Beyond Luck & Guesses:
Overcoming the High Cost of Worthless Op Risk Models

By Dennis Wenk,
Certified Organizational Resilience Executive, Symantec

Although risk assessments have been widely used for a number of years, many people fail to grasp the essential character of risk models relative to IT risks, and persist in the view that without "precise" data, model-based risk analyses are worthless. The term "precise" is used to suggest that there *should be* aspects of future risk events that could be known with accuracy in the present. Since mere mortals cannot *know* the course of future events with precision, some in IT reason that *nothing* can be deduced *in the present* about future risk-related losses. In fact, there are things that can be done today to keep the odds in our favor and a loss-expectancy risk model that economically quantifies operational risk not only will identify the serious risks but will provide the important cause-and-effect correlation needed to rationally evaluate risk-reduction tradeoffs.

Still others believe that it is '*Good Enough'* to evaluate operational-risks using outdated methods like the "Business Impact Analysis" or some sort of High-Medium-Low risk-scale. In many cases, these methods are not rich, statistically grounded, qualitative assessments but based on unreliable intuition about the likely facts and a close examination reveals inherent flaws in both approaches. Not only must the serious risks be identified from an ever increasing number of risks, but there are also a wide variety of risk-reduction choices with each choice presenting varying degrees of features, functions, and costs. The real problem for `traditional' approaches, like the BIA, is not that they are wrong, but that they offer no guidance on how to optimally improve the situation, do not account for existing mitigation actions and therefore, *many financial organizations spend a disproportionate amount of time and money on options without understanding the intrinsic IT-infrastructure risk or the economic returns back to the organization.*

By far the prevailing 'best practice' in IT today is to manage risk by instinct, heuristic and reaction, rather than by facts and proactive risk mitigation. Current 'Best Practices" have had a long running bias that favors a '***Better Safe than Sorry***" approach to dealing with risk. This approach enjoys widespread support particularly within the Business Continuity community and this bias has lead to an emphasis on process improvements such as contingency planning, crisis management, and other 'preparedness' activities that only address disastrous events.

The unintended consequence of this precautionary attitude is that operational aspects of IT have been systematically neglected: ***This might be the biggest blunder in business today***. While organizations directed their attention to preparedness and worse-case scenarios, a host of new risks were being introduced through the complexities of IT and its fragile infrastructure. Over time, as financial organizations became highly dependent on IT, they have also unknowingly accepted these new risks[1]. In today's technologically-dependent financial organizations, rather than simply focusing on planning/recovery activities to address catastrophic-events, best practices objectives should be directed toward minimizing the expected '*cost of economic-loss*' by means of implementing the optimal combination of risk-reduction measures. The likelihood that a financial organization will experience a catastrophic loss from a service interruption due to an IT problem is far greater than any service interruption being caused by some disaster or 'black swan' event.

[1] Power, Michael. "The risk management of everything." The Journal of Risk Finance 5 (2004): 58-65.

Introduction

Webster's defines *model* in the sense we are using it here as "A tentative description of a theory or system that accounts for all of its known properties." The techniques used to model risk are sometimes implemented as computer programs that embody a group of equations that describe the relationships between risk parameters, facilities to enter the input data into the equations, and other facilities to display in numerical and graphic form the results generated by the equations.

This suggests that there are two parts to constructing a model of risk:

(a) identification of the parameters that define risk losses, and
(b) Relationships between these parameters which determine the magnitude of losses.

In the discussion that follows we will focus exclusively on the modeling of what is commonly referred to in the world of business as *operational risks*. Operational risks are those risks associated with the operations of an organization. The Basel II Accord for International Banking defines Operational Risk as the, "Risk of loss from inadequate or failed internal; processes, people, and systems or external events". When Processes, People or Systems fail, whether it is from internal or external events, the losses can be substantial. Unlike credit and investment risks, operational risks are assumed to have only a negative (loss) outcome, alternately stated, there is no upside gain, we can only limit losses.

The association between operational risk and IT-infrastructure risk is that we live in a technology driven world. All possible business processes have been automated; automated to the point where Information Technology is deeply embedded in the operating fabric of the financial organization. The modern financial organization is now highly dependent on information technology. Simultaneously, and quite unintentionally, information technology has introduced new exposures which have deceptively seeped into every layer of the financial organization. These exposures have been created through a combination of the technologies complexity, its enmeshed interdependence, and the brittle nature of this infrastructure.

Since investing to reduce operational risks offers no expected-gain, people routinely make the erroneous assumption that there can be no quantifiable benefit, and therefore no return on investment (ROI). The first thing to recognize is that investing to reduce operational risk for IT infrastructure is a very different kind of investment decision. The vast majority of business investments expect a return; a gain. Capital is wagered with the expectation of a payoff, something greater than the original investment. The expected-gain is what makes taking the risk worthwhile. When investing to reduce an operational risk, however, there is no obvious payoff or gain; the best that can be expected is to prevent something bad from occurring, to avoid a loss that might be reasonably expected or an ***EXPECTED LOSS.***

Expected loss is not something subjective and abstract like the terms 'risk aversion', 'risk appetite' or 'risk tolerance' and it is not High-Medium-Low risk mapping. It is a familiar mathematical approach for quantifying risk using the expected value theory originally developed by Blaise Pascal and Pierre de Fermat. We can apply expected loss to a real business problem; the inherent risk of the IT-infrastructure. Expected-loss parameters include; numerical scales for risk probabilities, economic loss exposures, and vulnerability factors. Specifically, the number of occurrences per year of each operational risk will need to be determined or estimated, the potential for economic loss of each of the business processes per risk impact will need to be calculated or estimated, and the related vulnerability, zero to 100%, of each process to each risk calculated.

The "Big Question"

Traditionally risk management has been associated with the effects of catastrophic events such as, floods, fires, hurricanes, tornados, earthquakes, fires, terrorism and pandemics, and sometimes even the threat of a meteor strike is included. The question is not whether these events should be of interest to a business; the possibility of excessive losses to life and resources are quite obvious and it is self-intuitive that being unprepared for a catastrophe could cripple a business. The value, however, of operational risk management lies not in identifying risks and exposures; the value lies in determining the optimal ***'investment'*** to mitigate the most serious risks. In a technologically dependent world there are many potential loss-events and they occur with great frequency. What was once just a minor annoyance, like a software-bug, can now generate an economic loss similar to that of a flood or a fire (i.e. Knight Trading where a software bug caused a $440 million loss in about 30 minutes, that loss is three times the company's annual earnings.[2]). Just knowing that a large risk exists does not mean that all the tradeoffs of addressing the risk are also well known. As Robert Halm points out 'This leads to a paradox that is becoming increasingly recognized. Large amounts of resources are often devoted to slight or speculative dangers while substantial and well-documented dangers remain unaddressed'[3].

An important part of managing operational risk is the proper allocation and alignment of scarce-resources to the proper mitigation actions. The "Big Question" is how to *optimize* scarce resources *today*, to achieve the greatest reduction in *future losses*. The two components of that Big Question are: which risks are the important ones and what are the optimal risk-reduction actions. The traditional Business Impact Analysis (BIA) and qualitative High-Medium-Low Risk analysis offer little guidance in answering the Big Question.

In general, people are reluctant to trade REAL funds today for uncertain future events. The need to answer the Big Question arises from the fact that the available resources are not unlimited but scarce. This is why it is irrational to strive for zero losses, which has the goal of avoiding ALL future losses. Achieving absolutely zero losses would be infinitely expensive.

Thus, the dilemma faced by management is that IT systems are exposed to a wide range of possible 'future' risk-events, the occurrences of which will cause losses, large loses. The greater the reliance on technology, the larger the potential for loss becomes. In any highly-automated organization the potential for loss is due to the fact that IT can lose more transactions, of significantly greater value, at an unprecedented speed. A 2010 study by the Ponemon Institute estimates a whopping 2.84 million hours of annual data center downtime. Ponemon estimates that the average hour of outage costs about $300K per hour of outage; that translates into a total loss of $426 billion a year[4].

Given the "un-know-ability" of future risk events, how does management *optimize* the allocation of the limited available resources? Perfect foresight would make things simpler. For example, if we were, somehow, gifted enough to "***know***" that the next loss would come from a data-theft attack occurring next Wednesday at 9:00 AM then, armed with this prophecy, we could deploy every effort *now* to block the attack. This, however, is more in the domain of oracles and soothsayers

---

2 Hour, In Less than an. "Is Knight's $440 million glitch the costliest computer bug ever?" CNNMoney. 09 Aug. 2012. Cable News Network. 25 Jan. 2013 <http://money.cnn.com/2012/08/09/technology/knight-expensive-computer-bug/index.html>.

[3] Making sense of Risk: An Agenda for Congress, in Risks, Benefits, and Lives Saved 183, Robert Hahn, ed. (New York: Oxford University Press, 1996).

[4] Power'..., Emerson Network. "Calculating the Cost of Data Center Outages." Scribd. 21 Jan. 2013 <http://www.scribd.com/doc/56246462/Calculating-the-Cost-of-Data-Center-Outages>.

than that of rational managers. Unfortunately we can never "***know***" what will happen next. How, then, can we be most effective at answering the Big Question? The Big Question *must* be answered; that is to say successful financial organizations *must* make effective choices about risk mitigation *today* in anticipation of *future events*.

As previously noted, there is a tremendous, $426 billion opportunity for improving risk mitigation within the IT-infrastructure. Presumably, the objective is to come up with the optimum answer for the Big Question based on what is known today about the likelihood of future risk events. How best to keep the odds in our favor. With respect to this uncertainty of future events, it is interesting to note that future estimates are essential to every cost/benefit analysis, such as expectation of future profitability. These elements are also based on assumptions concerning uncertain future events. While events that occur in the future can only be estimated; they are neither certain nor are they known with absolute precision. While there are no "knowable" facts about the future, there are informed estimates that can be based on past experience. Our understanding of how present conditions and trends will affect future risk events can be reasonably calculated. Consequently, we must make informed estimates about future losses and then take investment action based on those informed estimates.

Business Impact Analysis and Risk Assessment

The Business Impact Analysis (BIA) was fundamentally designed to address catastrophic events. Its purpose is to determine how losses accumulate over time from a worse-case event. The BIA also provides the important details needed to respond to worse-case circumstances, details such as critical dependencies and resource requirements. These details are quite complicated and fundamental to responding to 'smoking-hole' IT disasters. Whether the destruction is caused by a fire or a meteor falling from the sky, knowing 'WHAT' to do when a big, adverse event happens is extremely important. Survival depends primarily on successfully reacting and responding to the perilous circumstances; therefore the clear solution is to identify the key dependencies and create recovery-contingencies for the most large catastrophic risks. Case in point to the dysfunctional nature of the BIA is Mr. Roger Sessions estimates that "Worldwide, the annual cost of IT failure is around USD 6.18 trillion (that's over USD 500 billion per month)". [5]

The BIA, however, has a significant limitation; it systematically ignores probability. Its focus is on the 'potential-for-loss' of a worse-case scenario, in other words for IT the cost-of-downtime. Loss potential is a good indicator of the magnitude of a worse-case event, but the focus on large loss-potentials can cause excessive worry over statistically small risks. This probability-neglect distorts perceptions about many serious threats to an organization, creating a bias toward the large events with large loss potentials. By simply ignoring probabilities the BIA makes the 'smaller' threats seem invisible, and what is out of sight is effectively out of mind, hence never addressed. This omission disregards the trickle-down effect that a minor disruption can have on a financial organization In today's technologically-dependent world, this lapse of judgment could destroy a financial organization just as easily as any catastrophe.

Two Israeli psychologists, Daniel Kahneman and Amos Tversky[6] conducted extensive research regarding how people manage risk and uncertainty. Their work was awarded the 2002 Nobel-prize in economics in what is now called 'Prospect Theory'. The ground-breaking research showed that intuitions about risks are highly unreliable and can produce severe and systematic errors by reinforcing the tendency people have to fear things that are not dangerous and not fear

[5] "Simple Architectures for Complex Enterprises." : The IT Complexity Crisis: Danger and Opportunity. 26 Jan. 2013 <http://simplearchitectures.blogspot.com/2009/11/it-complexity-crisis-danger-and.html>.
[6] Kahneman, Daniel and Amos Tversky, 1979. "Prospect Theory: Analysis of Decisions under Risk." Econometrical, Vol. 47, No.2.

things that impose serious risks. The cost-of-downtime and the BIA neither help identify causes nor help prioritize preventative actions; these are left for intuitive judgment as the BIA universally overlooks the causal relationship of risk. The BIA provides little value for controlling operational risks because its primary purpose is to respond and recover, not prevent. The fundamental weakness of the BIA is that it was never intended to treat a cause or a symptom of an operational risk, because its objective is to produce contingencies for worse-case circumstances to ensure an (eventual) re-start of the organization. It is an after-the-fact approach, and not a preemptive, proactive approach to strengthening operations.

An important part of managing operational risk is how to optimize scarce resources today to achieve the greatest reduction in future losses. The BIA promotes a 'better-safe-than-sorry" presumption, focusing attention on threat-events in which a single occurrence is irreversible, and could cause catastrophic losses, such as pandemics or terrorism. There is widespread support for this precautionary approach and the BIA. In fact the BIA has been the fundamental 'BEST Practice" found universally in every business continuity standard to date. Yet just knowing the size of a worse-case scenario does not go very far in helping to determine how to properly allocate our scarce resources. In fact, the guidance a BIA provides misleading information relative to the serious operational-risks and resources are often misdirected toward the wrong threat(s).

It should not surprise anyone that even though a catastrophic event presents the possibility of Avery large loss, rational business managers are reluctant to invest funds to mitigate these risks because these types of risks are both rare and highly unlikely. Rational managers realize they must invest wisely and all the risks need to be evaluated. Certainly there would be no need to worry about a rare pandemic or an unlikely meteor if simply making the next payroll is questionable. Still, the key to survival is allocating the appropriate resources to the "right" risks; while that may include planning cost-effective contingencies for a worse-case scenario, to be effective management needs more guidance regarding the tradeoffs than the BIA is designed for or is capable of delivering.

There is a sense that IT-infrastructure risks can be prevented, avoided or, at a minimum, mitigated. While recovery is vital to managing the IT-infrastructure, most financial organizations would much rather prevent IT-infrastructure failures than recover from them (it is well known that the 'cost of recovery' can be more than twice the 'cost of prevention'). So recovery activities, while critical, add unnecessary costs to IT operations. Preventative actions would be more beneficial that recovery activities. Being effective at managing IT-infrastructure risks requires more than just being prepared for the unexpected, more than just planning to respond to a bad situation, and it requires more than simply quantifying the cost of downtime. What is missing is measuring the quantifiable benefit to justify the required investment.

To get a sense of the size of the risk, the BIA assumes that a disaster has occurred, and estimates the financial consequences for each of the affected business systems. The assumption is that these quantitative estimates can serve to identify where remediation is most needed, however while the worse-case economic consequences might be the same whether the loss comes from a fire or a tornado, the remediation actions are quite different. A BIA is based on the potential for loss for the worse-case scenario and does not include specific threats. Since there is no 'cause-and-effect' relationship in the BIA, it cannot be used to provide the necessary cost-benefit support for either mitigation actions or recovery efforts. Investing money to prevent a loss from a tornado would not be wise if there is a higher likelihood of loss from a fire, even though both threats have worse-case results.

There are a wide variety of risk-reduction alternatives relative to IT-infrastructure such as server clustering, remote replication, virtualization, storage arrays, converged networks, etc. Any of these solutions or combinations increase performance, availability or reduce some type of threat to the IT-infrastructure. Implementation of a high availability solution routinely requires an investment for additional redundant hardware, making the solution very expensive. While the

only rational reason to invest in a high availability solution is the expectation that its benefits outweigh the costs, however for the BIA there is nothing to measure because the BIA does not collect cause-&-effect information. There are all sorts of things that could go wrong in a technologically complex, fragile and uncertain world, without cost-benefit balancing, an organization would go broke trying to reduce every single risk. To be successful we must devote the right resources to the right risks

There are lots of tradeoffs and variability on the risk reduction-side of the equation, as well. Risk reduction solutions will have varying degrees of cost and will either address specific dimensions of the operational risk or provide more efficient ways to respond after the risk-event has occurred i.e., one solution might prevent a systems crash while other solutions provide faster recovery after the crash has occurred.

To overcome the BIA-deficiency the majority of IT standards, guidelines, and 'best practices' favor combining the Business Impact Analysis (BIA) with a Risk Assessment (RA). The product of a BIA is an estimate of the worst-case consequence of risks to a business system based on an evaluation of the loss that results from an IT service-interruption and/or damage to or loss of data.

The RA process is defined as a list of threats to the systems and a list of the mitigation measures that address each of the listed threats. Any mitigation measures that have already been implemented, are deleted from the list and the remaining mitigation measures become recommendations for implementation along with such abstract terms as ***Risk Aversion***, ***Risk Tolerance*** and ***Risk Appetite***, as if forward-looking choices about IT operational risk can be rationally made using these intuitive, subjective feelings.

The above is a brief description of the RA process as proposed by its adherents. When one looks more closely at the process, serious difficulties become apparent. The fundamental defect is the complete disconnect between the BIA and the RA. Notice that actions taken, or not taken, regarding mitigation, as prompted by the RA, have no effect on the results of the BIA. As a result, there is no connection between proposed business continuity plan and threat mitigation. Note, however, that if threat mitigation reduces the impact of a threat, the requirement to recover from the threat is reduced. A BIA does not track this because it ignores threat parameters. Likewise, recovery plans based on a BIA has no impact on the results of an RA. The other major difficulty is the failure to take into account implementation costs. The absence of a particular mitigation measure is not a sufficient justification for its implementation. In short, BIAs and RAs cannot possibly provide trustworthy answers to the Big Question.

So what are we left with? When one examines the BIA/RA process in action, it becomes clear that a consultant's recommendations are merely based on intuition and subjective judgment; what some might call "informed guesswork." The consultant asserts that the BIA and the RA show what must be done without any need for a cost-benefit justification. There is no reason to assume that the recommendations will be an optimum use of available resources.

High-Medium-Low Risk Analysis:

In an effort to link threats and potential loss, some authorities, particularly government bodies, have endorsed a risk assessment technique that evaluates risk event occurrence rates and the resulting impacts as High, Medium or Low. Presumably the intention is to avoid the perceived difficulties in making quantitative assessments of risk parameters. Once all the business system loss potential / threat pairs have been evaluated, one consults a chart like the chart illustrated in Figure1.

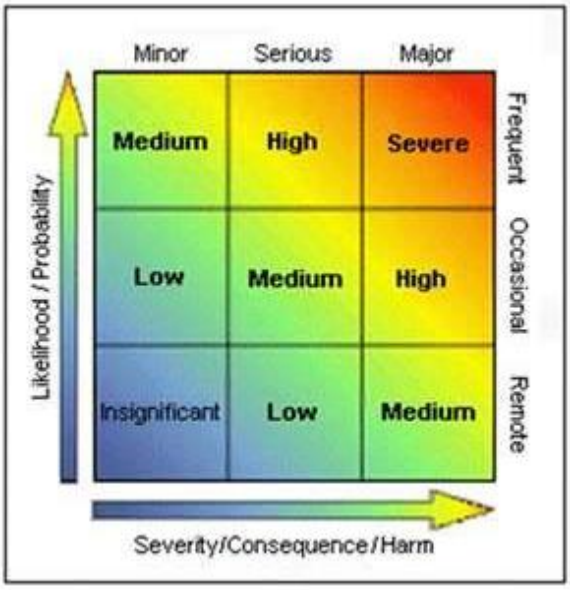

Figure 1: Examples of High-Medium-Low Risk Evaluation Charts

The evaluation chart in Figure 1 is a common way of attempting to prioritize the Threat-System pairs. Thus, if a threat occurrence rate is evaluated as High, and a system loss potential is evaluated as Medium, the priority action to mitigate the threat on the system is evaluated as High. Notice that the analyst is freed of the pesky task of developing quantitative estimates of the two quantities. What could be easier?

Theoretically, this technique would seem to be a reasonable way of prioritizing risk management actions. Unfortunately, there are three major problems. The first is that neither the cost nor the effectiveness of a proposed risk management measure is evaluated. The second problem is the three-step evaluation scale is completely inadequate to represent the real world. Finally, HML does not take into account the fact that a given threat will not necessarily have the same impact on all business systems. A single HML ranking cannot reflect this important factor.

Real world experience analyzing individual risk environments, using reasonable quantitative estimates, commonly yields risk occurrence rates that range over six or seven orders of magnitude during a typical risk analysis project. For example, a set of occurrence rates might be estimated to range from a hundred times a year (100/year) to once every ten-thousand years (1/10,000 years), covering a range of a million to one. In such a case, if we use a High-Medium-Low classification, the quantitative occurrence rates would translate into the 'HIGH' segment covering occurrences of 100 times per year to once-a-year, the 'MEDIUM' segment covering once-a-year to 1/100 years; and the 'LOW' segment covering 1/100 years to 1/10,000 years as shown in Figure 2.

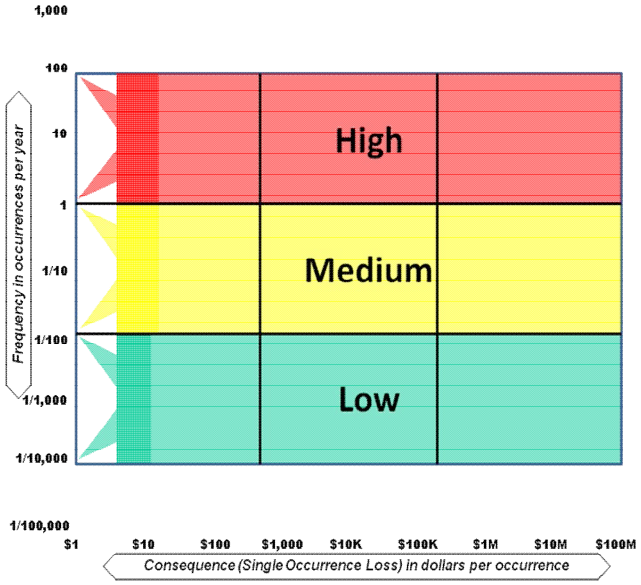

Figure 2: Range of Occurrence rates

Quantitative Loss Expectancy Model

The QLE model uses data collected about threats, functions/assets, and the vulnerabilities of the functions and assets to the threats to calculate the consequences. Consequences are the economic losses that result from occurrences of the threats that have been included in the model.

To date no better common denominator has been found for quantifying the impact of an adverse circumstance - whether the damage is actual or abstract, the victim a person, a piece of equipment or a business function - than monetary value. It is the recompense used by the courts to redress both physical damage and mental anguish. The economic impact is one which transfers directly to fiscal usage without any intermediate translation.

When a threat impacts a function or asset, it generates an economic loss, the consequence. The consequence is equal to the loss potential (loss potential is the worst case loss) of the function/asset multiplied by its vulnerability (0.0 to 1.0) to the threat. For instance, if a function or asset is subject to a seasonal threat that typically occurs 4 months out of a year, then the loss potential would be multiplied by .334. If, on the other hand the function or asset were subject to the threat throughout the entire year, then it would be multiplied by 1; or if some functions/assets were not subject to a threat then they would be multiplied by zero. The consequence for a threat is the aggregate of all of the individual consequences (single occurrence losses). The occurrence rate (frequency/probability/likelihood) and the consequence of a threat can be plotted as is shown in figure 3.

The vertical scale is the estimated occurrence rate of the threat, e.g., 10 times per year, once every five years, etc. The horizontal scale is the calculated consequence of an occurrence of the threat in monetary terms, e.g. $1.00 in losses if this occurs or $1M in losses if that occurs. Since the impact will be expressed in economic terms and fiscal matters are organized on an annual basis, a year is the most suitable time period to specify in expressing expected frequency of occurrence of threats. The resulting point represented by the red-dot on the plot in Figure 3 represents the threat event. Note that both scales are logarithmic.

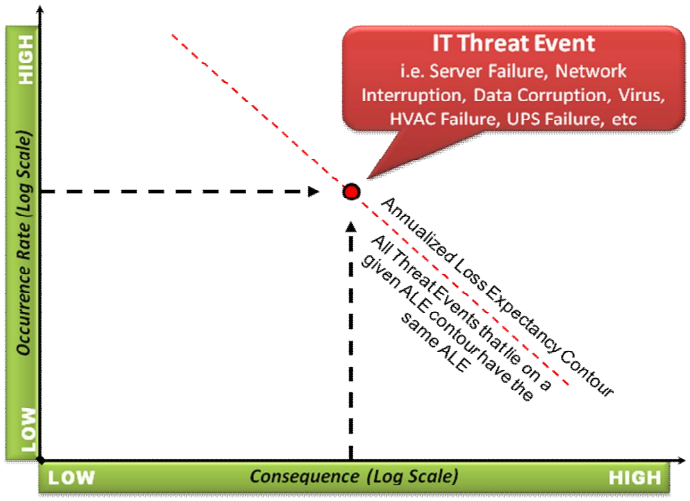

Figure 3: Threat occurrence rate/consequence & ALE Plots

A line has been added to the plot in Figure 3 to represent the Annualized Loss Expectancy (ALE) or Expected Loss of the threat event over the course of a year. For example, if the occurrence rate of a threat is 1/5 (once in five years), and the consequence (single occurrence loss) is $50,000, the threat's ALE is 1/5 X $50,000 or $10,000/year. Figure 3 shows the ALE line on which the threat event lies. Since both of the two scales are logarithmic, contours of constant ALE are straight lines. All threat events that lie on a given ALE contour have the same ALE. In other words, over the long term these threats will all trigger the same total losses. Expected Loss is what provides the economic cause-and-effect relationship that is deficient in the BIA which as a reminder only looks at Loss Potential, not ALE. Expected Loss ensures better priority setting of the serious risks and provides a business process to evaluate multiple alternatives using a standard ROI technique and cost/benefit analysis.

When all the threats included within the scope of a QLE (Quantitative Loss Expectancy) project are plotted, they will tend to lie on a band extending from the upper-left to the lower right like the example plot shown in Figure 4.

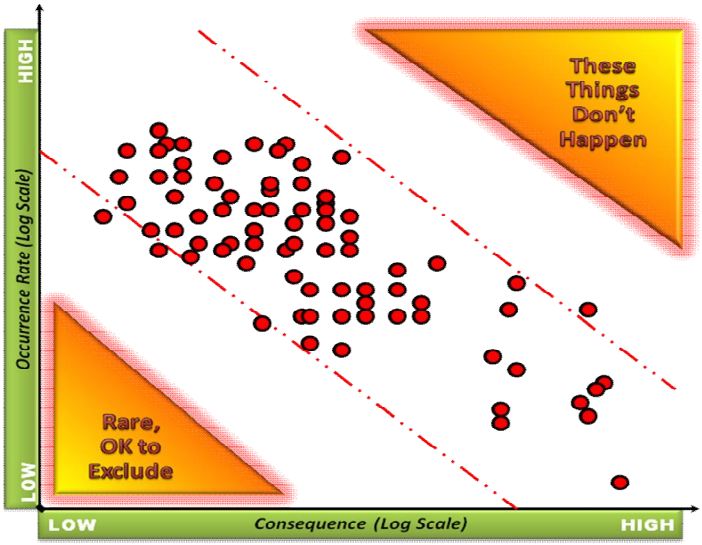

Figure 4: A Plot of a range of threats.

Threat events in the lower left corner are trivial losses, e.g., a $1 loss every million years, doesn't matter. Threat events won't occur in the upper right corner; this is the 'Doesn't Happen Zone' because these events just don't happen. If any organization experienced a $10-billion loss every day, it could not continue to exist. Said another way, life on Earth would be impossible if such large consequence events occurred all the time.

This threats-consequences plot can be used to classify the threat events in terms of the appropriate risk management strategy to apply to each of the threat events. This is the essence of the risk management process. It is possible to define an 'Irreversible or Catastrophic Loss' threshold, as illustrated in Figure 5. The value of the threshold may be based on a loss that would bankrupt a private sector organization, or generate an unacceptable drop in the share price, or for a government organization, the threshold may be set at the consequence that would require a supplemental appropriation. Presumably, if a threat lies to the right of this threshold, then action must be taken to either reduce its consequence, or its occurrence rate. Notice that an insurance policy that allows us to claim compensation when such a treat event occurs has the effect of reducing the economic consequence of the threat by "transferring" a financial portion of the risk to the insurer at the cost of the insurance policy's premium.

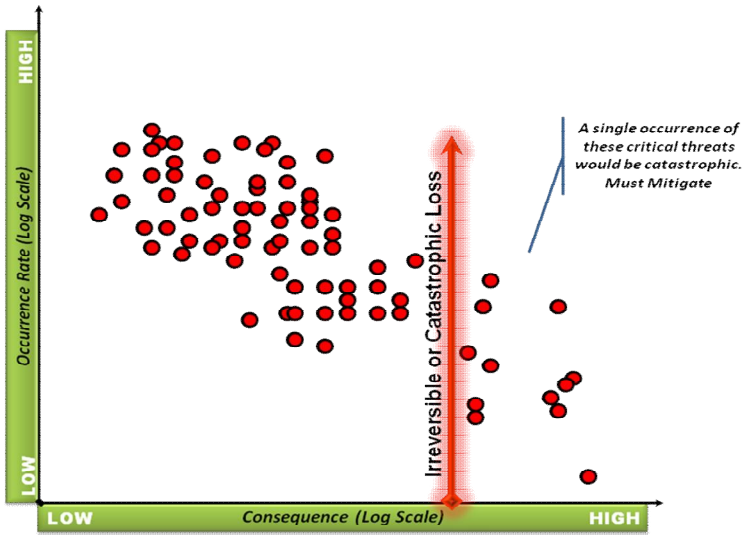

Figure 5: Irreversible or Catastrophic Loss Threshold

These large catastrophic events are so infrequent that it is often difficult to develop credible estimates of occurrence rate. Notice, however, that the QLE treatment of these threats is governed only by their consequences, which typically can be estimated with reasonable accuracy. Thus, the uncertainty of the occurrence rate estimate does not have a strong affect on the risk management decisions for threats in this zone. This is an important risk management concept that effectively answers concerns about the difficulty of estimating the occurrence rates of infrequent threats.

Applying Actual Loss Data to HLM Model

Using data extracted from an actual, real world QLE analysis, Figure 6 shows a computer-generated plot of IT threat-events depicting the results of the analysis using a realistic quantitative loss-expectancy risk model. The red dots represent about thirty threat events, plotted in terms of threat occurrence rate and their associated consequences. Although it may be impossible to know absolutely the loss-impact or the frequency of many threat-events, estimates *within an order of magnitude* are sufficiently accurate in most cases. The two scales are logarithmic so the contours in Figure 7 show constant straight dotted-lines, labeled in expected lost dollars per year or Annualized Loss Expectancy (ALE), which is the product of occurrence rate and consequence.

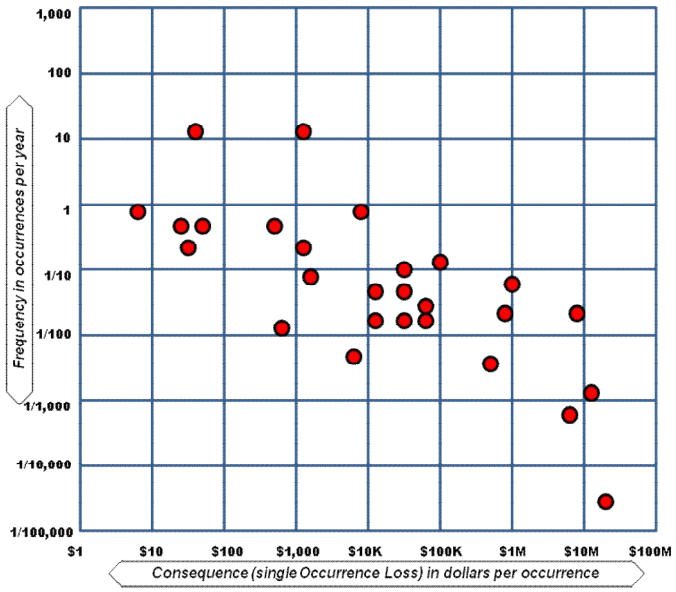

Figure 6: Plot of Quantitative Threat Analysis

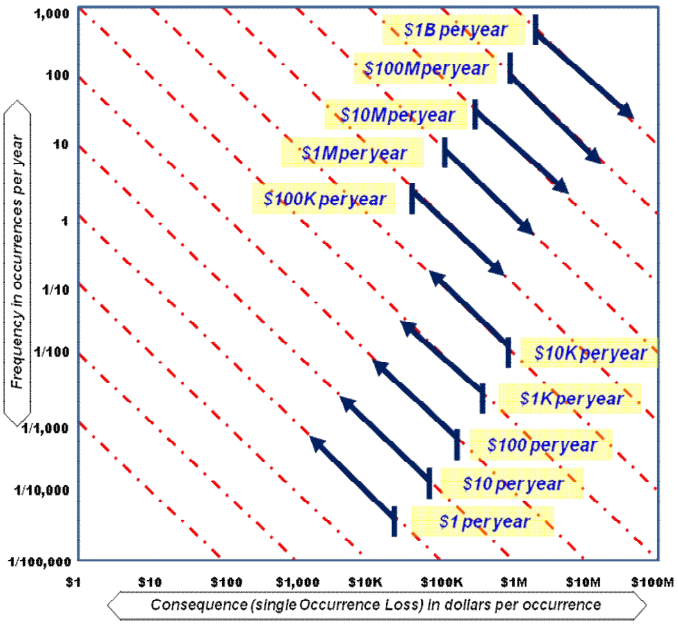

Figure 7: Annualized Loss Expectancy (ALE) Contours

To fully understand the limitation of the HML-model we inserted the very same threat-events in Figure 6 but plotted against a HML-Chart in Figure 8. It is obvious from this chart that HML is inadequate to satisfactorily represent the relatively simple range of threat values collected during an actual QLE analysis. While there are a large percentage of threats that have large ALE's, only a few are ranked as 'High' and were are no threats that ranked 'Severe'. The lack of any 'Severe' threats is not surprising since as was said previously "High Likelihood/High Consequence" events really don't occurrence, i.e. meteors don't crash into data centers on a daily basis, although I did review an article that had ranked "***IT-Failure***" in the "***High Probability/High Impact***" sector. I guess that CIO should consider a new career.

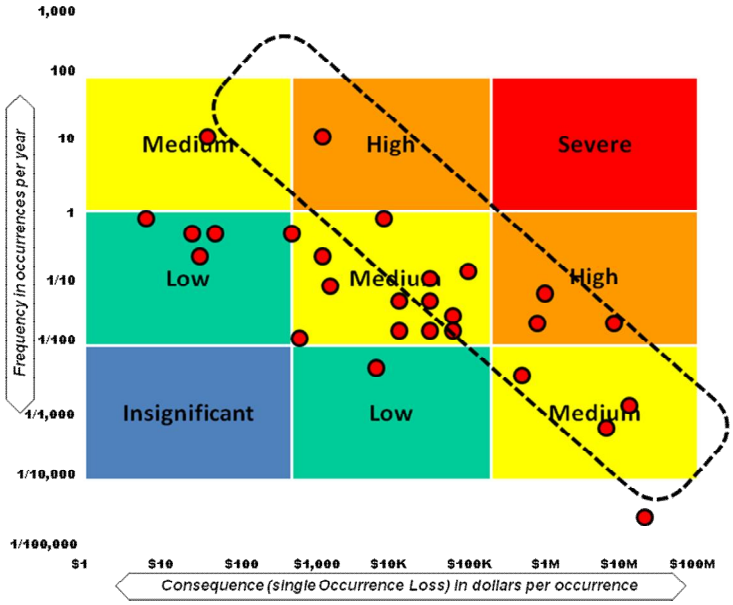

Figure 8: Plot of Threat events applied to the HML-chart

This raises two practical questions. If the quantitative estimate of an occurrence rate is estimated to be on or near the boundary between two of the classifications, for example, $1 per year using the example mapping above, should the risk be rated as Low or Medium? It seems likely that about 50% of a group of experienced risk analysts will answer Low, and the other 50% will answer Medium. This suggests that using an HML scale will result in uncertain prioritization of countermeasures for about half the risk events (those with one or both of the rankings near the boundaries) being considered. The white-arrows in Figure 9 identify the uncertain boundary threat-events.

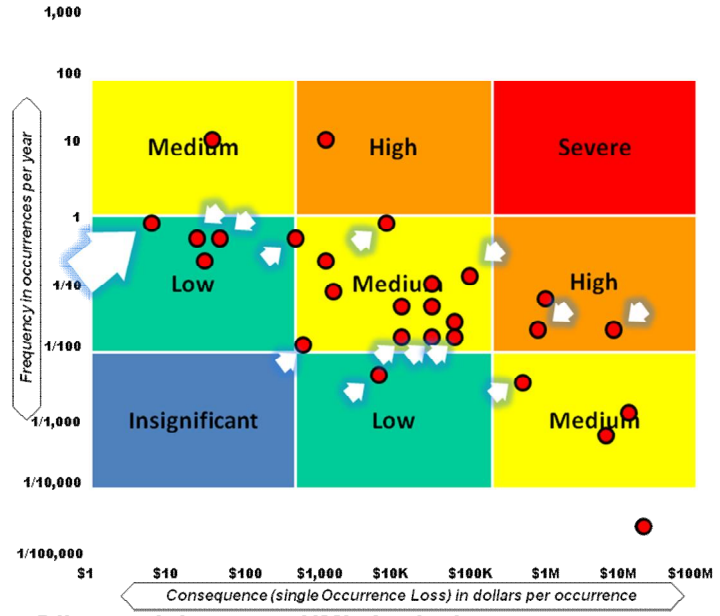

Figure 8: Boundary Dilemma Inherent to HML Analysis

Similar questions arise for two threats at opposite ends of a range. Consider this example: It seems reasonable to assume that we can, in most cases, make a credible estimate of an occurrence rate within an order of magnitude. For example, a quantitative estimate of 30 times

per year might be made for a risk event with an actual occurrence rate of 100 per year or 3 per year. Similarly, one might estimate an occurrence rate of 10 per year or Once-every-10 years for an event with an actual occurrence rate 1 time per year. Using an HML scale would result in both threats being evaluated as High, and therefore worthy of the same level of countermeasures, even though in about half such cases the occurrence rate of the two "High" threats would differ by from one to two orders of magnitude. See figure 10.

The other half of the risk loss equation is the potential for loss of the processes, data, and assets exposed to the threats. These too will have a similar wide range of values. Thus, for the same reasons given above, the application of an HML scale will result in unrealistic assessments. Since risk loss is related to the product of the two halves of the equation, we see that when an HML scale is applied to both factors we can have a situation where losses that differ by four orders of magnitude, $10,000 to $1, receive the same HML ranking. See figure 10.

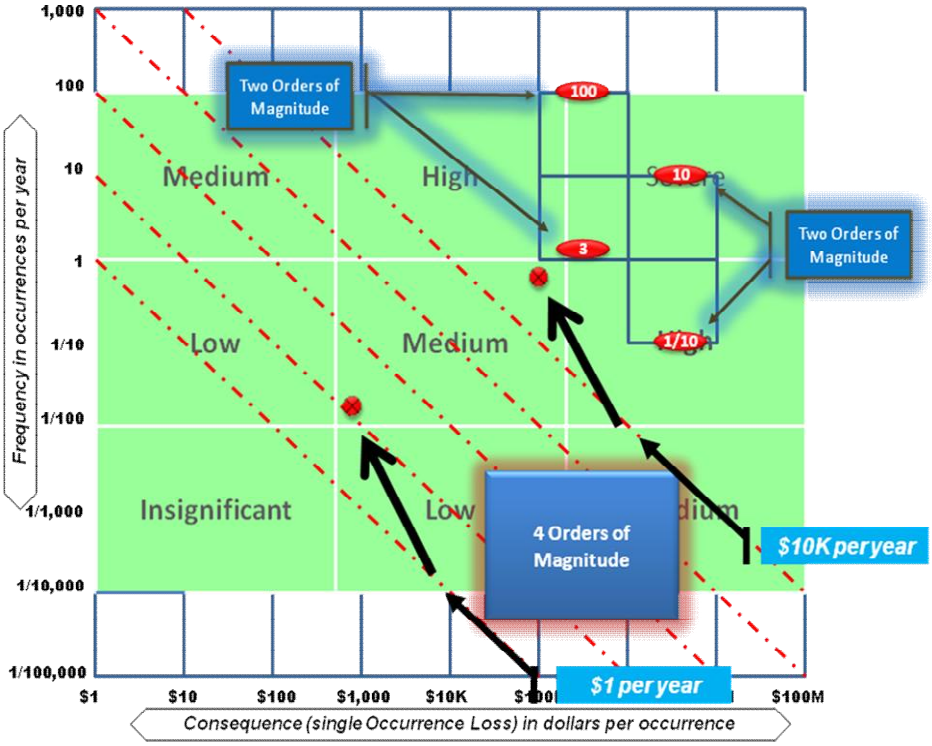

Figure 10: Possible Order of Magnitude Threat-Event Variances

Clearly, quantitatively based estimates will be a much more reliable and accurate basis for assessing risk. Critics of quantitative assessments often refer to them as subjective and therefore suspect. It is clear that HML estimates are much more subjective because of the difficulties cited above.

Evaluating Cost-Benefit:

To keep the odds in our favor we must economically-quantify the operational risks so that we can properly evaluate the many tradeoffs. The *only* rational reason to implement risk mitigation measures is to reduce future losses[7]. To properly evaluate a mitigation action's value we can use standard cost-benefit balancing. The cost-benefit of the proposed mitigation-solution is the expected-value of the '***reduction in future losses***' from the solution is compared to the cost to implement and maintain the solution. Bearing in mind our assumed order of a magnitude level of

[7] In some situations law or regulations may mandate certain risk management measures.

uncertainty in our quantitative estimates, how should we evaluate the following example results, quite typical of the real world?

- Mitigation/Solution #1: ROI = -95.0%
- Mitigation/Solution #2 ROI = 2.5%
- Mitigation/Solution #3 ROI = 3,500.0%

Mitigation/Solution #1 will still be a bad investment even if we have *understated* its loss reduction effect by an order of magnitude. Mitigation/Solution #2 appears to be marginal. If we have overestimated its benefit by a relatively small margin, for example 10%, well within our postulated error range, it will be a money loser. On the other hand Mitigation/Solution #3 appears to be a winner despite as much as a two-orders-of-magnitude *overestimation* of its effectiveness.

As with occurrence rate estimates, the wide range of results encountered in the real world makes ROI estimates useful even when based on uncertain estimates of risk parameters. In the situation above, we can confidently fund #3. Remaining funds, if any, can be applied cautiously to #2, but #1 is a clear looser.

Why Is HML Model Popular?

Given these serious defects is the HML model of risks, why does this approach receive so much support, particularly from government agencies? The answer appears to be rooted in the background and experience of the HML advocates. Most seem to be focused exclusively on the security of IT systems, and to have a background in government IT systems to the exclusion of business operations. This probably leads to several biases:

- Lack of experience in evaluating real-world risks and loss experiences quantitatively leads to the mistaken conclusion that such evaluations are infeasible, very difficult, or "too subjective".
- The losses experienced by third parties, i.e., citizen "customers", are ignored or downplayed when assessing the cost impact of risk events[8].
- The focus of the risk management is on "data" in the abstract rather than on processes and their data.
- A focus on IT-related risks, e.g., hackers, viruses, etc., to the exclusion of other risks, e.g., environmental and infrastructure.
- And finally, HML might be easy for an 'experienced expert' with 'years of industry knowledge' to conduct simply and quickly; but using such intuitive judgment as, 'I know from my 2-decades in this business and studying this topic that the risk of this event is always 'High', however there is no data to validate that the expert is actually correct!

The result is that governments tend to adopt IT risk management strategies with these characteristics:

- Risks are ranked using a data-focused HML model, and quantitative evaluations are passed over as infeasible.
- Risk mitigation measures are selected using HML-based guidelines without explicit regard to cost as long as the funding is available.
- Non-IT risks are downplayed or ignored.

The concentration on IT risks has another significant effect. The dominant IT risks are related to Internet access to IT systems, and are well known and well understood. Consequently the choice

[8] A risk analysis for a US government agency in which the agency took the position that the IT service interruption loss potential was determined solely by the loss of fees paid to it by other agencies for the use of its IT facility without regards to the dysfunction impact on end users!

of mitigation measures appears to be "obvious," not really requiring any analysis of risks. As a result, the application of limited risk management funds may be unbalanced and far from optimum[9].

A Valid Quantitative Risk Model:

This discussion shows that rational decisions regarding the trade-off relative to IT operational risk requires a valid quantitative risk model, which incorporates the relationships between risk parameters accurately. The stakes are too high relative to IT operational risk to leave it to subjective guesses or 'gut' feelings. The trade-offs between competing mitigation actions and counter-measures are often quite complex and can be very expensive, however, the consequence of doing nothing or doing the wrong thing can be enormous. My 17 years of practical experience of economically quantifying IT operational risk and using it to cost-justify IT investments suggest that any quantitative risk model must include the following parameters to be valid:

- Threat[10] occurrence rates – annualized.
- A set of service interruption durations – used to quantify threat and function parameters.
- Threat service interruption[11] profile – percentage of occurrences that results in each service interruption duration.
- Function, process or IT system loss potential[12] as a function of service interruption durations – monetary loss for each interruption duration.
- Function schedule – hours per day and days per year to model the effect of slack time on interruption losses.
- IT system data loss – monetary loss per hour of data lost when restoring stored data from a backup copy.
- Asset[13] and liability exposure[14] loss potentials – monetary loss per each worst-case threat event.
- Threat effect factor – percent of loss potential resulting from the occurrence of each treat with respect to each function, asset, and data loss ranging from 0% (no loss from an occurrence) to 100% (worst case threat-triggered loss).
- Mitigation measure effect on risk losses – reduction in threat occurrence rates, change in threat effect factors, and/or changes in threat outage duration profiles.
- Mitigation measure cost to implement in dollars.
- Mitigation measure cost to maintain – dollars per year.
- Remaining useful life of mitigation measure implementation – years
- Cost of money – percent per year.
- IT system recovery mode RTO and RPO performance parameters – hours.
- IT system recovery mode cost to implement for each IT system – dollars per year.
- The correct metric is Loss Expectancy not Loss Potential

[9] Several years ago the Board of Directors of a major bank discovered to its dismay that all 150 of its IT security people were working *exclusively* on Internet access risks.
[10] The term "threat" is used to refer to risk events that cause losses when they occur.
[11] Service interruptions are often referred to non-availability or denial of service. The resulting risk loss depends on the duration of the interruption, and so both threats and functions must be defined in terms of the defined set of durations
[12] "Loss potential" is a short hand term for the potential for loss, an inherent characteristic of functions and assets. For the purposes of the risk model loss potential is defined as the worst-case loss. This implies that there is a threat, which, when it occurs, will trigger a loss equal to the loss potential. Other threats will trigger lower or zero losses
[13] An asset may be physical property, stored data, or intellectual property.
[14] The term "liability exposure" is used to refer to the exposure to a claim for damages inherent in the character of the object of a risk assessment project.

A Key Benefit of Automation:

Apart from the obvious time saving when the detailed calculations required by a quantitative model are automated, the ease with which the calculations can be repeated has important benefits. An important aspect in the construction of a model is the aggregation of the real world. For example, one could postulate a wide range of threats when performing a risk assessment with relatively minor difference between some of the threats. Likewise, one can define a long list of functions and assets to be included in an analysis project, however, doing so greatly increases the time and cost to perform the data collection without a commensurate increase in the value of the analysis. This is because many of the fine details will have very little impact on the results, so the effort used to collect overly detailed data will have been wasted.

In the real world we find that the 80-20 rule applies to the input data. Typically about 20% of the input parameters will dominate the results. In other words, given the range of uncertainty in the parameters, it is a waste of time and money to *begin* at a high level of detail. Instead it is better to begin with an aggregate model, and let the initial results identify the important input parameters. If necessary these can be modified to add appropriate details, and the analysis easily repeated. Similarly, if a particular parameter with a low level of confidence is seen to be important to the results, it is easy to test its sensitivity by trying a range of values to determine the importance of the low confidence level estimate.

Conclusions:

The inherent disconnects between Business Impact Analyses and Risk Assessments negate cost-benefit balancing and therefore the combination of BIAs and RAs cannot answer the Big Question: how to optimize the allocation of limited resources to mitigate operational risk. In practice BIAs and RAs are nothing more than “informed guesswork” in disguise. Extensive research has documented that intuition and subjective judgment produce systematic errors regarding risk and in reality, BIAs and RAs are effectively based on intuition and subjective judgment.

Additionally, there is no reason to conclude that the use of High-Medium-Low scales to evaluate threats and assets will yield more credible results than quantitative estimates. *The exact opposite is true* for three reasons.

1. Based on the fact that real world risk parameters typically span a range of a million or more to one, quantitative estimates with an uncertainty range of as much as ±50% will still be more than an order of magnitude more precise that HML-based rankings.
2. HML rankings cannot be used to estimate cost-benefit or ROI, which are *essential* to optimizing the allocation of finite risk management resources. Indeed, if an ROI is greater than about 50%, a not uncommon finding, implementation of the risk mitigation measure will be supported even if the underlying risk reduction has been overestimated by half an order of magnitude.
3. The HML scale introduces unavoidable ambiguities for risks and loss potentials at or near the H-M and M-L boundaries. Quantitative estimates avoid these ambiguities.

Net: HML simply lacks the granularity of quantitative risk modeling and relies ‘intuition & experience’ rather than data. This is a profound error.

It seems likely that support for the use of HML risk ranking is based on the mistaken belief that it is impossible to develop credible quantitative risk estimates. That belief, however, is imaginary as real world experience shows that there are a wealth of data on which to base quantitative risk estimates with a degree of precision and confidence sufficient to support sound risk management decisions. We need to apply the appropriate level of discipline relative to the complexity of the situation.

While precise information is extremely valuable, even small amounts of reliable data is infinitely better that than relying on subjective value judgments when it comes to making sound decisions about managing IT-infrastructure risk. Data is increasing available. There is a surprising amount of information from which to make realistic calculations and estimates about IT infrastructure and operational risks. As Immanuel Kant said “We Have a Duty - Especially Where the Stakes Are Large - to Inform Ourselves Adequately about the Facts of the Situation.” All in all, it is far better to use empirical data than rely on intuitive, subjective judgments.

We must make informed estimates about future losses and then take action based on the estimate. The underlying model, however, must be constructed to accurately portray all of the pertinent risk parameters. To develop a proper proportional response we must inform ourselves accurately about the facts of the situation. To keep the odds in our favor we must economically-quantify the operational risks of the IT-infrastructure so that we can properly evaluate the many tradeoffs, arriving at the optimal solution for our organization.

Bibliography

Bernstein, Peter L. Against the gods: The remarkable story of risk. New York: John Wiley & Sons, 1996.

Carr, N.G. "IT doesn't matter." IEEE Engineering Management Review 32 (2004): 24.

"Carr revisited." Information Age (London, UK) 10 June 2005.

"Continuity Insights." Focus On The Largest Source Of Risk: The Data Center. 20 Jan. 2013 <http://www.continuityinsights.com/articles/2012/06/focus-largest-source-risk-data-center>.

"CORAÂ® Cost-of-Risk Analysis by International Security Technology, Inc." CORAÂ® Cost-of-Risk Analysis Software. 20 Jan. 2013 <http://www.softscout.com/software/Project-and-Business-Management/Risk-Management/CORA--Cost-of-Risk-Analysis.html>.

Hour, In Less than an. "Is Knight's $440 million glitch the costliest computer bug ever?" CNNMoney. 09 Aug. 2012. Cable News Network. 25 Jan. 2013 <http://money.cnn.com/2012/08/09/technology/knight-expensive-computer-bug/index.html>.

"How Many Data Centers? Emerson Says 500,000." Data Center Knowledge RSS. 21 Jan. 2013 <http://www.datacenterknowledge.com/archives/2011/12/14/how-many-data-centers-emerson-says-500000/>.

Kahneman, Daniel. "A perspective on judgment and choice: Mapping bounded rationality." American Psychologist 58 (2003): 697-720.

Kahneman, Daniel. Thinking, fast and slow. New York: Farrar, Straus and Giroux, 2011.

Lewis, Nigel Da Costa. Operational risk with Excel and VBA: Applied statistical methods for risk management. Hoboken, NJ: Wiley, 2004.

Power'..., Emerson Network. "Calculating the Cost of Data Center Outages." Scribd. 21 Jan. 2013 <http://www.scribd.com/doc/56246462/Calculating-the-Cost-of-Data-Center-Outages>.

Power, Michael. "The risk management of everything." The Journal of Risk Finance 5 (2004): 58-65.

Rescher, Nicholas. Luck: The brilliant randomness of everyday life. New York: Farrar, Straus and Giroux, 1995.

"Simple Architectures for Complex Enterprises." : The IT Complexity Crisis: Danger and Opportunity. 26 Jan. 2013 <http://simplearchitectures.blogspot.com/2009/11/it-complexity-crisis-danger-and.html>.

"Software Testing Lessons Learned From Knight Capital Fiasco." CIO. 14 Aug. 2012. 25 Jan. 2013 <http://www.cio.com/article/713628/Software_Testing_Lessons_Learned_From_Knight_Capital_Fiasco>.

Sunstein, Cass R. Laws of fear: Beyond the precautionary principle. Cambridge, UK: Cambridge UP, 2005.

Sunstein, Cass R. Probability neglect: Emotions, worst cases, and law. [Chicago, Ill.]: The Law School, University of Chicago, 2001.

Sunstein, Cass R. Risk and reason: Safety, law, and the environment. Cambridge [England: Cambridge UP, 2002.

Tversky, A., and D. Kahneman. "The framing of decisions and the psychology of choice." Science 211 (1981): 453-58.

Tversky, A., and D. Kahneman. "Judgment under Uncertainty: Heuristics and Biases." Science 185 (1974): 1124-131.