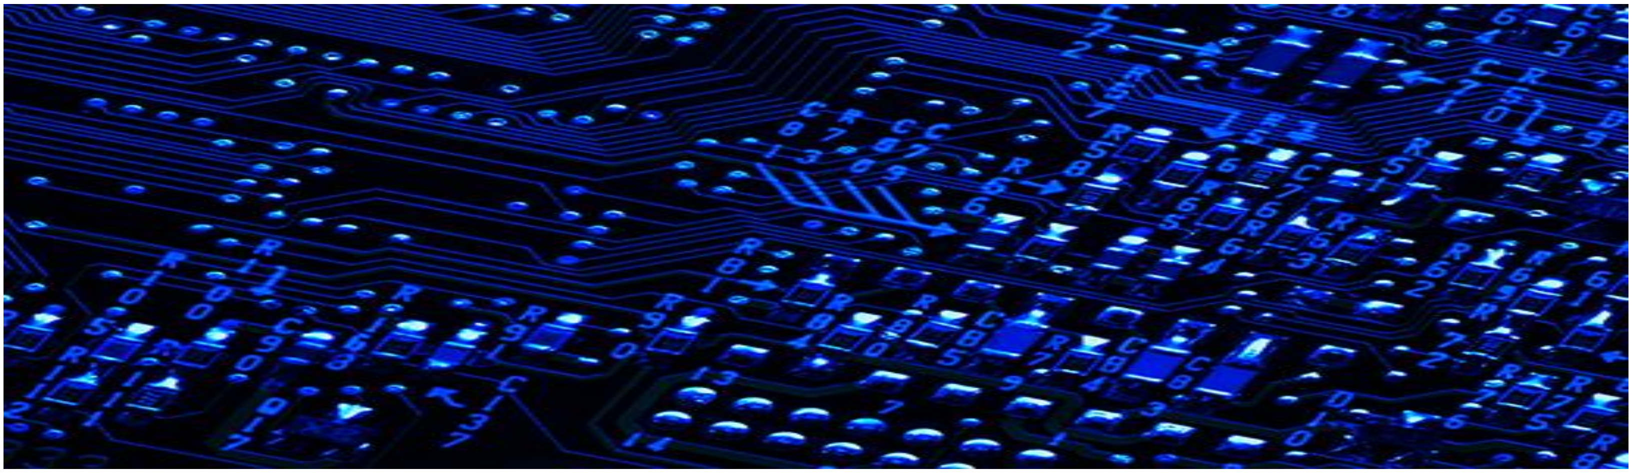

HDR Opinion Survey of Quantitative Risk Assessment Methods

What Cybersecurity Professionals Believe, What Methods They Use, and What Should Change About It.

July 25, 2016

PREPARED BY
HUBBARD DECISION RESEARCH, INC. | 2S410 CANTERBURY CT. | GLEN ELLYN, IL 60137
www.hubbardresearch.com

Table of Contents

Executive Summary ... 4

Background ... 4

THE SAMPLE ... 6

- Survey Participants ... 6
- Industries Represented ... 6

THE RESULTS ... 8

- Frameworks and Methods Used ... 8
- Attitudes Regarding the Use of Quantitative Methods ... 9
- A Connection Between Attitudes and Statistical Literacy ... 11
- Statistics Quiz Performance vs. Self-Assessment ... 16
- Possible Relationships to Breach Frequency ... 17

CONCLUSIONS & RECOMMENDATIONS ... 20

- Summarizing The Survey ... 20
- Suggested Actions ... 21

ABOUT THE AUTHORS ... 22

Appendix A: Survey Questions Details and Distribution of Responses ... 25

- Background Information on Participants and Their Organizations ... 25
- Attitudes Regarding Quantitative and Qualitative Methods ... 32
- Statistical Literacy Questions ... 34

Appendix B: A (Brief) Discussion of the Statistical Methods Used ... 41

Executive Summary

Hubbard Decision Research conducted a survey of 171 cybersecurity professionals from around the globe to determine the state of quantitative and qualitative methods being used to assess risk in the cybersecurity industry. The survey also included a measure of the attitudes towards quantitative methods, the frequency of data breaches, as well as a short test of statistical literacy. Results showed:

- The majority of respondents favored the continued incorporation of quantitative methods in risk analysis.
- However, actual methods used in the industry show that quantitative methods are used significantly less often than "qualitative" methods.
- Furthermore, a correlation emerged between opinions regarding the use of statistics, and respondent's relative statistical skills.
- There is some indication that the use of quantitative methods is associated with lower breach frequency.
- Results also suggest that greater exposure to quantitative measures will breed a more positive outlook towards their adoption in common practice.

Background

Cybersecurity-related threats are a growing concern in every industry. According to one researcher, the current maximum breach size of 200 million [records] is expected to grow by 50 percent over the next five years.[1] Given the magnitude of the potential costs of a security breach, it is important to understand how corporations are determining the risk of such an event occurring, and why they use the methods they use. Since this survey is meant to support management

[1] Wheatley, S., Maillart, T., & Sornette, D. (2015). The Extreme Risk of Personal Data Breaches & The Erosion of Privacy. *arXiv preprint arXiv:1505.07684*.

decisions, it would also be appropriate to recommend specific actions if justified in the findings.

To that end, Hubbard Decision Research sought to gauge the use of and attitudes towards different methods used in cybersecurity risk assessment. The focus of this survey was a comparison of methods we can broadly describe as "quantitative" and "qualitative." For the purposes of this survey, we mean "quantitative" to refer to methods which involve actually computing the probability of various impacts. Qualitative methods do not attempt to explicitly assess probabilities. Instead, they convey relative values with verbal scales for likelihoods (e.g. "likely", "very unlikely" etc.) and impact ("low", "extreme", etc.). Qualitative methods may also represent these categories with an ordinal point scale (e.g., 1 to 5) but without explicitly using probabilistic methods.

Along with the specific types of methods being used, the survey assesses attitudes towards different methods, familiarity with statistical concepts in the cybersecurity profession, a self-assessment of the understanding of these concepts, and the frequency of breaches. We also gathered background information about the cybersecurity professionals themselves and the organizations for which they work.

THE SAMPLE

Survey Participants

One hundred seventy-one industry professionals were recruited from a variety of channels, including the Information Systems Audit and Controls Association (ISACA), the Society for Information Risk Assessment (SIRA), as well as four of the largest cybersecurity groups listed on the professional social network, LinkedIn. Here are a few relevant observations about the participants:

- The majority of participants identified themselves as analysts, engineers, or individual contractors. 24% identified themselves as top information security personnel (e.g., CISO, VP), and 20% identified themselves as managerial staff (e.g., mid-level manager).

- 93% of participants in the survey had at least 3 years of experience in cybersecurity and 64% had 10 more years.

- The majority also had at least one professional certification in information security – the most common being CISSP (46%) followed by CISM (23%).

- Only 9% stated they always worked in information security and 70% said they had worked in other areas of information technology (the remainder previously worked in fields unrelated to information technology).

Industries Represented

Survey participants represented over 17 different industries including services that specialize in information security. Specifically, the responses reveal the following:

- 28% of the sample came from consulting services that specialize in information security, whereas 16% of the sample provide information technology consulting services that do not deal with information security specifically.

- The majority of participants came from organizations whose main products and services are not related to cybersecurity, such as banking/insurance/credit services (29%), government (9%), healthcare

(8%) manufacturing (8%), telecommunication (5%), education (6%), and retail goods (4%).

- 62% stated they worked for organizations that had at least 1000 employees and 67% worked where there were six or more individuals specifically assigned to cybersecurity.

- 66% said they were at least sometimes required to provide financial justifications for the cost of security and 24% said they were always required.

A full list of all survey questions and details of the distribution of responses can be found in the Appendix A.

THE RESULTS

Frameworks and Methods Used

Part of the survey investigated the frameworks and methods used to assess cybersecurity risk. We use the term "framework" to refer to a broadly defined approach often promoted by standards organizations. Within these frameworks will be multiple specific methods such has how a likelihood is determined and represented.

- The National Institute of Standards and Technology (NIST) appeared to be the most popular security risk management framework with 59% of participants endorsing its use in their organization.

- Other popular frameworks included an International Organization for Standardization (ISO) Standard (49%), a proprietary framework (30%), and the Online Web Application Security Project (OWASP) (29%). 15% of our participants reported their organization does not use a security risk management framework.

- Specific methods varied even among those using the same framework. More than 60% of our survey participants reported using ordinal matrices – also called "heat maps" or "risk matrices."

- Generally, risk assessment methods became less popular as they grew more mathematically rigorous. For example, only 13% of respondents used Monte Carlo simulations. *Figure 1* breaks down the representation of different risk assessment methods among the sample.

- Only 34% of survey participants felt that their firm dedicated an adequate amount of time to risk assessment related to cybersecurity while 53% reported that the amount of time spent by their firm was too little.

- The majority of survey participants reported sharing internal data regarding security incidents with other entities outside their organization. In most cases, these entities were affiliated with the government. Overall, 43% reported sharing data from security incidents with select colleagues in their respective field or industry while 21% of participants reported sharing data from security incidents with other firms.

Figure 1: Proportion of Organizations Using Various Information Security Risk Assessment Methods

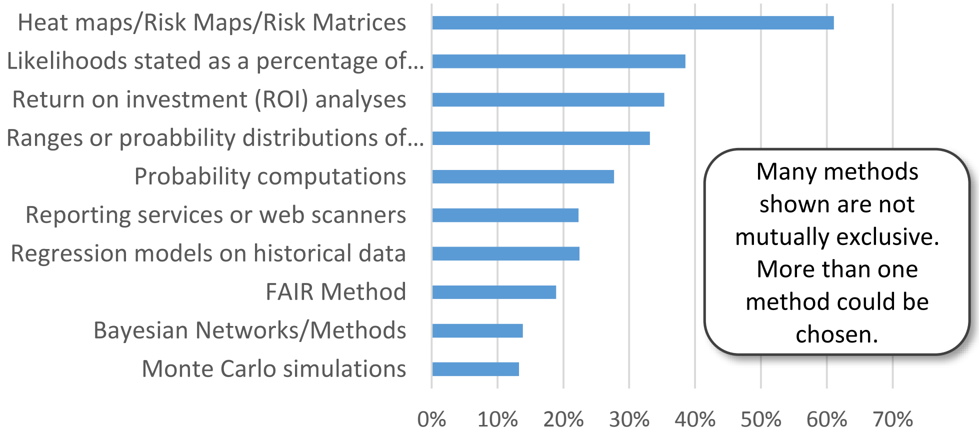

Many methods shown are not mutually exclusive. More than one method could be chosen.

Attitudes Regarding the Use of Quantitative Methods

In addition to investigating what methods and frameworks are currently being used by organizations, we asked individual participants about their own attitudes towards the use of quantitative vs. qualitative methods in the cybersecurity space. The participants were given 18 statements regarding these attitudes with which they would agree, disagree or say they had no opinion. The responses were coded as being either "favorable" or "unfavorable" toward quantitative methods which were also interpreted, respectively, as "unfavorable" or "favorable" toward qualitative methods.

For example, agreeing with the statement "Cybersecurity should eventually adopt more sophisticated probabilistic methods based on actuarial methods where it has not already done so." was coded as favorable toward quantitative methods. On the other hand, agreeing with statements like "Information security is too complex to model with probabilistic methods." was coded as unfavorable to quantitative methods. The total number of responses favorable and unfavorable toward quantitative methods were added up for each participant. A complete list of these questions and the distribution of responses is in the Appendix A. A summary of the results follows:

- We found that 86% had a generally positive attitude toward the adoption of quantitative methods – that is, they had more responses that favored quantitative methods than responses that did not. In fact, the median among the participants was to answer 13 out of 18 questions in a way that favored that favored quantitative methods over softer methods. To further illustrate this point, a clear majority (76%) agreed with the statement:

> **"Cybersecurity should eventually adopt more sophisticated probabilistic methods based on actuarial methods where it has not already done so."**

- Although a minority had more responses less favorable to quantitative methods, none of the 171 participants answered all 18 questions in a manner that were less favorable to quantitative methods. There were, however, four who chose the pro-quantitative response in all 18 questions.

- While most agreed with pro-quantitative statements, most also allowed for the possibility that softer methods add some value. For example, 64% agreed with the statement:

> **"Commonly used ordinal scales help us develop consensus for action."**

- More years of experience was typically associated with more favorable attitudes towards quantitative methods. For example, 34% of those who responded more favorably to quantitative methods had over 16 years of experience whereas only 20% of those who opposed quantitative methods had that much experience.

- Not surprisingly, we also found there is a strong correlation between participants' attitudes and the methods they use in their organizations. For example, of those using Monte Carlo simulations, 86% were pro-quantitative. Similarly, people who were pro-quantitative made up 73% of the individuals who use Bayesian methods.[2] (*Figure 2*).

[2] Note that we make no claim about a causal relationship, here. It could be that those who support quantitative methods convince their organizations to adopt these tools or those that use these tools later gain a more favorable view of quantitative methods.

Figure 2: Attitudes towards Quantitative Methods by Methods Used

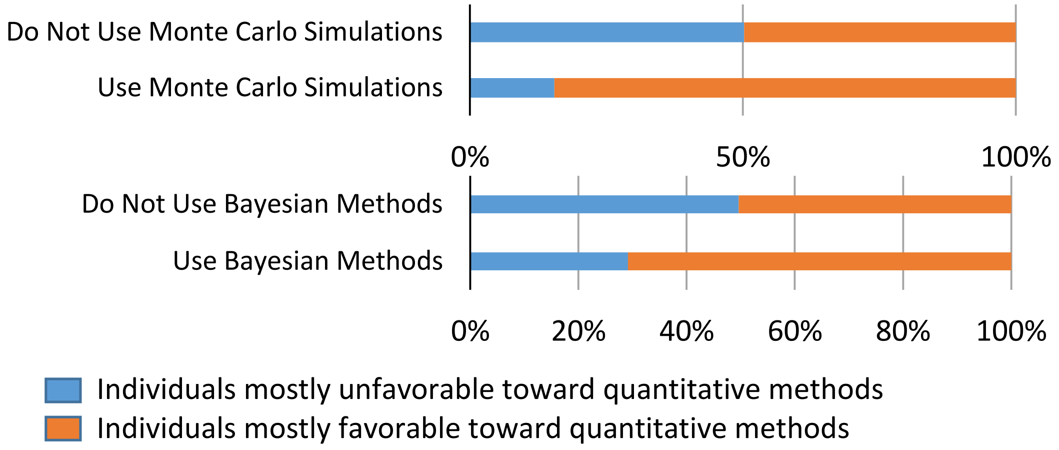

A Connection Between Attitudes and Statistical Literacy

To put the opinions about the use of quantitative methods in context, it would be important to assess what the respondents actually know about basic concepts in probability and statistics. Previous research with students found that attitudes about statistics were related to a limited understanding of statistics.[3,4,5] We were interested in whether we would observe something similar with professionals who mostly already have college degrees and many of which would have successfully completed at least one statistics course (although any formal study would often be many years in the past).

As part of the survey, participants were given a short, multiple choice test consisting of 10 questions about probability and statistics. Some of the questions were based on questions used in previous research to evaluate common misconceptions. Other questions were simply testing understanding of basic concepts like “statistical significance.” Some of the questions involved a simple calculation but, since they were multiple choice, if the respondent was within an order of magnitude or directionally correct they would get the correct answer.

[3] Lalonde, Richard N.; Gardner, Robert C. (1993). Statistics as a second language? A model for predicting performance in psychology students. *Canadian Journal of Behavioural Science/Revue canadienne des sciences du comportement*, 25(1),108-125.

[4] Schutz, P.A., Drogosz, L.M., White, V.E., & Distefano, C. (1998). Prior knowledge, attitude, and strategy use in an introduction to statistics course. *Learning and Individual Differences*, 10(4), 291-308.

[5] Perepiczka, M., Changler, N., Becerra, M. (2011). Relationship between graduate students' statistics self-efficacy, statistics anxiety, attitude toward statistics, and social support. *The Professional Counselor*, 1(2), 99-108.

Of the 171 respondents, 133 completed the quiz. Below is a summary of key findings:

- There is a statistically significant correlation between negative opinion about quantitative methods in cybersecurity and poorer performance on the quiz.
- Those who had more negative views toward quantitative methods were also much more likely than others to skip the statistics quiz.
- There is evidence that a specific type of misconception may explain much of the resistance to quantitative methods and poor performance on the quiz.

Figure 3 below shows the distribution of correct answers out of 10 questions on the quiz.

Figure 3: Variation in Statistics Literacy Quiz Performance

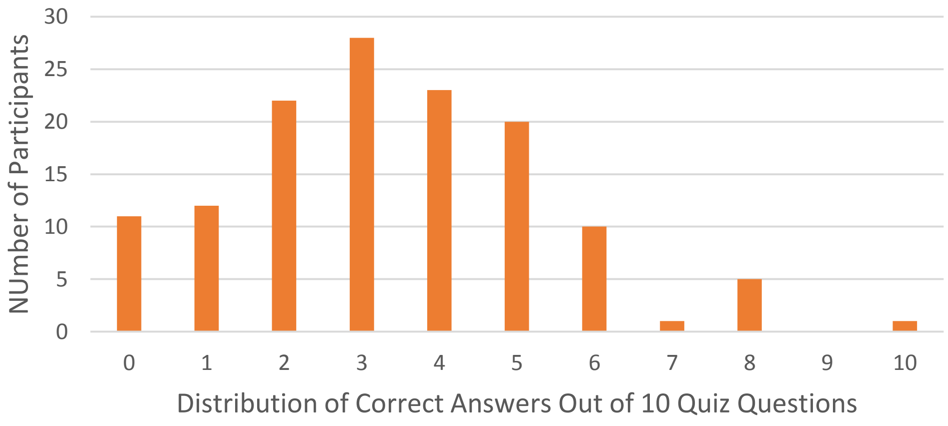

While ten questions may at first seem too few to be useful, the actual statistical analysis of results says otherwise. First, some individual differences in performance were so large that even a 10 question quiz is sufficient to detect a real difference, making a “flatter” distribution of scores than we would expect by chance alone. In other words, if there were no real differences in skills and all variation was due to chance in this small sample of questions, then we would not have observed as many scores of “0” and as many scores above “5” as we did. Second, and more importantly, our objective is not to get an accurate assessment of each individual’s understanding of these concepts. We are primarily concerned with larger patterns across all participants which are still apparent even when individual test error is large.

The median score was 3 out of 10 and the average was 3.4. Given the number of possible answers in each question, randomly picking any answer other than "I don't know" would have averaged 2.27 out of 10 correct, which is only slightly less than the median. The average is pulled up by the occurrence of more high scores than chance alone would allow. For example, there were 60 out of 133 that that scored 4 or higher while merely guessing would have allowed for only 21 scores of 4 or higher. So, even this small quiz shows evidence of understanding of statistical concepts by at least some cybersecurity experts.

These may seem like low scores but, again, the quiz focused on questions that were intended to detect common misconceptions about probabilities and statistics. In fact, it is interesting that, overall, in 6 of the 10 questions the most common answer (not necessarily the majority) chosen for each question was the correct answer and in 2 of the 10 the correct answer was the majority answer.

Of the 4 questions where the correct answer was not chosen more often than any of the other answers, 3 of them were based on previous published research assessing particularly widely held but incorrect beliefs about fundamental concepts in statistics and probability.[6,7] The cybersecurity experts did slightly better than the general public on those questions. As we would expect, there is evidence of understanding of key concepts of probability and statistics by these professionals.

When we compare the quiz performance to attitudes toward statistical methods, our results were consistent with previous research indicating a relationship between attitudes and skills. People who favored the use of quantitative methods generally performed better on the quiz than those who were against the use of quantitative methods. The results are statistically significant.[8] For a discussion of the methods we used, see Appendix B.

We also observed that those who were more experienced in cybersecurity performed slightly better on the quiz. Also, more experience together with higher quiz performance was an even better predictor of positive attitudes

[6] Conjunction fallacy, hospital fallacy and birthday problem. Tversky, Amos; Kahneman, Daniel (1974), "Judgments Under Uncertainty: Heuristics and Biases" (PDF), *Science* 185 (4157): 1124–1131.

[7] Tversky, A. and Kahneman, D. (1982) "Judgments of and by representativeness". In D. Kahneman, P. Slovic & A. Tversky (Eds.), *Judgment under uncertainty: Heuristics and biases*. Cambridge, UK: Cambridge University Press.

[8] Attitudes towards quantitative methods were aggregated such that pro-quantitative items were coded positively and anti-quantitative items were coded negatively. Thus attitudes could range from -18 (extremely anti-quantitative) to 18 (extremely pro-quantitative). This correlation with the number of correct responses on the quiz was statistically significant, *with* $p < 0.01$.

toward quantitative methods than quiz performance alone.[9] For example, if someone did above average on the quiz and had more than 16 years of experience then they were more likely to support the use of quantitative methods than most and even more likely than most who did just as well on the quiz. (This might be surprising to those expecting that it was the younger, less experienced, who were more likely to embrace methods different from what is currently used.)

This result is even more pronounced when we limit the analysis to questions that are more clearly pro-quantitative or anti-quantitative.[10] For example, agreeing with "Ordinal Scales are uninformative and add error to security decision making" is clearly a contradictory position with the statement "Ordinal scales must be used because probabilistic methods are not possible in cybersecurity" However, agreeing with statements like "Ordinal scales are better than nothing" could be held even by those that were generally more accepting of quantitative methods.

To illustrate this finding, *Figure 4* shows the distribution of participants who were at or below the median on the subset of pro-quantitative responses vs. those above the median and those at or below the median quiz score vs. those above the median. As the chart shows, those who were less favorable (than the median on this subset of questions) were less likely to get a quiz score above the median. Conversely, the majority of respondents who scored above the median in statistics literacy embraced the transition to quantitative methodology.

Figure 4: Literacy versus Attitude Median Split

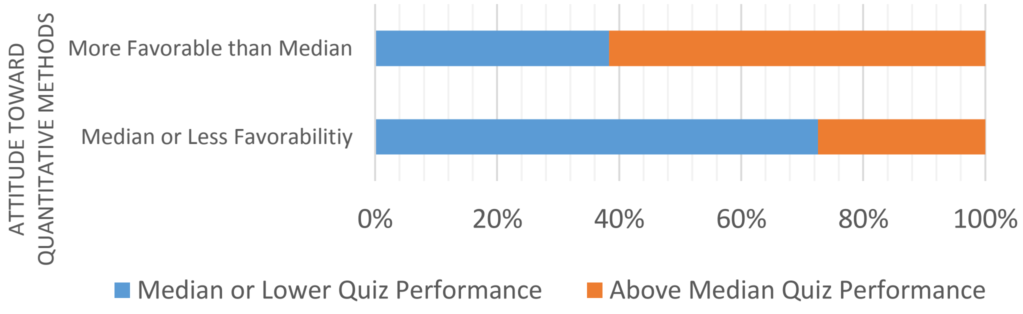

[9] Using hierarchical linear regression to predict attitudes towards quantitative methods, the addition of numbers of years worked in cybersecurity explained an additional 8% of variance above performance on the statistics quiz, $\Delta R^2 = 0.08$, $F(1, 129) = 10.88$, $p = 0.001$.

[10] Leaving out more neutral statements produced p<.001 as opposed to p<.01 using all questions.

Remember, the median score on the quiz was getting a correct response on 3 out of 10 questions, which is only slightly higher than what we would expect to see if they were randomly choosing any answer other than “I don’t know.” The fact that so many did worse than randomly guessing indicates that certain misconceptions are common – that is, they are more likely to choose incorrect answers than they would have simply by guessing.

We also found that 19% of survey participants who scored below the median on supporting quantitative methods elected to skip the statistics quiz in the survey. This is double the amount of pro-quantitative participants who elected to skip the quiz. This effect becomes even more pronounced when we compared the least pro-quantitative quartile with the most pro-quantitative quartile, where the first group was 3.3 times as likely to skip the quiz as the latter group. The different response rates for the statistics literacy quiz may have, if anything, resulted in an underestimation of the relationship between statistics literacy and acceptance of quantitative methods. Some participants may have opted out of the quiz only because it was perceived to be time consuming, however, it also seems like a reasonable possibility that those who knew they would perform very poorly on the quiz were less likely to take it.

The item on the attitudes-towards-quantitative-methods section that was the best predictor of performance on the quiz may be revealing about the source of both unfavorable attitudes toward quantitative methods and poor quiz performance. The survey item in question is whether the respondent agreed or disagreed with the following statement:

> **“Probabilistic methods are impractical because probabilities need exact data to be computed and we don't have exact data.”**

24% of those who responded agreed with this statement. Agreeing with this statement had a statistically significant correlation to poor quiz performance.[11] Those who disagreed with the statement scored above the median on the quiz three times as often as those who agreed. Also, those who agreed with this statement were 2.4 times as likely to skip the quiz altogether as those who disagreed with the statement.

The statement actually represents a common but incorrect understanding of probabilistic methods. In fact, it contradicts an important school of thought in statistics related to the use of Bayesian methods which quantify the uncertainty of an observer both before and after receiving new data. These methods are

[11] While the correlation may seem modest (r=0.27), even this level of effect would have been unlikely due to chance alone (p-value<.002).

often most useful where data is particularly limited or less than perfectly reliable. In other words, such methods are useful specifically *because* we lack "exact data," not in spite of it.

Statistics Quiz Performance vs. Self-Assessment

Participants were also asked for a self-assessment of how well they understand statistics relative to their peers in cybersecurity and we found that those who performed poorly on the quiz may have unrealistic self-assessments in statistical literacy. There is a tendency for individuals to overrate themselves in a variety of skills relative to their peers (driving skills, intelligence, etc.) and this is referred to in psychological literature as the Dunning-Kruger effect.[12] We were curious whether this tendency applies to cybersecurity professionals and whether this is related to attitudes or skills.

We found a slight tendency for the entire group to overrate themselves in statistics skills but those who scored worse on the quiz were more likely to overrate their skills. 63% of those who fell below the median in statistics literacy wrongly identified themselves as having at or above average proficiency in statistics (A slight tendency to overate themselves also appears in self-assessments of general cybersecurity skills. For example, 17% of the respondents rated themselves in the top five percent of cybersecurity professionals.)

In the quiz, the "Not enough information is provided to answer the question" was used 19% of the time even though this was not actually the correct answer on any of the questions. Many of the poor performers on the quiz used this answer frequently. Respondents may have believed that they could have answered it if only they were given more information. However, in subsequent conversations with some respondents who chose this answer frequently, none were actually able to articulate what specific information would be required to compute the correct answer or how they would have computed the answer given that information. Likewise, since 66% of all questions were answered incorrectly, it would appear that a group with realistic assessments of their own skills would have used the "I don't know" response very often – in fact, it was used only 13% of the time.

The overall tendency to overrate their own statistical knowledge is about the same for those who are more pro-quantitative than the median as it is for those who are less favorable toward quantitative methods. However, some indication

[12] Kruger, J., & Dunning, D. (1999). Unskilled and unaware of it: how difficulties in recognizing one's own incompetence lead to inflated self-assessments. *Journal of Personality and Social Psychology*, *77*(6), 1121.

of a difference does appear when we look at more extreme cases of opinions against the use of quantitative methods. Of the 15 most anti-quantitative respondents, 6 (40%) gave themselves a higher rating than their quiz results would justify whereas only 28% of those who were more favorable to quantitative methods overrated themselves. The size of the sample and the effect is too small for a statistically significant result but we can also compute that it is at least more likely than not that those who are the most unfavorable toward statistical methods overrate their knowledge in that area more than those with more favorable views.[13] If we observed this with these few data points in isolation, we will simply say that these results are inconclusive. But this does point to an interesting possibility when combined with previous evidence such regarding common misconceptions about probabilistic methods and the Dunning-Kruger effect. Not only is there a shortcoming of knowledge of basic statistical concepts but probably also lack of awareness that there is any such shortcoming.

Possible Relationships to Breach Frequency

Although the purpose for the survey was not to assess the influence of risk assessment methods on the probability of a security breach, there are two findings in this regard worth mentioning.

- Those who said they attempt to compute probabilities of cybersecurity events appear to experience fewer breaches than those who don’t compute probabilities. However, important caveats apply to this finding.

- Those who rely only on threat reporting services and web scanning for risk assessments are more likely to experience breaches than those who do not. The same caveats apply to this finding as above.

The first finding above is supported by the fact that those who answered NO to “We are able to compute probability of various levels of losses for the organization” also were significantly more likely to claim they experienced a breach compared to those who answered YES. Using the same kind of method we used to assess the chance of overrating one’s skills with small samples, we cross-referenced participants who reported breaches to whether they answered

[13] While not conclusive (p-value=0.14), a Bayesian analysis appropriate for small samples shows that it is more likely than not that overrating one’s statistics literacy is more common among those with unfavorable views toward quantitative methods. The Bayesian analysis explained in Appendix B shows an 85% chance this is true.

YES or NO to the above question to produce two distributions of the estimate of breach frequency for each group.[14] (*Figure 6*)

Figure 6: Likelihood of Breach Rate by Use of Probabilistic Methods

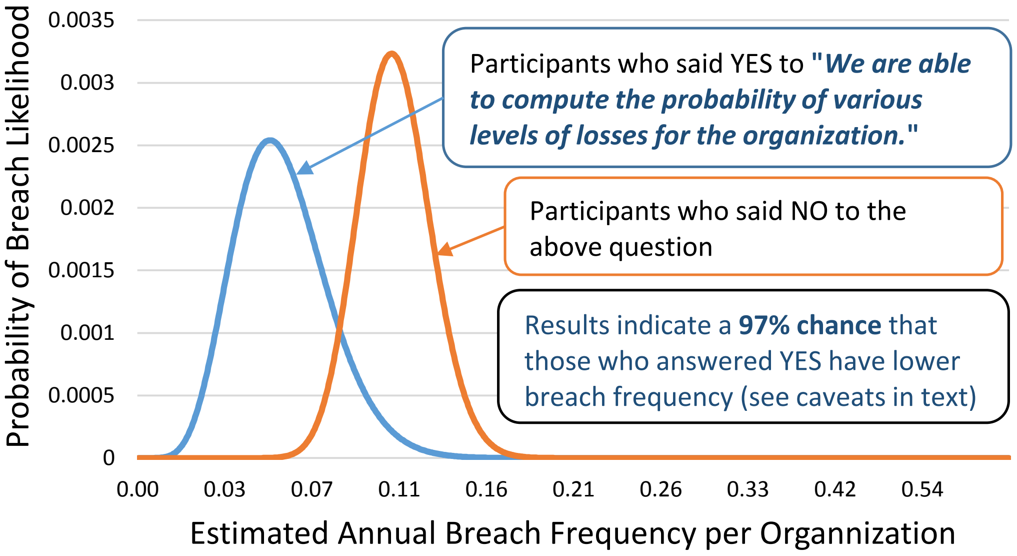

Participants who said YES to "***We are able to compute the probability of various levels of losses for the organization.***"

Participants who said NO to the above question

Results indicate a **97% chance** that those who answered YES have lower breach frequency (see caveats in text)

These two distributions represent our uncertainty about the breach frequency of each group based on responses. The two distributions show an overlap indicating it is possible that there is no difference between the groups. But simulation of possible breach frequencies using these distributions indicates a 97% chance that the breach frequency is higher for the group that states they do not compute probabilities of events.

Although these are important findings to note, keep in mind that the survey questions were not designed to be specific enough to confirm this relationship and we cannot make a claim of a cause-and-effect relationship. For example, what is considered a "breach" may not be consistent. However, if understanding of the meaning of "breach" varied randomly, then that alone would not explain the observed differences in responses. Another possible conflating factor is when the breach occurred relative to when methods were adopted. It could be possible that methods changed after a breach occurred and that would have

[14] Of the 118 who reported they do not compute these probabilities, 34 stated they had a breach in the previous 36 months. Of the 44 who said they do compute probabilities, 6 stated they had a breach in 36 months. Using a beta distribution to estimate population frequencies with these small samples, we start with uninformative uniform priors and update the distributions with the observed responses. The two distributions are compared in a simulation to show how often one would have a higher frequency than the other.

bearing on these results. (However, this kind of reporting bias could work either way.)

We will simply note that this is an interesting observation which would not, in isolation, be sufficient to indicate that quantitative methods are preferable. However, this finding is far from being an isolated observation - it is entirely consistent with other research that shows that the adoption of quantitative methods actually improves estimations and decisions.[15,16,17,18] Previously published research points out serious shortcomings of ordinal scales, qualitative evaluations of likelihoods and heat maps.[19,20,21]

On a side note, we also found that we make two additional observations about breach frequency. (Note however, the same caveats would apply to this observation as the previous observation about breaches as a function of computing probabilities.):

- Exclusively basing risk assessments on vulnerability reporting services and web scanners had a significant influence on the probability of a breach. That is, when the participant responded YES to “Security metrics purely based on vulnerability reporting services or web scanners,” the expected annual breach rate was nearly double as compared to participants who responded NO. Using the same method as the one used above (the comparison of beta distribution estimates of breach frequency) we find a 99% chance that those who used only reporting services or web scanners as metrics experience more breaches than those who use other methods as well.

- There were about twice as many firms who had breaches but did not report them as firms who had breaches and reported them (17% vs. 8%). This would have implications for any study which uses publically available breach data to assess breach risk. In this study, however, we counted both reported and unreported breaches when comparing breach rates to methods used.

[15] Paul E. Meehl, *Clinical versus Statistical Prediction; a Theoretical Analysis and a Review of the Evidence*, University of Minnesota Press, 1954.

[16] William M. Grove et al., "Clinical versus Mechanical Prediction: A Meta-Analysis," *Psychological Assessment* 12, no. 1 (2000): 19-30.

[17] William M. Grove et al., "Comparative Efficiency of Informal (Subjective, Impressionistic) and Formal (Mechanical, Algorithmic) Prediction Procedures: The Clinical-Statistical Controversy," *Psychology, Public Policy, and Law* 2 (1996): 293-323.

[18] R.M. Dawes, et al., "Clinical versus Actuarial Judgment," *Science* (1989), doi: 10.1126/science.2648573.

[19] L. A. Cox Jr., “What’s Wrong with Risk Matrices?” *Risk Analysis* 28, no. 2 (2008): 497– 512.

[20] P. Thomas, R. Bratvold, and J. E. Bickel, “The Risk of Using Risk Matrices,” *Society of Petroleum Engineers Economics & Management* 6, no. 2 (April 2014): 56–66.

[21] D. V. Budescu, S. Broomell, and H. Por, “Improving Communication of Uncertainty in the Reports of the Intergovernmental Panel on Climate Change,” *Psychological Science* 20, no. 3 (2009): 299–308.

CONCLUSIONS & RECOMMENDATIONS

Summarizing The Survey

The purpose of this survey was to gather information regarding the prevalence, acceptance, implementation, and knowledge of quantitative methods among the information security community as it relates to assessing risk of potential data breaches.

- Quantitative methods that use probabilities (like Monte Carlo simulations, Bayesian methods, statistical analyses of past events, etc.) are adopted only by a minority. In other words, the most popular methods currently in use to assess risks are not what the majority would prefer.

- There appears to be strong support for adopting more quantitative methods but limited understanding and certain misconceptions about probabilistic methods may be a source of resistance.

This would not be an issue if there were no significant difference in the performance of quantitative and qualitative methods. But, given the previously mentioned evidence, we cannot assume this. Indeed, we see clues that probabilistic methods are better than qualitative, consistent with previous research. We can say (with some caveats already mentioned) that we find some strong statistical evidence that the use of quantitative methods may reduce breach risks through better risk management. While this survey was not designed to detect this effect, the findings are at least consistent with other published research showing that probabilistic methods outperform expert intuition alone and that the most popular qualitative methods introduce errors even to intuition.

Most survey participants (67%) work for organizations with six or more individuals in information security. This means that even for those who are more supportive of quantitative methods, the majority still work for organizations where at least some on their team are more negative toward quantitative methods. The existence of at least some resistance in most organizations may hinder broader adoption of quantitative methods. It is also the case that the majority of those who are pro-quantitative still see some benefit for softer methods that are inconsistent with proper mathematical solutions. It is also true that even the supporters of more quantitative methods in some responses will show a conflicted stance toward quantitative methods. The fact that they see some benefit for some softer methods may indicate that they see less benefit in completely moving away from purely qualitative methods.

Suggested Actions

The growing importance of the issue of cybersecurity means that we should not take the effectiveness of any risk assessment method for granted – whether it is quantitative or qualitative. Based on the evidence so far, it appears more likely that firms which objectively investigate the performance of their methods based on empirical research will find that they should move toward more quantitative methods based on probabilistic models. If that is the case, then these findings point toward some specific strategies for making this transition.

First, **directly address the elephant in the room with training.** The fact that is that the skepticism of the minority who are less favorable toward quantitative methods is based at least partly on misinformation about quantitative methods. According to these results, it may not just be poor understanding of statistical concepts that gets in the way of more acceptance of quantitative methods, *but the incorrect belief that they do understand quantitative methods.* Such individuals will believe they understand quantitative methods well enough to have an informed position on their effectiveness, feasibility or relevance to cybersecurity risk assessment when in reality their beliefs are based on common misconceptions. Training will be required to overcome beliefs that contradict the actual foundations of probabilistic models as well as the beliefs that softer methods somehow alleviate the complexities of assessing risks in cybersecurity. While we cannot conclude that merely being exposed to more quantitative methods *causes* an increase acceptance of them, our findings do indicate that exposure and acceptance are strongly related. Firms should at least consider this possibility that exposure through training will improve acceptance.

Second, **consider leveraging the more experienced in the team to explain this.** Although not true in every case, in most cases we see that it is the less experienced that have more negative views toward quantitative methods and the more experienced who indicate and interest in the adoption of more quantitative methods. It will also usually be the case in firms with multiple cybersecurity specialists that interest in using more quantitative methods is actually the majority opinion. They may simply not be aware of how to implement such methods in practice.

Finally, **adopt a scientific approach to measuring the approach to your own risk assessment methods.** That is, start with what the existing published research says about the measured performance of different methods and develop a method based on those findings. Then, track the performance of your method using the methods that the (now-trained) staff are familiar with for assessing probabilities of rare events with limited information.

Summary

- **Advanced quantitative methods make up the smallest percentage of methods being used in the cybersecurity industry.**
- **Pro-quantitative individuals use more sophisticated quantitative methods than anti-quantitative individuals.**
- **More experience in cybersecurity is associated with positive attitudes towards quantitative methods.**
- **Individuals with opinions less favorable to quantitative methods are more likely to avoid the statistics quiz in the survey.**
- **Positive attitudes towards quantitative methods was associated with higher statistical literacy.**

ABOUT THE AUTHORS

Dr. Jim Clinton is a Senior Quantitative Analyst at Hubbard Decision Research. He brings to HDR a background in quantitative psychology and years of experience conducting experimental research in the areas of cognitive and social psychology. His psychological experiments entail advanced quantitative methods such as multinomial logistic regression and linear mixed-effects modeling. His work has been funded by the National Science Foundation, and he is an active member of the Psychonomic Society, Midwestern Psychological Association, Association for Psychological Science, amongst many others. His work has been published in peer-reviewed journals such as Projections, featured in

blogs about cognition and marketing, and he has served as a reviewer for peer-reviewed journals such as Memory & Cognition, Media Psychology, and Discourse Processes.

Douglas Hubbard is the inventor of the Applied Information Economics (AIE) method and founder of Hubbard Decision Research (HDR). He is the author of *How to Measure Anything: Finding the Value of Intangibles in Business*, *The Failure of Risk Management: Why It's Broken and How to Fix It*, and his latest book *Pulse: The New Science of Harnessing Internet Buzz to Track Threats and Opportunities*. In addition to his books, Mr. Hubbard has been published in *Nature*, *The IBM Journal of R&D*, *OR/MS Today, Analytics*, *CIO*, and *Information Week.*

For over 28 years, Mr. Hubbard has applied quantitative methods for measurement and risk analysis for large, critical projects, investments and other management decisions for a variety of industries and government agencies. His AIE methodology, has received critical praise from The Gartner Group, The Giga Information Group, and Forrester Research. He is a popular speaker at IT metrics & economics conferences all over the world. Previous to starting his own consulting business, he was in Management Consulting Services in Coopers & Lybrand. He has an MBA from the University of South Dakota.

Andrew Triplett is a Senior Quantitative Analyst at Hubbard Decision Research. He brings to HDR a background in social and quantitative psychology, with a Masters in Applied Social Psychology from Loyola University Chicago. He is now working on his Doctorate in Applied Social Psychology at Loyola University Chicago. Mr. Triplett has years of experience in conducting experimental research in cognitive and social psychology. Specifically, his research focuses on topics such as aggression, savoring, and emotion and cognition, all of which entail the use of a variety of advanced quantitative methodology. Andrew is an active member of the American Psychological Association, the Association for Psychological Science, the Midwestern Psychological Association, amongst many others. Mr. Triplett has articles in review for Aggression and Violent Behavior and Applied Research in Quality of Life and serves as a peer reviewer for Aggression and Violent Behavior and Oxford University Press.

Appendix A: Survey Questions Details and Distribution of Responses

Background Information on Participants and Their Organizations

Title within organization	Frequency	Percentage
Analyst, engineer or other individual contractor	49	29%
Mid-level manager (manage other analysts)	34	20%
CISO/VP (top information security person)	41	24%
Other	46	27%

Number of full-time employees in cyber/information security in organization	Frequency	Percentage
1 to 2	32	19%
3 to 5	20	12%
6 to 10	23	14%
11 to 20	16	10%
21 to 50	23	14%
51+	49	29%
I don't know	4	2%

Industries represented	Frequency	Percentage
Information security consulting services	48	28%
IT consulting services, not specifically information security	27	16%
Other professional consulting services	13	8%
Insurance, finance, banking, credit services	49	29%
Manufacturing, refining, production	13	8%
Resources extraction (mining, oil & gas, forestry, etc.)	3	2%
Retail goods (clothing, hardware, furniture, electronics, general merchandise, etc.)	6	4%
Retail food service, hospitality recreation, or entertainment	2	1%
Telecommunications	9	5%
Transportation or Delivery including road, air, rail, and sea	4	2%
Media, publishing, broadcasting	3	2%
Government (including national, state, municipal)	15	9%
Healthcare services including hospitals, ambulatory, nursing homes, and HMOs	13	8%
Education	11	6%
Agriculture, specifically animal and crop production	1	1%
Construction	3	2%
Wholesale, Warehousing distribution	4	2%

Number of years worked in information security	Frequency	Percentage
0 to 3	12	7%
3 to 9	49	29%
10 to 15	54	32%
16 to 20	30	18%
20+	23	14%

Total years of experience	Frequency	Percentage
0 to 3	2	1%
3 to 9	16	10%
10 to 15	26	16%
16 to 20	41	25%
20+	82	49%

Participation in sharing internal data	Frequency	Percentage
We have shared security incidents data with government entities.	78	46%
We have shared security incidents data with other firms.	35	21%
We have shared security incidents data with select colleagues in our industry or field.	73	43%
I don't know	39	23%

Total number of employees in firm	Frequency	Percentage
1 to 10	17	10%
11 to 50	11	7%
51 to 200	15	9%
201 to 500	11	7%
501 to 1,000	12	7%
1,001 to 5,000	28	17%
5,001 to 20,000	38	23%
20,001 to 50,000	11	7%
More than 50,000	26	15%

Organization specialization	Frequency	Percentage
My organization specializes in providing cybersecurity services.	47	28%
My organization does not specialize in providing cybersecurity services, but I work in a cybersecurity capacity within my organization.	95	57%
Neither is accurate.	26	15%

Security risk management frameworks used	Frequency	Percentage
NIST	101	59%
OCTAVE	6	4%
an ISO Standard	83	49%
ISACA's Risk IT	27	16%
OWASP	49	29%
Proprietary framework	51	30%
No security risk management framework	25	15%

Frequency of executive demands for financial arguments for cyber-security/information security investments	Frequency	Percentage
Always	40	24%
Usually	40	24%
Sometimes	30	18%
Rarely	26	15%
Never	6	4%
I don't know	26	15%

Experience prior to information security	Frequency	Percentage
I have always been in information security.	14	9%
I was in other areas of information technology prior to specializing in information security (e.g., software development, network admin, etc.) but never outside of information technology.	113	70%
I was in a field unrelated to information technology.	34	21%

Occurrence of a data breach in the last 36 months	Frequency	Percentage
No we did not experience a data breach in the last 36 months.	123	75%
Yes, and it was publicly reported.	14	8%
Yes, but it was not reported.	28	17%

Certifications held in information security	Frequency	Percentage
CISSP	79	46%
CISM	39	23%
GIAC	18	11%
Security+	12	7%

CEH	13	8%
Other	89	52%

Information security methods used	Frequency	Percentage
Security metrics program.	108	65%
Security metrics purely based on vulnerability reporting services or web scanners.	37	22%
Security events assigned likelihoods stated as a percentage probability (instead of a scale with values like "high", "medium", "unlikely", etc.)	64	39%
Potential losses communicated as ranges or probability distributions of dollar amounts (instead of a scale with values like "high", or "insignificant", etc.)	55	33%
Compute the probability of various levels of losses for the organization.	46	28%
Monte Carlo simulations for assessing information security risks.	22	13%
Bayesian networks or other Bayesian methods for assessing information security risks.	23	14%
Regression models on historical data to assess information security risks.	37	22%
Heat maps/risk maps/ risk matrices to communicate risks.	102	61%
FAIR method.	31	19%
Analytic Hierarchy Process (AHP) to assess information security risks.	15	9%
Prioritize controls based on computed return on investment (the value of the risk mitigation compared to the cost of the mitigation).	59	35%

Perceived time spent on risk assessment related to cybersecurity	Frequency	Percentage
Too much time.	4	3%
Too little time.	81	53%
About the right amount.	52	34%
I have no opinion on the matter.	17	11%

Self-perception in cyber security/information security skills compared to peers	Frequency	Percentage
I am one of the best. I rate myself in the top 5% of my peers.	26	17%
I would rate myself in the top 30% but not as high as the top 5% of my peers.	61	39%
I understand it about as well as my peers. I put myself somewhere around the median of my peers on this topic.	45	29%
I would rate myself in the bottom 30% but not as low as the bottom 5% of my peers.	11	7%
My understanding of it is much less than my peers. I rate myself in the bottom 5% of my peers.	3	2%
I prefer not to answer.	9	6%

Attitudes Regarding Quantitative and Qualitative Methods

Statement	#Agree	% Agree	# Disagree	% Disagree	# Don't Know/No Opinion	% Don't Know/No Opinion
Evaluating risk using ordinal scales (such as "high" or "medium" or "low") improves decisions about information security.	74	48%	63	41%	18	12%
Ordinal Scales are uninformative and add error to security decision making.	75	49%	60	39%	19	12%
Ordinal scales must be used because probabilistic methods are not possible in cybersecurity.	28	18%	102	66%	24	16%
Ordinal scales used in most information risk assessments are better than quantitative because they are easy to understand and explain.	47	31%	85	55%	22	14%
Commonly used ordinal scales help us develop consensus for action.	99	64%	30	20%	25	16%
Information security risks can be assessed using quantitative methods.	134	87%	6	4%	14	9%
For assessing information security risks, ordinal scales are better than nothing.	114	75%	20	13%	19	12%
Probabilistic methods are impractical because probabilities need exact data to be computed and we don't have exact data.	37	24%	100	66%	15	10%
Each security event and each organization is different so quantitative statistical methods cannot be applied.	17	11%	123	80%	13	9%
There is no way to calculate a range of the intangible effects of major risks like damage to reputation.	21	14%	114	75%	18	12%

Statement	#Agree	% Agree	# Disagree	% Disagree	# Don’t Know/No Opinion	% Don't Know/No Opinion
Management and users won't understand the quantitative methods output.	38	25%	96	63%	19	12%
Quantitative methods don't apply because risk is ultimately subjective.	15	10%	122	80%	16	10%
I have argued against the use of quantitative, probabilistic methods in cybersecurity risk assessments.	21	14%	107	70%	25	16%
Quantitative methods don’t apply in situations where there are human agents that act unpredictably.	19	13%	112	74%	21	14%
An expert using quantitative probabilistic methods will do better risk assessments than an expert using intuition alone.	97	63%	21	14%	35	23%
Ordinal scales or qualitative methods alleviate the problems with quantitative methods.	44	29%	75	49%	34	22%
Information security is too complex to model with probabilistic methods.	13	9%	116	76%	23	15%
Cybersecurity should eventually adopt more sophisticated probabilistic methods based on actuarial methods where it has not already done so.	115	76%	13	9%	23	15%

Statistical Literacy Questions

Imagine you have a portfolio of systems, many of which are targets of a particular type of attack. How many attacks would have to be witnessed in this portfolio of systems in order for you to have a "statistically significant" sample size of the frequency of such attacks?	Frequency	Percentage
Hundreds of data points are required to be "statistically significant."	7	5%
Less than 200 but more than 30 data points are required to be "statistically significant."	24	18%
It is possible for less than 10 data points to produce a "statistically significant" result.*	51	39%
Not enough information is given to adequately answer the question.	35	27%
I don't know.	14	11%
***correct answer**		

Assume that in 2013 a risk analyst assessed a 5% chance that a security breach would occur for her organization in 2014 which would result in compromising personal health data of employees. Assume that such an event did occur for this organization in 2014. Which of the following is true?	Frequency	Percentage
The original probability assessment of 5% must have been incorrect.	2	2%
The original probability assessment of 5% was correct.	13	10%
A single instance does not by itself indicate whether the 5% probability was correct or not.*	87	66%
Such an event proves that probabilities are not appropriate for cybersecurity.	5	4%
There is insufficient information to answer the question.	17	13%
I don't know.	8	6%
***correct answer**		

Assume that you have a portfolio of systems for which you have observed no security events in the past year that resulted in a monetary or productivity loss, which of the following statements are true?	Frequency	Percentage
If no events were observed, then we have no data about the likelihood of these events.	3	2%
The fact that no events were observed tells us something about the likelihood of these events.*	50	38%
One year is not long enough time to gather enough observations to make an inference.	6	5%
Since some events may not have been observed, the lack of observed losses tells us nothing.	41	31%
There is insufficient information to answer the question.	22	17%
I don't know	9	7%
***correct answer**		

Assume X and Y are events with probabilities equal to 50% and 20%, respectively. These events are positively correlated. Which of the following is true?	Frequency	Percentage
The chance of both events occurring is 10%.	26	20%
The chance of both events occurring is greater than 10%.*	37	29%
The chance of both events occurring is the sum of the probabilities regardless of the correlation.	16	12%
The chance of both events occurring is less than 10%.	5	4%
Insufficient information is given to answer the question.	17	13%
I don't know the answer.	28	22%
***correct answer**		

Assume that a breach has occurred in 10 out of 61 companies in industry X and 3 times out of 97 in industry Y in the last two years. Assuming there will be no change in the risk over the next two years (e.g., no corrective action is taken, threats do not change, etc.) which of the following is true?	Frequency	Percentage
Insufficient data is provided to determine relative likelihoods of these breaches.	21	16%
The breach is more likely to happen again in industry Y than industry X.	6	5%
The breach is more likely to happen again in an industry X firm than in an industry Y firm.*	89	70%
I don't know.	12	9%
***correct answer**		

Assume the probability of an event, X, occurring in your firm sometime in 2016 is 20%. The probability of this event goes to 70% if threat T exists. There is a 10% probability that threat T exists. Which of the following statements is true?	Frequency	Percentage
If the threat T does not exist, the probability of the event X must be less than 20%.*	21	16%
If the event X does not occur, then T does not exist.	4	3%
Given that the event X occurs, the probability that threat T exists must be greater than 50%	32	25%
There is insufficient information to answer the question.	33	26%
I don't know.	39	30%
***correct answer**		

John is a software developer and has a Master's degree in computer science. His peers consider him to be diligent and he enjoys his work. He is interested in information security. Which of the following is more likely? (Assume that neither is 100% likely and neither is 0% likely)	Frequency	Percentage
John enjoys skydiving.*	20	16%
John enjoys skydiving and is pursuing a CISSP certification.	19	15%
There is insufficient information to answer the question.	72	56%
I don't know.	18	14%
***correct answer**		

Assume the HR departments from two cybersecurity firms give a cybersecurity proficiency test to new job applicants. It is known that only half of all applicants nationwide pass this test. Firm A gives this test to 20 applicants per week on average. Firm B gives the same test to 200 applicants per week on average. Which firm is more likely to have more than 60% of applicants in a given week pass the test?	Frequency	Percentage
Firm A*	26	20%
Firm B	7	5%
They both have the same probability (within about +/-5%).	74	58%
Insufficient information is given to answer the problem.	8	6%
I don't know.	13	10%
***correct answer**		

A firm of 20,000 employees has decided to assign randomly generated numeric passwords for each employee. The password will be 6 digits long, allowing for up to 1 million unique passwords. What is the chance that at least two employees will have the same password?	Frequency	Percentage
2% or less	55	43%
over 2% and less than 40%	13	10%
equal to or over 40% and less than 60%	3	2%
equal to or over 60% and less than 98%	6	5%
98% or more*	16	13%
There is insufficient information to answer the question.	7	6%
I don't know.	27	21%
***correct answer**		

A survey conducted by a university was testing a hypothesis that organizations with more CISSP certified security professionals had fewer breaches. Suppose they found no statistically significant result. Which of the following statements is true?	Frequency	Percentage
The survey had no bearing on the chance of whether the hypothesis was true.	21	17%
The survey results are proven to be a random fluke.	1	1%
The survey showed the hypothesis was not true.	25	20%
Statistical significance is not the same as whether the hypothesis is true given the observations.*	53	42%
There is insufficient information to answer the question.	13	10%
I don't know.	14	11%
***correct answer**		

Self-perception in probabilities and statistics compared to peers	Frequency	Percentage
I understand statistics and probability very well. I studied it and I remember the concepts well. I rate myself in the top 5% of my peers.	4	3%
I understand statistics and probability better than most. I had some exposure to the math behind it and I believe I recall the key concepts. I would rate myself in the top 30% but not as high as the top 5% of my peers.	25	20%
I understand statistics and probability about as well as my peers. I put myself somewhere around the middle of my peers on this topic.	65	51%
4My understanding of statistics and probability is below average compared to my peers. I would rate myself in the bottom 30% but not as low as the bottom 5% of my peers.	19	15%
My understanding of statistics and probability is much less than my peers. I rate myself in the bottom 5% of my peers.	10	8%
I prefer not to answer.	5	4%

Opinion of the relationship between "threat" and "capability"	Frequency	Percentage
Unless a threat has a capability to attack, we don't identify it as a threat.	33	27%
Threat and capability are independent - we can identify a threat when there is no capability and we can identify a capability when there is no threat.	62	50%
We never identify specific capabilities, but we identify threats.	14	11%
We don't specifically identify threats or capabilities.	2	2%
I have no opinion on the matter.	12	10%

Assessment strategy in the likelihood of an event	Frequency	Percentage
We assess a specific probability for an event occurring (e.g., There is a 5% chance of this occurring).	26	21%
We use a verbal scale (e.g., "extremely unlikely", "unlikely" etc.) but we are given ranges for the meaning of each word (e.g., "unlikely" is defined as a probability of 5% to 20%, a frequency of once in decade, etc.).	35	28%
We use a verbal scale but specific probabilities or frequencies are not defined.	19	15%
We use an ordinal point scale (e.g., 1 to 5) but we are given ranges for the meaning of each word (e.g., "unlikely" is defined as a probability of 5% to 20%, a frequency of once in decade, etc.).	26	21%
We use an ordinal point scale (e.g., 1 to 5) but specific probabilities or frequencies are not defined.	17	14%

Opinion of the relationship between "vulnerability" and "threat"	Frequency	Percentage
We do not identify a vulnerability unless there is a realistic threat.	8	7%
Vulnerability and threat are independent - if there is a vulnerability, but no threat to exploit it, we still identify vulnerability.	98	80%
We never identify specific vulnerabilities, but we identify threats.	2	2%
We never identify specific threats, but we identify vulnerabilities.	5	4%
We never identify specific threats nor do we identify specific vulnerabilities.	2	2%
I have no opinion on the matter.	7	6%

Appendix B: A (Brief) Discussion of the Statistical Methods Used

We used two different methods as a test of statistical validity. The most common practice used in scientific research is to compute a "p-value." This is the chance that we would see the observed outcome or something even more extreme if it were merely a random fluke. For example, if we were testing whether a coin is more likely to land on heads and we get 60 heads out of a 100 flips, we would compute the chance of 0.028 that a fair coin could have given us 60 or more heads out of 100. The 0.028 is the p-value.

If the p-value is very small, then we conclude that it was unlikely to observe that result by chance alone. The threshold we compare to the p-value is a "significance level" and if the p-value is less than the significance level then we declare the results are "statistically significant." It is common in psychology literature to use significance levels of no higher than .05 but even stricter standards (requiring p-values to be less than significance levels of .01 or even .001) are preferred.

Another method used is a type of "Bayesian" analysis where we use the data to update a prior probability belief. In this survey, we applied this to problems of estimating a population proportion such as whether there is a higher percentage of firms in one group experiencing a breach than in another.

In this case, we started with a "robust" or "uninformative" prior which makes the least assumptions before we look at the data. A belief that allows for a lot of uncertainty for a prior is simply that the population proportion must be somewhere between 0% and 100% where every point on this range is equally likely. This is a very conservative assumption since a population proportion must, by definition, be between 0% and 100%. With a statistical method based on the "beta distribution" we update this prior probability distribution with the data from the survey. This allows us to estimate what the true proportion based on the sample of the population. We can then compare two such distributions to compute the chance that one group is really different from another or if it were just a fluke.

Note that the p-value and the Bayesian method answer related but different questions. The p-value is the chance we would observe that data (or something even more extreme) if it were a fluke. The p-value is not the chance that a hypothesis is true (just the chance you would see those results by chance if the hypothesis weren't true). The Bayesian method, however, allows us to actually compute the chance a claim is true, but it requires that we state a prior probability.