FAIR about it

On the surface, this taxonomical blurring of the things that make up our risk landscape may seem to be no big deal. After all, most of the time won't a conversation's context provide clarity about what's being said? Perhaps. However I can't count the number of times I've seen confusion and miscommunication within a conversation between two professionals because one of them is using the term "risk" to mean a threat when the other person is assuming the term "risk" means vulnerability, etc. Yes, it tends to get figured out eventually, but it contributes to inefficiency and is a hallmark of a profession whose act is not together.

### Perhaps even more importantly...

...you can't measure "a risk" -- at least not most of the "risks" you see in these lists. In fact, out of the list above only one (data leakage) is what I would call a risk scenario that can be evaluated in terms of frequency and magnitude of loss. All of the other "risks" in the list are really elements (subcomponents, if you will) that, until combined with other elements, don't make up scenarios that can be measured meaningfully. As a result, we're unable to meet the second objective given at the start of this post -- helping management understand the importance of the issues.

But wait a minute. Didn't I state earlier that by including these issues in a top-ten list we had at least implicitly provided information regarding their importance? Yes I did, and no we haven't -- not meaningfully anyway.

### Choices, and what does "high risk" mean, anyway?

Running a business requires constantly balancing resources against opportunities, operational expenses, and risk issues of various sorts. In most cases the opportunities (at least in the commercial world) and operational expenses are understood quantitatively. Likewise, what we might think of as more "traditional" business risk issues (e.g., credit, investment, etc.) also tend to be expressed quantitatively. Unfortunately, when all we do is slap a qualitative label on something we call "a risk", management is left to wonder where, really, it fits within the portfolio of all the other things they have to choose to deal with. And, because many of the "risks" in our lists can't individually be measured meaningfully, it can become very uncomfortable for us if/when management pushes back and asks us to explain/defend our risk ratings.

### I have to wonder...

...whether a lot of us intuitively and subconsciously recognize the problems I've described above, and that the resulting uneasiness is part of the reason some of us believe risk is so problematic. My own experience has been that having a clear and logical taxonomy of the elements that drive risk makes it a much easier (but still not trivial) problem to tackle, and allows us to do a much better job of helping management understand where infosec risk issues stand relative to the other things on their plates. And with that clearer understanding, I find that management is more receptive to our conclusions and recommendations. It also significantly reduces the communications challenges amongst those who are operating from the same set of definitions.

## Getting Loss Right

(<http://riskmanagementinsight.com/riskanalysis/?p=817>)

One of the most common questions I'm asked is, "*Where do the monetary values come from that are used in FAIR?*" It's an excellent question, and not a surprising one given that there's been very little credibility in the values our profession has used historically (when any are used). Before discussing where FAIR's loss values come from let's consider why our profession has struggled with this issue in the past.

I believe there are two key reasons why monetary loss estimates have been viewed with so much skepticism:

- Estimates have tended to focus on worst-case outcomes without recognition that most events don't come close to realizing a worst-case result
- The estimates usually come from information security professionals, and frankly, we're about the last people on earth who should be making those estimates

You might be asking yourself why "sparse empirical data" isn't included in the list. Well, yes, that's been an issue too, especially when so much of the information that is publicly available comes from extreme events. In the grand scheme of things though, I believe the two issues listed above are more fundamental.

### Chicken Little isn't welcome here

One of my mottos in life is to learn from other people's mistakes whenever possible, and I had just such an opportunity while I was the CISO at a large financial institution. We had engaged a consulting firm to perform a BIA, and they were presenting their results to the CEO and his cabinet. Now, these consultants weren't rookies, but they were ripped apart by the CEO when they presented their report. Why? Because their loss analyses were superficial and focused solely on worst-case outcomes. Amongst the withering feedback he provided them with were questions about how the heck he was supposed to use this information to make business decisions.

When all we provide are worst-case loss statements to management we've almost certainly provided an inflated view of risk. When we can provide them with information about how likely a worst-case outcome is and what more probable outcomes look like, that's when we're providing value. That's when they're able to make risk-aware business decisions.

### Most of us are mostly clueless

There's not a lot that needs to be said here. It's simple -- the vast majority of information security professionals are NOT qualified to estimate most business loss. We don't have the necessary visibility into key legal, market, revenue, and other business considerations. The good (?) news is that most of us recognize we're clueless, so we either ignore loss magnitude, we focus on worst-case, and/or we just slap a red/yellow/green label on it and hope nobody pushes back on why we labeled it that way.

I recently got a personal reminder about my own cluelessness while performing a risk analysis for a client. This particular client provides (amongst other things) back-end processing for many other companies, so one would think that availability outages would pose significant legal and reputation liability. (Lost revenue wasn't an issue because they had a highly captive user base -- i.e., revenue would be delayed but not lost.)

To be FAIR about it

My challenge was that I couldn't seem to get time on the legal person's calendar to obtain specific information about contractual liability tied to their SLA's. No problem, I thought. I'll just input some estimates of my own based on my understanding of their business model and my many years of experience so that I could continue with other parts of the analysis. (And no, I wasn't going to use the results in the report without first vetting the data with the lawyers.)

As one might expect, the results of my analysis reflected significant legal/contractual exposure associated with availability outages. However, when I finally met with the lawyer I learned that:

- In most cases, their service agreements included no availability guarantees
- Where availability guarantees did exist, it would take armageddon to miss them, and
- Even then, they were only liable for actual damages

Hadn't expected that. So much for my years of experience...

### To be FAIR about it

Now, back to the original question. Done properly, loss estimates used in FAIR come from appropriate business subject matter experts within the organization (e.g., legal, marketing, etc.) and, where available, from credible industry sources. That probably sounds logical, but then the question comes up, "*Am I going to have to speak to one or more business colleagues every time I do a risk analysis?*" If true, then we could probably expect to get a small handful of analyses done per year because they'd get tired of talking to us.

Fortunately, most of the loss data is highly reusable. For example, you may perform many risk analyses associated with potential credit card number compromise, and the loss data from one should be transferrable to others. Consequently, you can have the conversation with your business colleagues once about this type of event and then only go back to them periodically to ensure that some material change in the loss landscape hasn't occurred.

This then begs the question of size. Odds are, not all credit card scenarios you'll analyze involve the same number of credit cards, so what then? Do you just extrapolate up or down from whatever numbers you got for the first credit card scenario you performed? Nope, because loss magnitude doesn't change linearly based on the size of the event.

### Loss Tables

In order to leverage loss data that is reusable across scenarios of different sizes, I use what I refer to as Loss Tables. These tables describe loss magnitude ranges for scenarios of different sizes that can be used as a reference during analysis (an example of a table related to fines/judgments from credit card compromise is shown below).

To be FAIR about it

Number of CC's	Minimum	Most likely	Maximum	Rationale
1	\$0	\$0	\$2,500	Maximum value reflects the remote potential for some private action.
10	\$0	\$0	\$5,000	Same as above
100	\$0	\$0	\$10,000	Same as above
1k	\$0	\$0	\$20,000	Same as above
10k	\$0	\$25,000	\$100,000	Most Likely value reflects relatively minor actions by states attorney generals and/or small fines or settlements with the credit card companies and affected financial institutions.
100k	\$25,000	\$250,000	\$1,000,000	Minimum value represents an assumption that some fines/judgments/settlements are inevitable. Maximum value represents significant state attorney general actions, as well as an increased potential for VISA and/or Mastercard fines.
1M	\$250,000	\$2,500,000	\$10,000,000	Same as above

You only have to sit down with your business colleagues once to populate your tables and document the rationale behind them, and then you should be well-prepared to analyze most of the scenarios you'll face.

Depending on what your business is, you may want to create different loss tables for different type events (e.g., PCI, HIPAA, availability, etc.) and different categories of loss (e.g., response costs, fines and judgments, etc.).

## Benefits

The most obvious benefits to this approach are that the estimates should be much more accurate and will automatically have a much higher level of buy-in and credibility with your stakeholders. You'll also have a common reference others in the organization can use to perform their analyses, which will improve consistency. A less obvious benefit is that it is remarkable how much you can learn about the business (and risk) by having these focused conversations with your business colleagues.

If you've been through [FAIR training](#) and aren't using loss tables, please shoot me an e-mail and we can discuss how to get you started.

## More Than Just Numbers

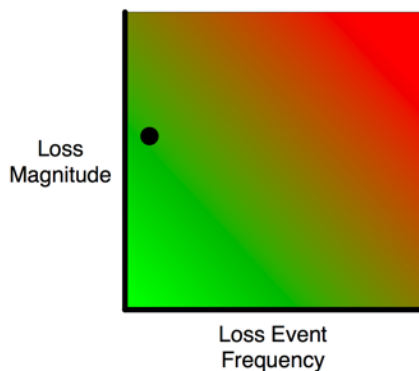
(<http://riskmanagementinsight.com/riskanalysis/?p=861>)

Many people believe that FAIR focuses strictly on quantitative risk statements, but they couldn't be further from the truth. The numbers simply allow us to recognize conditions and convey information better than we could do in any other way. Sometimes, however, numbers don't tell the whole story.

In this post I'll describe two conditions defined within the FAIR framework that help us to ensure management understands the nature of some risk scenarios that would be very difficult to describe quantitatively or qualitatively.

### Fragile conditions

Suppose we have a scenario where the threat landscape is very active but, due to a single extremely effective control, we actually have a very low probability of loss. If we were to plot this condition as a point on an X-Y chart, it might look something like this:



Now, if all we provided management was this point on a chart, there's a decent chance they'd be fine with it. After all, the frequency is low and the magnitude isn't outlandish. What isn't conveyed in the chart however, is the fact that if the single control fails, the point moves rapidly to the right -- i.e., there is no "grace period" or window of time in which we might avoid compromise. The threat event frequency is just too high.

In order for a decision-maker to make a well-informed decision about how to manage the risk scenario, they need to understand both the amount of current risk as well as the implications associated with the condition's fragile nature. With this information they may decide to introduce another layer of protection (defense-in-depth) and/or apply measures that make the control more robust and less likely to fail. Or, of course, they may decide to do nothing, but at least it would be an informed choice.

### Unstable conditions

Another scenario can exist where threat activity is inherently low but we have few or no resistive measures in place -- i.e., our vulnerability is high. Here again, the point on an X-Y chart would look just like the fragile condition above, and management might not be too concerned. What the numbers don't

To be FAIR about it

tell us though, is that we're essentially rolling the dice every day and counting on bad things not happening. We aren't actively managing the situation.

Here again, by letting management know about the unstable nature of the scenario, they're able to make an informed decision about their control options.

Another important aspect of unstable conditions is that in some cases the lack of preventative controls may be construed as an absence of due diligence by external stakeholders -- particularly if something bad happens.

### Why it matters

Many of us would intuitively recognize the nature of these conditions when evaluating a scenario, so you may be asking what the big deal is about formalizing their definition. Well, because it's difficult to convey these conditions quantitatively or qualitatively, what tends to happen is that people "adjust" the assigned risk level for scenarios like these so they'll land in the high-risk category -- essentially equating these scenarios to scenarios where the loss event frequency is actually high. Unfortunately, in doing so they misinform their decision-makers. The fact is, these conditions are importantly different from scenarios where the frequency/likelihood of loss are high, and management needs to recognize this difference and decide accordingly.

## It's still a choice

(<http://riskmanagementinsight.com/riskanalysis/?p=899>)

This post is prompted by an “enthusiastic debate” about regulatory compliance I had recently with another gentleman in our profession.

I'd love to take a poll of infosec professionals to find out how many of them adhere strictly to speed and other traffic laws when they drive. Why? Because many of these are the same people who state with conviction that, when a law/regulation exists regarding information protection, an organization **MUST** comply. While we might wish that were true, the fact is that compliance is **ALWAYS** a choice. It's just another risk decision; usually a trade-off of some sort. Does the organization prefer to accept the risk associated with potentially being caught and facing legal and other losses, or would they prefer to accept the costs and business impact associated with complying.

The other consideration in play is the fact that many laws are open to interpretation. I've been in plenty of meetings where the ambiguity in law is leveraged in decision-making. Not in a malicious, bwah-ha-ha sort of way, but in a legitimate “How do we best manage the cost and risk associated with running a business?” sort of way. And for those who'd argue that's a terrible thing, I'd bet a close look at some of your own decisions will find a little “harmless interpretation” of the law from time to time.

Of course, some people might argue that you can't compare speeding, tail-gating, and rolling through stop signs with the damage that can occur from a breach of credit cards or other PII. I beg to disagree. I believe the risk associated with automobile accidents resulting from even relatively simple carelessness or thoughtlessness is significant.

The point is, when we adopt the premise that laws/regulations somehow eliminate choice and decision-making, we're being naive, and this naiveté comes across pretty glaringly to many of the business professionals we serve and support. It's just another example to them of the infosec geek lacking perspective and viewing our very grey world in black-and-white terms.

## CVSS Review

(<http://riskmanagementinsight.com/riskanalysis/?p=909>)

I recently had the privilege of being a guest on the Securabits podcast and, during the session, was asked about other frameworks. I mentioned CVSS (Common Vulnerability Scoring System) in my answer and said I thought it had some serious problems as an analysis and measurement tool (however I also said there were good things about it). Given time constraints, I didn't go into detail in the podcast about what I thought was good or less-good about CVSS. That's what this post is about -- to clarify and share my thoughts regarding CVSS (version 2.0).

In the interest of keeping this post to a manageable length I'll constrain my observations to what I believe are the most important strengths and weaknesses of CVSS.

First, I have to acknowledge that what NIST and CMU have tried to accomplish with CVSS is both admirable and difficult. I can only imagine the debates that must have taken place during its development regarding tradeoffs that needed to be made in order to come up with a practical result. I also believe there's value in CVSS, even as it is today. That said, like any other model or framework there's always room for improvement. More importantly, like any other tool, its limitations should be well understood so that decisions based on it are made with both eyes open.

### What CVSS aims to be

The CVSS guide mentions three key benefits the framework is intended to provide:

- **Standardized vulnerability scoring** -- essentially, a common means of measuring "vulnerabilities". I think the framework accomplishes this objective for technical vulnerabilities because it does, in fact, provide a standard against which technical vulnerabilities can be scored. Enough said.
- **An open framework** -- i.e., a framework where scoring includes rationale so that the results don't have to be accepted on blind faith. As described further on, I think the framework hits this target in some respects, and misses completely in others.
- **Risk prioritization** -- i.e., a means of understanding the significance of vulnerabilities so that they can be compared and, thus, prioritized. Here again, in some limited respect CVSS accomplishes this objective. Overall though, as a CISO or other decision-maker, CVSS would not provide me with the information I need to make well-informed risk decisions.

### An open framework

Great idea -- a framework where justification is provided for the scores/measurements being used. And for the variables a user makes choices about within CVSS (e.g., Exploitability) there is some basic descriptive rationale in the selection matrix. Unfortunately, CVSS equations are also chock-full of weighted values, none of which appear to have clearly documented basis.

For example, the Base Equation multiplies Impact by 0.6 and Exploitability by 0.4. In other words, someone decided that Impact was always 20% more important than Exploitability. What's the rationale for that? In fact, by my count there are five weighted constants in the base equation alone. Six more weighted values (eleven total) if you include the fact that each Base metric will be given a value that appears to be arbitrarily assigned (e.g., For Confidentiality Impact the score will be 0.0, 0.275, or 0.660 depending on whether the vulnerability is assigned "None", "Partial", or "Complete" for that metric).

The other CVSS equations use weighted values in a similar fashion. Perhaps there are well-documented and thought-through rationale for each of these, but I haven't found them.

In my experience weighted values are rarely well-justified. Furthermore, they tend to be very sensitive to specific conditions/assumptions. For example, someone might argue that strong authentication is a more important control than logging. After all, "an ounce of prevention..." Consequently, it might be tempting to "weight" authentication's value higher than logging. Unfortunately, the logic breaks down if the scenario is focused on privileged insiders as the threat community -- i.e., people who are supposed to have access. In that scenario strong authentication isn't a relevant control at all and logging is much more important.

Unless there's good rationale for weighted values, they introduce ambiguity, limit the scope of where the analysis can be applied, and can in some cases completely invalidate results. At the very least, if weighted values are going to be used, some well-reasoned rationale should be provided so that users can make an informed choice about whether they agree with the weighted values.

## Effective risk prioritization

As a decision-maker, two of the fundamental inputs to any decision are "*What's the likelihood/frequency of bad things happening?*" and "*How bad are they likely to be if they do happen?*". These are the two values that, taken together, provide me with the loss exposure information I need in order to prioritize effectively. So, in order for CVSS to be an effective aid in risk-informed prioritization it has to provide useful information on both of those parameters.

CVSS tries to hit both targets, but falls short. With regard to frequency/probability of loss, CVSS focuses on the likelihood of attacker success from a couple of different angles, but never addresses the frequency/likelihood of an attack occurring in the first place. Without that metric, the likelihood of attacker success simply does not provide enough information for me to understand the frequency/likelihood of loss. CVSS may be trying to address the likelihood of attack through its Access Vector metric which, it could be argued, implies that the farther away an attacker is from the target, the less likely an attack might be. No argument with the logic (if that is in fact what the metric is supposed to represent), but there are a lot of assumptions built into that, including an assumption that the attacker isn't an insider.

From a loss magnitude perspective, the Base Metrics include Confidentiality, Integrity, and Availability references but these are actually measuring something pretty different. In a longer post at a later date I might describe a way in which these CVSS metrics could be used in a very interesting way, but that would make this post WAY too long.

CVSS's Environment Metrics try to include additional loss magnitude considerations. Besides being very qualitative, there appear to be some significant logic flaws in the approach. For example, the Target Distribution metric is essentially a measure of "surface area" (i.e., how many systems could be affected). One problem with this is that there are many scenarios where a single critical or highly sensitive system/asset is exposed (i.e., a small Target Distribution) but gross exposure exists. The way CVSS math works, this exposure would be unaccounted for. Something else to keep in mind is that Target Distribution is also a key consideration in loss event frequency (it may be even more important there in many respects), which isn't accounted for at all in CVSS.

Setting aside the points above, prioritization of CVSS ratings against anything outside of CVSS isn't practical because CVSS uses an ordinal scale. You can't usefully compare something that was measured on a 1-to-10 ordinal scale against something that was measured in monetary values or, for that matter, in a different 1-to-10 scale.

## Math

I've blogged before about the problems associated with using math on ordinal scales, so I won't belabor the point here. Suffice it to say that it just doesn't stand up to scrutiny. That said, if the user recognizes that the results are pretty much meaningless for anything but comparing one CVSS value against another, then I guess no harm, no foul.

## Bottom line

For all I know, the people who put CVSS together already thought through all of this (and the other problems within CVSS that I haven't talked about here) and decided that what they came up with was the only practical result given the constraints they faced and their objectives. Nothing wrong with that. Trade-offs are inevitable. It is important though, for users of the tool to have a realistic and accurate understanding of its capabilities and limitations.

CVSS seems like a decent way to measure and compare technical deficiencies ("vulnerabilities") against one another from a "*Very roughly) how much weakness does each vulnerability introduce relative to all of the other vulnerabilities measured using CVSS?*" perspective, which can be useful information. What it doesn't provide is meaningful information about how these vulnerabilities stack up in the bigger picture -- i.e., "*How important are these vulnerabilities relative to the other concerns I have to consider spending resources against?*" In other words -- "*How much do/should I care about the findings?*" In order to be useful in answering these questions, CVSS would have to evolve considerably.

Speaking of evolution... RMI has on the drawing board a potential alternative to CVSS that we believe will be both practical and more effective in characterizing the risk associated with vulnerabilities. Stay tuned!