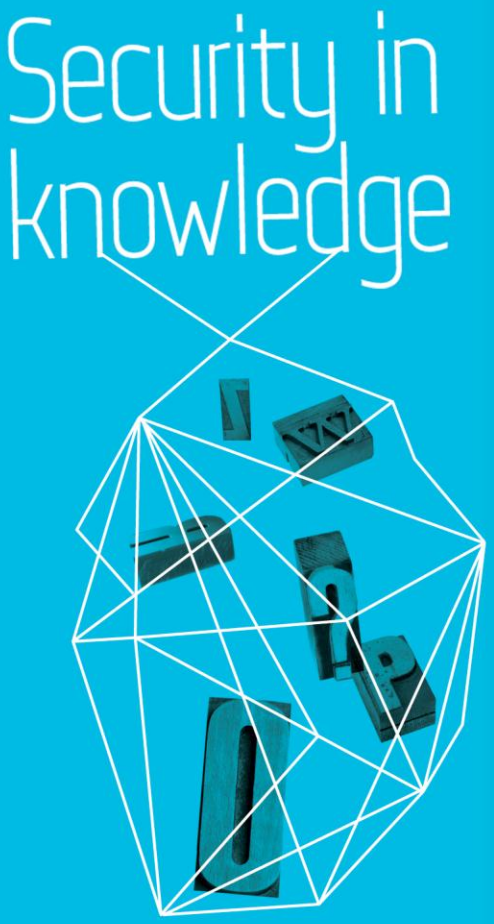

The Probability of Exploit

Predictive Analytics & Security Management

Doug Hubbard
Hubbard Decision Research, CEO

Richard Seiersen
Kaiser Permanente, Director Information Security

Session ID: GRC-W21

Session Classification: General Interest

Prelude:
The Truth About
Current Security Models

Prelude: True, Not True or Irrational?


+

+

÷3

=

Prelude: True, Not True or Irrational

OWASP Inherent Risk Rating Worksheet

		Threat Agent Factors				Vulnerability Factors						Business Impact Factors							
ID	**Risk Description**	Skill level	Motive	Opportunity	Size	Ease of discovery	Ease of exploit	Awareness	Intusion detection	**Likelihood**	**Likelihood Rating**	Patient safety	Regulatory	Revenue	Productivity	**Impact**	**Impact Rating**	**Risk score (L * I)**	**Risk rating**
1	Risk 1	4	9	4	5	9	9	6	3	**6.8**	**H**	2	0.5	2	0.1	**1.2**	**H**	**7.8**	**High**
2	Risk 2	6	1	0	2	3	5	4	3	**3.3**	**L**	0.1	0.5	0.5	0.1	**0.3**	**M**	**1.0**	**Low**
3	Risk 3	6	9	7	9	3	3	6	8	**7.1**	**H**	10	0.5	0.1	0.1	**2.7**	**E**	**18.9**	**Critical**
4	Risk 4	1	4	7	6	3	5	6	9	**5.7**	**M**	2	0.5	0.5	0.1	**0.8**	**M**	**4.4**	**Moderate**
																	Overall Assessment Risk	**32.2**	

E =	Extreme
H =	High
M =	Medium
L =	Low
N =	Negligible

- For each "risk" likelihood sums and divides 8 ordinal scores.
- Impact does this for 4 ordinal scores.
- Risk Score is the sum of these two factors
- The overall security assessment is the sum of the "Risk Scores"

Example Only

Prelude: We Believe (ALL) of These Are Equivalent!

Inherent Risk Rating Worksheet

		Threat Agent Factors				Vulnerability Factors						Business Impact Factors							
ID	**Risk Description**	Skill level	Motive	Opportunity	Size	Ease of discovery	Ease of exploit	Awareness	Intusion detection	**Likelihood**	**Likelihood Rating**	Patient safety	Regulatory	Revenue	Productivity	**Impact**	**Impact Rating**	**Risk score (L * I)**	**Risk rating**
1	Risk 1	4	9	4	5	9	9	6	3	**6.8**	**H**	2	0.5	2	0.1	**1.2**	**H**	**7.8**	**High**
2	Risk 2	6	1	0	2	3	5	4	3	**3.3**	**L**	0.1	0.5	0.5	0.1	**0.3**	**M**	**1.0**	**Low**
3	Risk 3	6	9	7	9	3	3	6	8	**7.1**	**H**	10	0.5	0.1	0.1	**2.7**	**E**	**18.9**	**Critical**
4	Risk 4	1	4	7	6	3	5	6	9	**5.7**	**M**	2	0.5	0.5	0.1	**0.8**	**M**	**4.4**	**Moderate**
																Overall Assessment Risk		**32.2**	

E = Extreme
H = High
M = Medium
L = Low
N = Negligible

=

Prelude: & Why Should You Care?

.....**Bad Models**, Over Confidence & Placebos

Summary of response distributions of 5-point scales

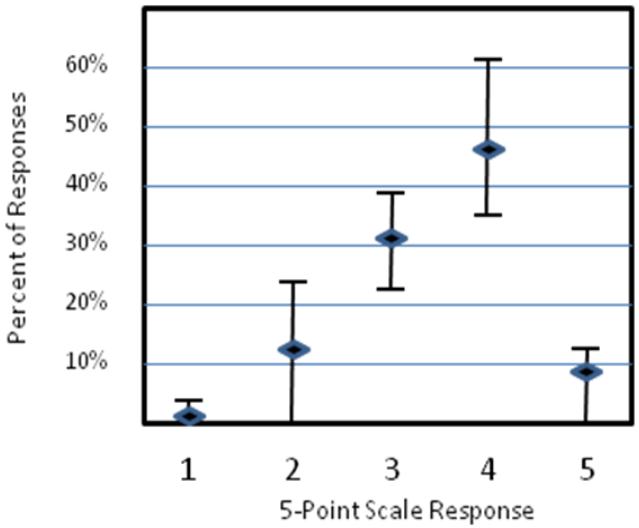

Hubbard D., Evans, D "Problems with scoring methods and ordinal scales in risk assessment" *IBM Systems Journal: Special Issue on Risk Management*, Oct 2009

- Scales obscure (rather than alleviate) the lack of information and create an illusion of communication (Budescu)
- The rounding effect of scales adds unexpected error (Cox).
- Arbitrary partitions have unexpected effects on scoring behavior (Fox)
- Scales introduce an error of "assumed ratios". (Hubbard, Evans)
- Clustering of responses amplifies all of the previously mentioned errors (Hubbard, Evans)

A Bad Model

Prelude: & Why Should You Care?

.....Bad Models, **Over Confidence** & **Placebos**

Gathering more information makes you feel more confident but, at some point, begins to reduce decision quality while confidence continues to increase.

- Collecting more than a few data points on horses makes expert handicappers worse at estimating outcomes (Tsai, Klayman, Hastie)
- Interaction with others only improves estimates up to a point, then they get worse (Heath, Gonzalez)
- Collecting more data about investments makes people worse at investing (Andreassen)
- Collecting more data about students makes counselors worse at predicting student performance. (Andreassen)
- An experiment with a popular method called "Analytic Hierarchy Process" shows confidence increased whether decisions are improved or degraded. (Williams, Dennis, Stam, Aronson)

We should *assume* increased confidence from analysis is a "placebo". Real benefits have to be measured.....we will show you how.

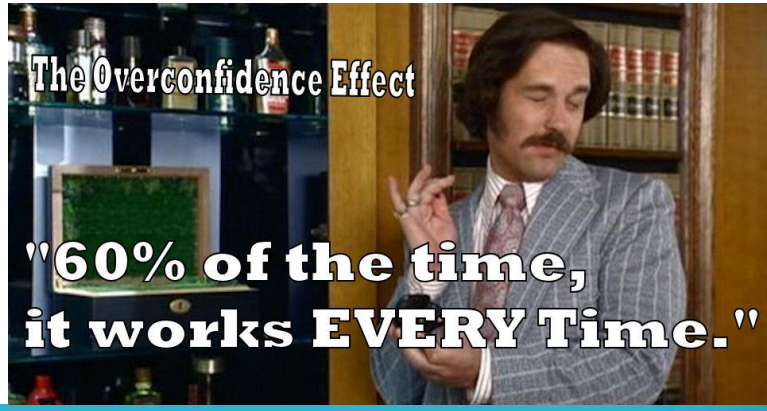

What We Will Cover

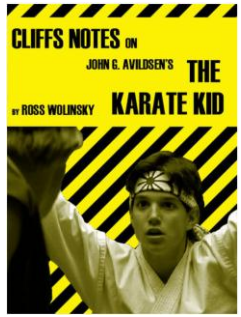

- Current Model Problems & Solutions:
  - The Cliff Notes
- Fixing The Model
  - The first step in recovery is admitting you have a problem
  - Calibrating the experts……to be like bookies
  - Streamlining Scores……increasing consistency
  - Speaking in probabilities……real numbers, real math
  - Putting it all together…….Monte Carlo Simulations to the rescue

Current Model Problems and Solutions (The Cliff Notes)

Model Problems and Solutions: The Cliff Notes

Questions We Should Be Able To Answer

- Here are some risks plotted on a "typical heat map".

- Suppose mitigation costs were:
  - Risk 1: $725K - **High**
  - Risk 2: $95K - **Low**
  - Risk 3: $2.5M - **Critical**
  - Risk 4: $375K - **Moderate**

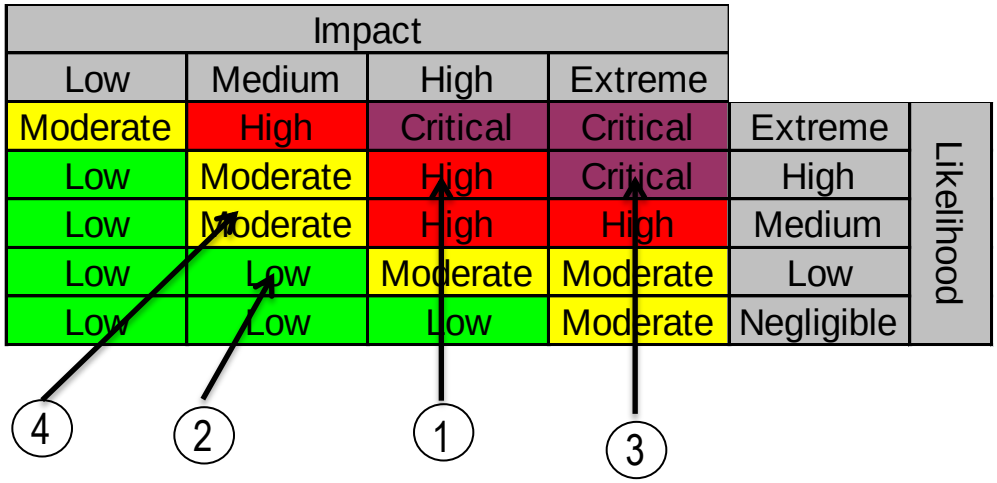

- What mitigations should be funded and what is the priority among those?

Example Only

Model Problems and Solutions: The Cliff Notes

Decision Making Supported Quantitatively

- If risks and mitigation strategies were quantified in a meaningful way, decisions could be supported.
- In order to compute an ROI on mitigation decisions, we need to quantify likelihood, monetary impact, cost, and effectiveness

Risk	Likelihood / Yr	Impact / Yr	Mitigation Effectiveness	Mitigation Cost / Yr	Mitigation ROI	Action
Risk 1	37%	$2M to $40M	95%	$725K	725%	Mitigate
Risk 2	11%	$50K to $400K	100%	$95K	-80%	Track
Risk 3	34%	$5M to $80M	90%	$2.5M	329%	Monitor
Risk 4	29%	$500K to $20M	98%	$375K	437%	Mitigate

- The optimal solution would be to mitigate Risks1 & 4 first.
- If you have the resources, then mitigate Risk 3.
- Risk 2 is not worth fixing.

Example Only

Model Problems and Solutions: The Cliff Notes

Experts Vs. Quantitative Models

Paul Meehl assessed 150 studies comparing human experts to statistical models in many fields (predicting football games, the prognosis of liver disease, etc.).

"There is no controversy in social science which shows such a large body of qualitatively diverse studies coming out so uniformly in the same direction as this one. [With hardly] a half dozen studies showing even a weak tendency in favor of the [human expert], it is time to draw a practical conclusion."

Philip Tetlock tracked a total of over 82,000 forecasts from 284 political experts in a 20 year study covering elections, policy effects, wars, the economy and more.

"It is impossible to find any domain in which humans clearly outperformed crude extrapolation algorithms, less still sophisticated statistical ones."

Fixing The Model…
Quantitatively

Fixing The Model: The Problem

Inherent Risk Rating Worksheet

		Threat Agent Factors				Vulnerability Factors						Business Impact Factors							
ID	Risk Description	Skill level	Motive	Opportunity	Size	Ease of discovery	Ease of exploit	Awareness	Intrusion detection	Likelihood	Likelihood Rating	Patient safety	Regulatory	Revenue	Productivity	Impact	Impact Rating	Risk score (L * I)	Risk rating
1	Risk 1	4	9	4	5	9	9	6	3	6.8	H	2	0.5	2	0.1	1.2	H	7.8	High
2	Risk 2	6	1	0	2	3	5	4	3	3.3	L	0.1	0.5	0.5	0.1	0.3	M	1.0	Low
3	Risk 3	6	9	7	9	3	3	6	8	7.1	H	10	0.5	0.1	0.1	2.7	E	18.9	Critical
4	Risk 4	1	4	7	6	3	5	6	9	5.7	M	2	0.5	0.5	0.1	0.8	M	4.4	Moderate
																		Overall Assessment Risk	32.2

E =	Extreme
H =	High
M =	Medium
L =	Low
N =	Negligible

- The existing model starts with subjective scores for likelihood and impact
- Like most "popular" risk management models, scores are "added" to get ratings for likelihood and impact and a total "risk score".
- The scale used indicates relevant order, not actual units of measure (i.e. not real numbers), which introduces ambiguity; "adding" or "multiplying" adds further obscurity to the items being measured and compromises the risk assessment process.

Example Only

Fixing The Model: Subjectivity

Human expertise is an important input and it is hard to completely automate. But there are certain types of errors in human judgment we know how to measure and control for:

- Overconfidence – Their chance of being right is much less than they believe
- Influence by irrelevant factors – Factors like the order in which you consider projects, whether it is a 5-point scale or a 10-point scale, or how much other people in the room smile all affect your judgments
- Inconsistency – When given the same sets of problems to evaluate, experts have a hard time giving the same answers; also, their memory is reconstructed so that they believe they always had one preference when in fact they didn't

"Experience is inevitable. Learning is not."
Decision science researcher Paul Schoemaker

Fixing The Model: Ambiguity Exercise

- What has a higher overall risk?
  - Two risks of "medium" impact and "high" likelihood
  - One risk of "high" impact and "medium" likelihood
  - Four risks of "low" impact and "medium" likelihood
- What is the probability of a low, medium, and high likelihood? Is this a probability per year, per month, ever, etc?
- On a 1 to 5 scale, what is more likely?
  - At least one of three events of a "2" likelihood
  - An event of "4" likelihood

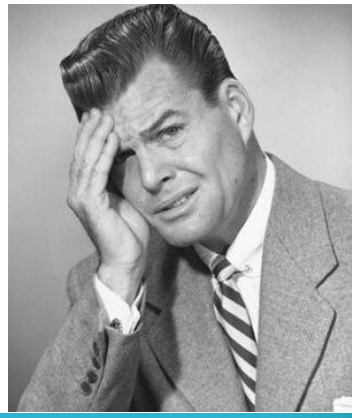

Fixing The Model: The Problem…ambiguity

- Using scores like 'Low,' 'Medium,' 'High,' etc., adds ambiguity, as people have wildly different interpretations of these labels for risk.*
- When asked to estimate the likelihood of different risks, participants were consistent with other research on the use of these scales (see chart below).
  - For example, SME's may equate a Likelihood of 'Extreme' with a 5% chance of occurrence, while also equating a Likelihood of 'High' to a 40% chance of occurrence; this is inconsistent as higher probability should equate to a higher Likelihood rating.
- Using unambiguous quantitative estimates (e.g., chance of occurrence per year, 90% confidence interval for monetary impact) enables us to better assess and manage key Company risks.

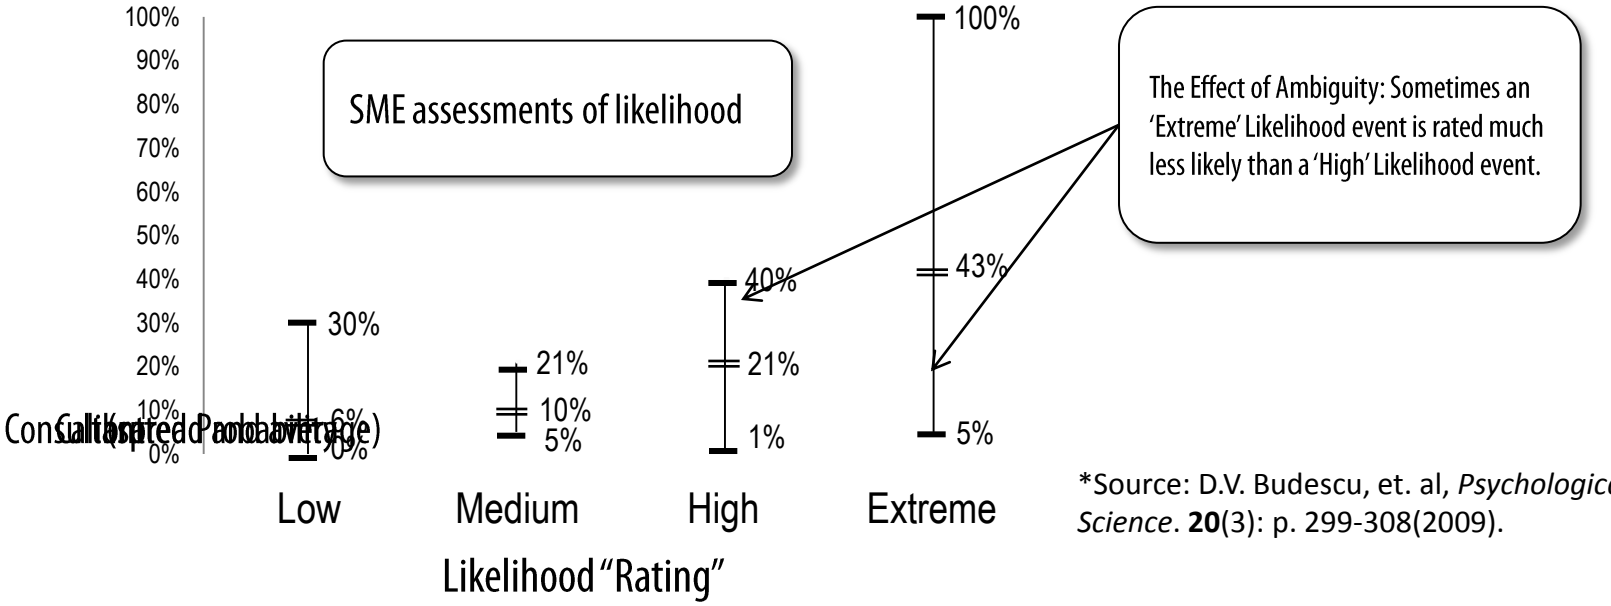

*Source: D.V. Budescu, et. al, *Psychological Science*. **20**(3): p. 299-308(2009).

Fixing The Model: Calibration Exercise

For the initial calibration test, you have two types of questions:

- For the questions that ask for a range, provide an upper and lower bound that you are 90% certain contains the correct answer
  - Napoleon Bonaparte was born what year? _____ to _____
  - What is the average weight of an adult male African elephant (tons)? _____to_____

- For the true/false questions, circle true or false and then circle the percentage that best represents your confidence in your response
  - Brazil has a larger population than Spain. True/False_____ Confidence:____%
  - A hockey puck will fit in a golf hole. True/False_____ Confidence:____%

Fixing The Model: Calibrating Experts…to be like bookies

- Research shows that assessing subjective probabilities and ranges is a teachable skill.
- One of the calibration tasks for participants in training was estimating 90% Confidence Intervals.
- At first, they were like most people – extremely overconfident.
- By fourth test, almost everyone could estimate 90% confidence intervals that contained the actual value about 90% of the time.

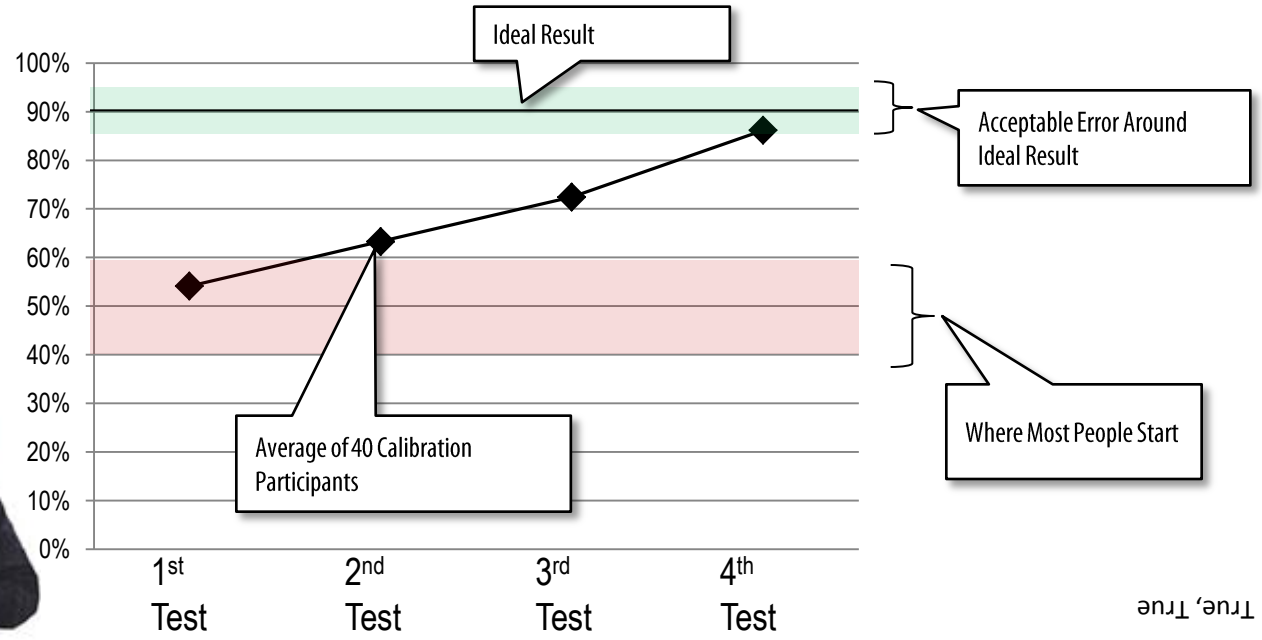

A: 1769, 3.5 tons, True, True

Fixing The Model: Increasing Consistency

…..getting rid of redundant scores

Ease of Exploit predicts Skill Level

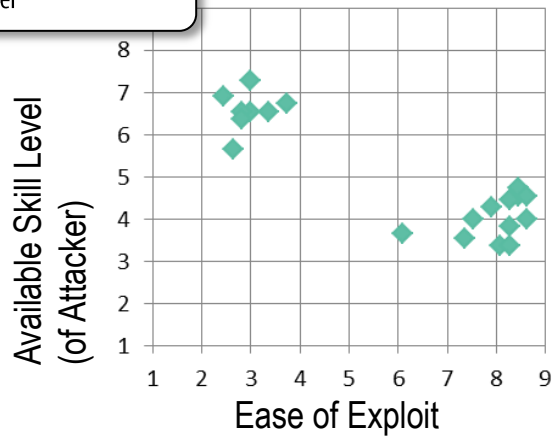

Ease of Discovery predicts Awareness

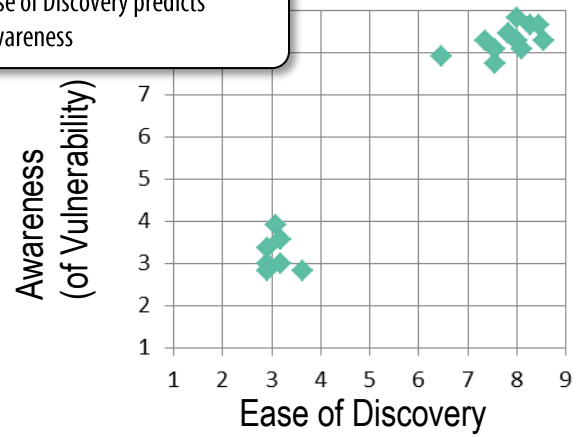

- Via independent tests, 11 SME's were asked to assign the 8 original factors and assess annual probability for 20 specific scenarios.
- We specified that likelihood mean *chance per year of >1 occurrence.*
- Irrelevant factors – Motive and Intrusion Detection had almost no correlation to the estimated annual probability.
- Redundant factors – Available Skill Level and Awareness were almost entirely predictable based on other scores (within +/- 1 point).
- The best model used only 4 of 8 factors.
- Scores are still used at this stage, but are minimized and converted directly to a meaningful probability.

Fixing The Model: Speaking in Probabilities

. . . .and further increasing consistency

- We developed a model that estimates the average probability scores of the security consultants.
- Comparing the new model we developed to actual security consultant group assessments shows that our model is more consistent than using security consultant probability assessments alone.
- Reducing inconsistency is an immediate benefit of the new model.

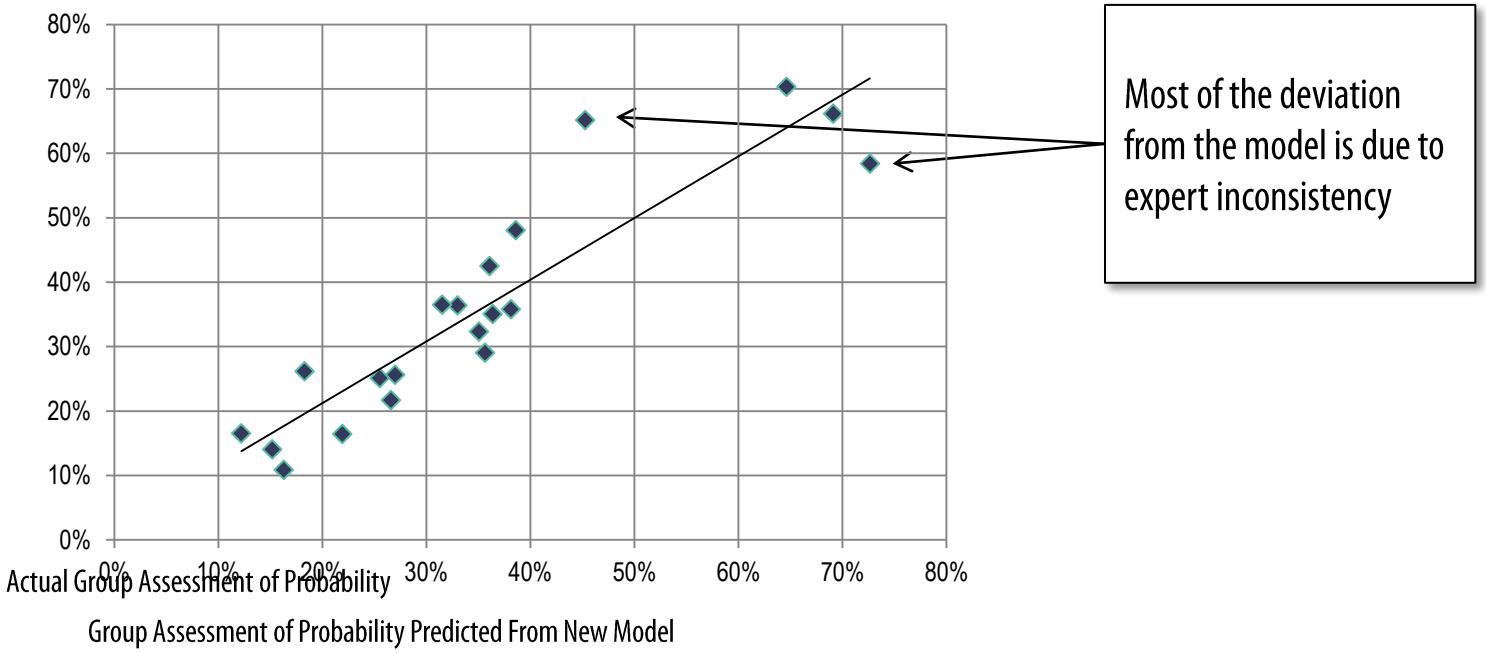

Putting It All Together... *(Monte Carlos To The Rescue)*

Putting it All Together – Monte Carlo Simulations

Likelihood

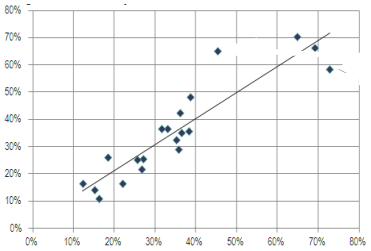

Business Impact

Mitigation Cost

Monte Carlo Simulation

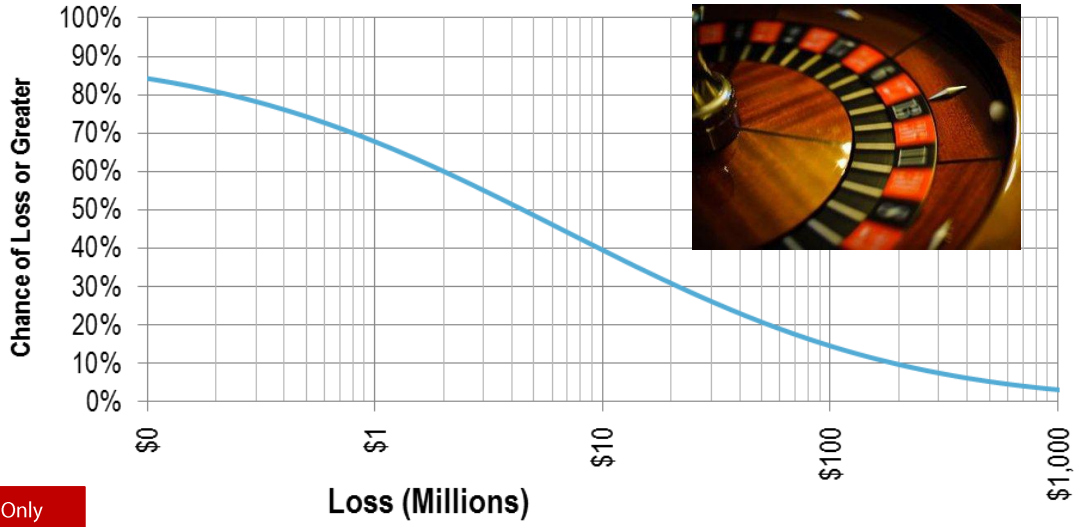

Monte Carlo Simulation

A computer simulation that seeks to determine the likelihood of various scenarios by running multiple simulations using random variables. The results of the Monte Carlo simulation show the most likely outcomes.

Used extensively in Physical sciences, Engineering, Computational biology, Applied statistics, Games, Design and visuals, Finance and business and Telecommunications

Example Only

Putting it All Together: Communicating Security Risks Quantitatively

- If we can express uncertainty of individual risk events quantitatively, we can answer quantitative risk questions like "What is the chance a given risk will result in more than $5 million in losses in a given year?"
- The curve on this chart is based on calibrated estimates and the new model
- It represents the chance that a loss for a given a risk equal to or greater than some amount will occur in a given year.
- These can also be added up into Project and Portfolio risk in a meaningful way.

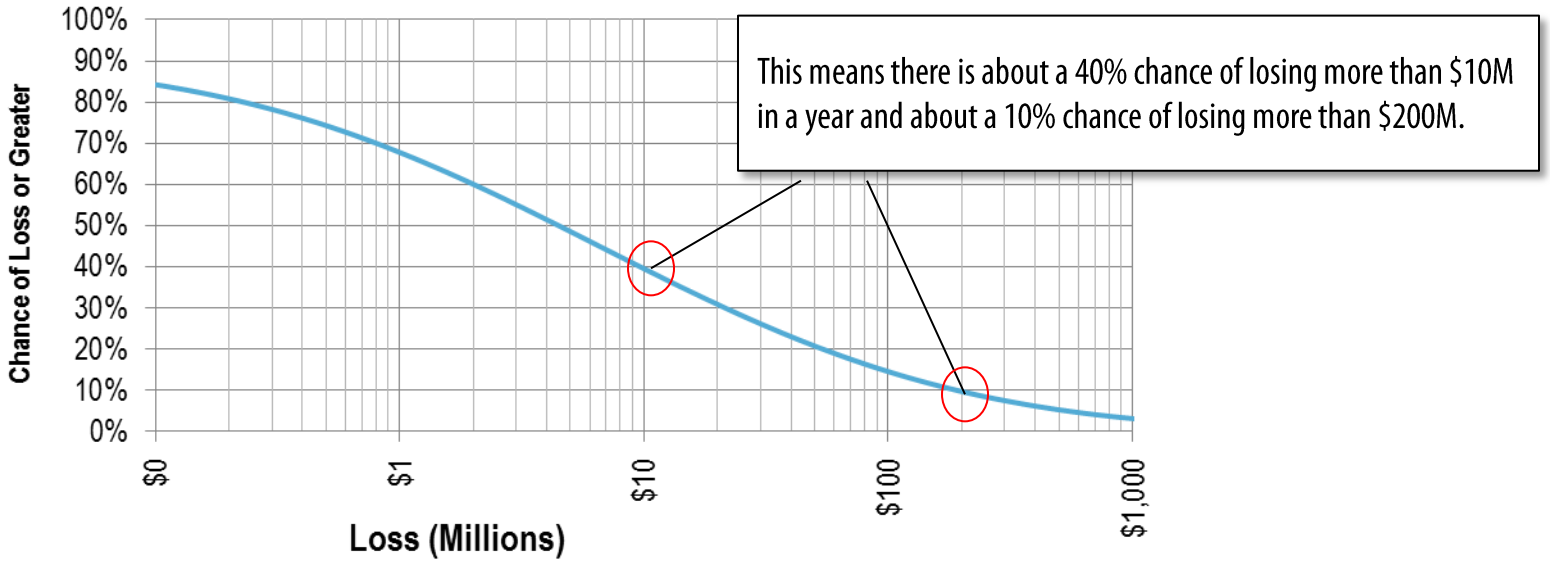

Example Only

Putting It All Together: Aggregating Security Risk

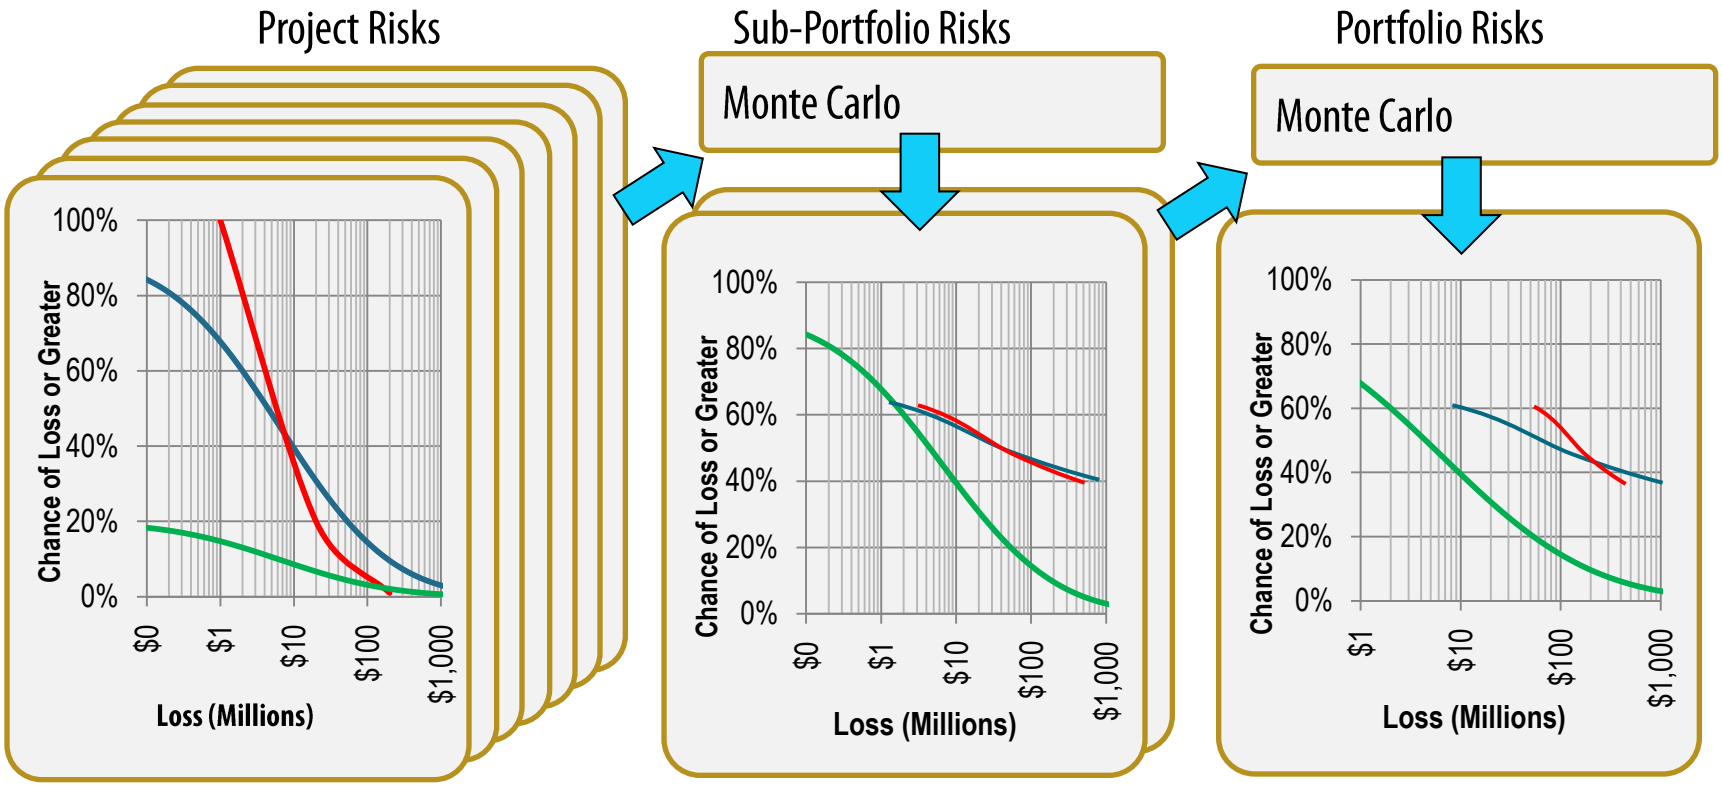

- A Monte Carlo simulation is used to roll Project Risks into Sub-Portfolio Risks into Portfolio Risks
- Each level can have an acceptable risk boundary defined by management

Putting It All Together: Technical Details

Risk ID-->	A	B	C
Risk Description	**DB Access**	**File Access**	**Network Access**
Opportunity	6	6	3
Size of Attack Agent	4	5	4
Ease of Discovery	4	7	7
Ease of Exploit	7	7	8
Probability per year	27%	35%	26%
Impact LB	$ 3,000,000	$ 100,000	$ 100,000
Impact Median	$ 30,000,000	$ 911,043	$ 1,732,051
Impact UB	$ 300,000,000	$ 8,300,000	$ 30,000,000
Extreme Impact scenario (1% of years)	$ 362,198,485	$ 11,731,027	$ 36,899,388
Impact Mean (if event occurs)	$ 79,905,469	$ 2,245,237	$ 7,784,358
log mean	7.477121255	5.959539046	6.238560627
Expected loss per year	**$ 24,688,697**	**$ 969,4**	
Mitigation cost per year	$ 200,000	$ 400,0	
Mitigation effectiveness	95%	9	
Return on Mitigation	**11627%**	**11**	
Mitigation Action Taken? 1 = yes	**1**		
Prob per year, post-mit	1.331%	3.50	
Expected loss per year, post-mitigation	**$ 1,068,820**	**$ 80,1**	

All Risks can be added to compute Project Risk before and after mitigation relative to acceptable risk.

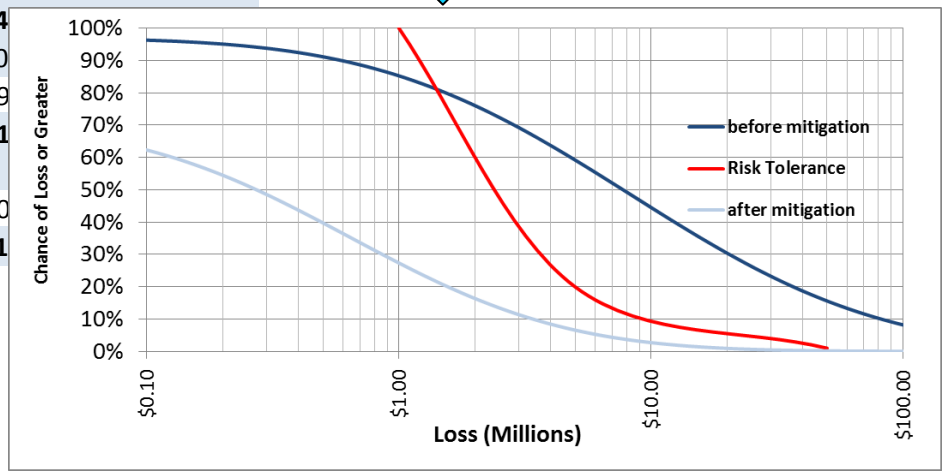

Blue shaded areas are compute from inputs provided in white background areas

Example Only

Putting It All Together: New Quantitative Model in Action

What action to take if only $3.5M available for mitigation?

	Expected Loss / Yr	Mitigation Cost	Mitigation Effectiveness	Mitigation ROI	Action
DB Access	$24.7M	$800K	95%	**2,832%**	**Mitigate**
Physical Access	$2.5M	$300K	99%	**727%**	**Mitigate**
Application Authentication	$4.4M	$600K	95%	**602%**	**Mitigate**
Data in Transit	$2.3M	$600K	95%	**267%**	**Mitigate**
System Patch	$2.8M	$800K	95%	**230%**	**Mitigate**
Network Access Control	$2.3M	$400K	30%	**74%**	**Mitigate**
File Access	$969K	$600K	90%	**45%**	**Monitor**
Web Vulnerabilities	$409K	$800K	95%	**-51%**	**Track**
System Configuration	$113K	$500K	100%	**-77%**	**Track**

$3.5M

Example Only

Putting It All Together: Project Loss Probability Curves

What is risk exposure after applying available mitigations?

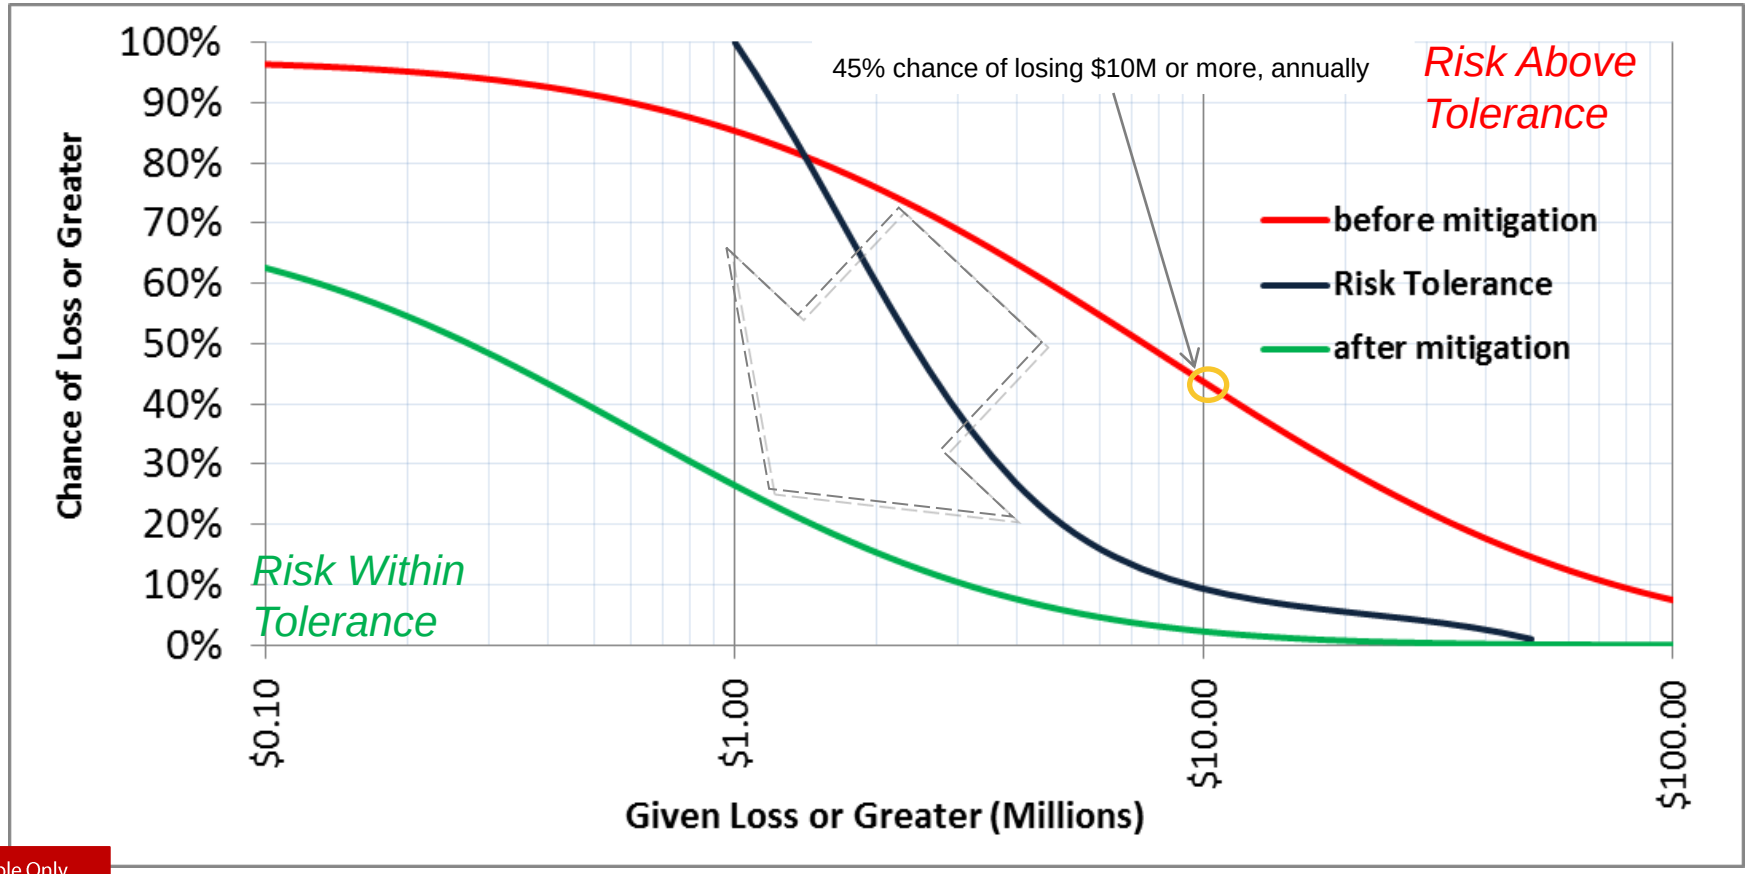

Example Only

Putting It All Together: BI Applied

Security Assessment & Analytics Services

Quantitative Loss Analytics

Note: All probabilities are predicated on successfull exploit of each environment as well as calibrated expert loss analytics.

Loss Category	Probability Per Year	Expected Loss/Yr	Mitigation Cost/Yr	Mitigation Effectiveness	Return On Mitigation (ROI)	
Application Authentication	43%	$4,431,310.38	$600,000.00	95%	602%	Mitigate
Data in Transit	26%	$2,319,492.00	$600,000.00	95%	267%	Mitigate
DB Access	27%	$24,688,696.51	$800,000.00	95%	2832%	Mitigate
File Access	35%	$969,468.78	$600,000.00	90%	45%	Monitor
Network Access	26%	$2,319,492.00	$400,000.00	30%	74%	Mitigate
Physical	28%	$2,506,265.07	$300,000.00	99%	727%	Mitigate
System Config	36%	$112,914.78	$500,000.00	100%	-77%	Track
System Patch	30%	$2,779,445.44	$800,000.00	95%	230%	Mitigate
Web App Vuln	34%	$409,154.11	$800,000.00	95%	-51%	Track

✓ Keep Only ✕ Exclude

Action Translation:	**Monitor**
Loss Category:	**File Access**
Expected Loss/Yr:	**$969,468.78**
Extreme Impact:	**$20,827,780.11**
Mitigation Cost/Yr:	**$600,000.00**
Mitigation Effectiveness:	**90%**
Probability Per Year:	**35%**
Return On Mitigation (ROI):	**45%**
Number of Records:	**175**

Findings Mid-Level Detail

Findings Cate..	Vulnerabilities	Existing Controls	Description	Required Remediation	
	authentication data (multiple databases i..	Database access is password controlled butp..	unauthorized access to data and/or PHI (e.g. steal password hashes secret q/a etc.)"	passwords that are properly protected."	
File Access	"Inadequate password length complexity or uniqueness. (multiple instances) See Appendix G.2 G..	Some missing passwords same as username default passwords many trivially guessable.	"Unauthorized access to data administrative control of servers and infrastructure resources. Attacker can steal PHI destroy data and/or cause denial of service. Compromised root account (Appendix G.2) a..	Ensure all passwords comply with KP policies for length complexity uniqueness.	Extreme
	"Missing / inadequate controls for access to application source code. See Appendix G.21"		"Unauthorized Access to Member Data (PHI) Denial of Service Data Corruption. Disclosure of embedded access credentials access control algorithms ability to analyze and discover controls and potential vulnerabilities."	"Ensure access credentials (e.g. ftp accounts) are properly secured. Ensure use of strong passwords that comply with KP policy for length and complexity."	High
	"Unsecured files and file shares containing PHI (multiple instances validated min 26K records) See Appendix G.7 G..	Accessible by domain-authenticated users.	Unauthorized Access to Member Data (PHI)	Ensure all PHI is secured with appropriate least-privilege controls.	Extreme
	"Unsecured PII See Appendix G.6"	Accessible by domain-authenticated users.	Unauthorized Access to PII	Ensure all PII is secured with appropriate least-privilege controls.	

Share Remember my changes Download

Local intranet 100%

Example Only – Not Real Data

Calculation Detail

Risk ID-->	A	B	C	D	E	F	G	H	I	Project
Risk Description	**DB Access**	**File Access**	**Network Access**	**Web App Vuln**	**App Authen**	**System Patch**	**Physical**	**Data in Transit**	**System Config**	
Opportunity	6	6	3	5	7	3	6	3	5	
Size of Attack Agent	4	5	4	5	5	5	7	4	5	
Ease of Discovery	4	7	7	8	8	8	3	8	8	
Ease of Exploit	7	7	8	7	8	8	4	7	8	
Probability per year	27%	35%	26%	34%	43%	30%	28%	26%	36%	**97%**
Impact LB	$ 3,000,000	$ 100,000	$ 100,000	$ 100,000	$ 100,000	$ 100,000	$ 100,000	$ 100,000	$ 100,000	**$ 341,724**
Impact Median	$ 30,000,000	$ 911,043	$ 1,732,051	$ 565,685	$ 1,732,051	$ 1,732,051	$ 1,732,051	$ 1,732,051	$ 223,607	**$ 8,605,090**
Impact UB	$ 300,000,000	$ 8,300,000	$ 30,000,000	$ 3,200,000	$ 30,000,000	$ 30,000,000	$ 30,000,000	$ 30,000,000	$ 500,000	**$ 133,349,880**
Extreme Impact scenario (1% of years)	$ 362,198,485	$ 11,731,027	$ 36,899,388	$ 4,137,352	$ 54,984,362	$ 41,634,906	$ 38,890,679	$ 36,899,388	$ 570,736	$ 503,298,821
Impact Mean (if event occurs)	$ 79,905,469	$ 2,245,237	$ 7,784,358	$ 985,235	$ 7,784,358	$ 7,784,358	$ 7,784,358	$ 7,784,358	$ 252,029	
log mean	7.477121255	5.959539046	6.238560627	5.752574989	6.238560627	6.238560627	6.238560627	6.238560627	5.349485002	6.934755394
Expected loss per year	**$ 24,688,697**	**$ 969,469**	**$ 2,319,492**	**$ 409,154**	**$ 4,431,310**	**$ 2,779,445**	**$ 2,506,265**	**$ 2,319,492**	**$ 112,915**	**$ 40,536,239**
Mitigation cost per year	$ 800,000	$ 600,000	$ 400,000	$ 800,000	$ 600,000	$ 800,000	$ 300,000	$ 600,000	$ 500,000	$5,400,001
Mitigation effectiveness	95%	90%	30%	95%	95%	95%	99%	95%	100%	
Return on Mitigation	**2832%**	**45%**	**74%**	**-51%**	**602%**	**230%**	**727%**	**267%**	**-77%**	6.86
Mitigation Action Taken? 1 = yes	**1**	**1**	**1**	**0**	**1**	**1**	**1**	**1**	**0**	
Prob per year, post-mit	1.331%	3.503%	18.026%	33.951%	2.169%	1.500%	0.275%	1.288%	36.081%	70%
Expected loss per year, post-mitigation	**$ 1,068,820**	**$ 80,162**	**$ 1,548,407**	**$ 409,154**	**$ 170,833**	**$ 117,769**	**$ 21,460**	**$ 100,953**	**$ 112,915**	**$ 3,630,473**

Example Only

Selected Sources

- Tsai C., Klayman J., Hastie R. "Effects of amount of information on judgment accuracy and confidence" *Org. Behavior and Human Decision Processes,* Vol. 107, No. 2, 2008, pp 97-105
- Heath C., Gonzalez R. "Interaction with Others Increases Decision Confidence but Not Decision Quality: Evidence against Information Collection Views of Interactive Decision Making" *Organizational Behavior and Human Decision Processes,* Vol. 61, No. 3, 1995, pp 305-326
- Andreassen, P." Judgmental extrapolation and market overreaction: On the use and disuse of news" *Journal of Behavioral Decision Making*, vol. 3 iss. 3, pp 153-174, Jul/Sep 1990
- Williams M. Dennis A., Stam A., Aronson J. "The impact of DSS use and information load on errors and decision quality" *European Journal of Operational Research,* Vol. 176, No. 1, 2007, pp 468-81
- Knutson et. al. "Nucleus accumbens activation mediates the influence of reward cues on financial risk taking" *NeuroReport*, 26 March 2008 - Volume 19 - Issue 5 - pp 509-513
- A small study presented at Cognitive Neuroscience Society meeting in 2009 by a grad student at U. of Michigan showed that simply being briefly exposed to smiling faces makes people more risk tolerant in betting games.
- Risk preferences show a strong correlation to testosterone levels – which change daily (Sapienza, Zingales, Maestripieri, 2009).
- Recalling past events that involved fear and anger change the perception of risk (Lerner, Keltner, 2001).