Federal Reserve Bank of Philadelphia

Discussion Paper

Quantifying Cyber Risk in the Financial Services Industry

Larry Santucci
Federal Reserve Bank of Philadelphia
Consumer Finance Institute

DP 18-03
November 2018

Consumer Finance Institute

Quantifying Cyber Risk in the Financial Services Industry

Larry Santucci*
Federal Reserve Bank of Philadelphia

November 2018

Abstract

The Consumer Finance Institute hosted a workshop in February 2017 featuring James Fox, partner and principal at PricewaterhouseCoopers (PwC) and a leading authority on cybersecurity in the financial services industry. He discussed the importance of measuring cyber risk, highlighted some challenges that financial institutions face in measuring cyber risk, and assessed several leading cyber-risk management methodologies. Fox also provided some recommendations for bank exams and insights into how federal agencies might begin to quantify systemic cyber risk. This paper summarizes Fox's presentation and is supplemented by additional research.

Keywords: cyber risk, cybersecurity, risk appetite, risk quantification

JEL codes: G28, G32, K24, L14

* Larry Santucci is a senior research fellow in the Consumer Finance Institute at the Federal Reserve Bank of Philadelphia, Ten Independence Mall, Philadelphia, PA 19106-1574; larry.santucci@phil.frb.org.

I. Cybercrime and the Global Financial System

A rash of million dollar cybercrimes during the past three years has brought renewed attention to potential vulnerabilities in the global financial system, an unimaginably complex ecosystem of businesses, consumers, governments, and regulators. In February 2016, a group of hackers used an international payments messaging system to move $81 million from the Bank of Bangladesh's account at the Federal Reserve Bank of New York to accounts in the Philippines and Sri Lanka.[1] It was one of the most costly bank robberies in history.[2] The thieves employed some tactics that had also been deployed in a series of cyberattacks during 2015, including a $12 million theft from an Ecuadorean bank, a $6 million theft from a Russian bank, and an unsuccessful attempt to rob a small Vietnamese bank.[3] Then, in 2018, a corrupt employee of India's second-largest bank used similar tactics in an attempt to steal almost $2 million. The Mexican banking system lost another $18 million that year in a similar fashion.[4] These events demonstrate that cybercriminals have successfully accessed the global financial system via a single point of entry, gained access to one or more connected entities, and created enough chaos

[1] Joshua Hammer, "The Billion-Dollar Bank Job," *New York Times Magazine* (May 3, 2018).

[2] "SWIFT Action: Preventing the Next $100 Million Bank Robbery," PwC (June 2016), https://www.PwC.com/us/en/industries/financial-services/financial-crimes/library/swift-bangladesh-robbery-2016.html

[3] Gavin Finch, "Ecuador Bank Says It Lost $12 Million in Swift 2015 Cyber Hack," *Bloomberg* (May 20, 2016).

[4] Sudarshan Varadhan, "India Bank Hack 'Similar' to $81 Million Bangladesh Central Bank Heist," *Reuters* (February 19, 2018). The Russian central bank also reported that a Russian bank had been the victim of a $6 million cybertheft via the same payments messaging system during 2017. See Jack Stubbs, "Hackers Stole $6 Million from Russian Bank via SWIFT System: Central Bank," *Reuters* (February 16, 2018).

to disrupt the financial system. To paraphrase an executive vice president at the Federal Reserve Bank of New York, the financial sector is under attack.[5]

In another global cybersecurity incident in 2017, hundreds of thousands of computers worldwide were frozen by a massive ransomware attack that also compromised some Russian bank systems.[6] The so-called WannaCry malware exploited a known vulnerability in a Microsoft communication protocol to spread a type of self-propagating ransomware over the public Internet and through internal networks.[7] While the creators of WannaCry do not appear to have specifically targeted the financial services industry, other attacks have done just that. On December 8, 2010, a coordinated distributed denial of services (DDoS) attack disrupted both the MasterCard and Visa corporate websites.[8] In September 2012, several U.S. banks, including Bank of America, Citigroup, and Wells Fargo, experienced DDoS attacks to their online banking sites.[9] Consumers were unable to log in online to their banking accounts, send payments, or transfer money. No bank accounts were breached and no money was stolen, but the attacks again demonstrated a vulnerability in the financial system. Fintech companies and other nonbank financial companies are no less a target: An October 2016 DDoS attack disrupted several prominent websites, including PayPal, preventing some customers from making payments.[10]

---

[5] Sarah Dahlgren, "The Importance of Addressing Cybersecurity Risks in the Financial Sector," remarks at the OpRisk North America Annual Conference, New York City (March 24, 2015).

[6] Bill Chappell, "WannaCry Ransomware: What We Know Monday," *NPR* (May 15, 2017).

[7] John Miller and David Mainor, "WannaCry Ransomware Campaign: Threat Details and Risk Management," *FireEye* (May 15, 2017; update 3 on May 17, 2017); Bradley Mitchell, "Network Protocols," *Lifewire* (June 19, 2018).

[8] Aaron Smith, "MasterCard, Visa Targeted in Apparent Cyberattack," *CNN Money* (December 8, 2010).

[9] Nicole Perlroth and Quentin Hardy, "Bank Hacking Was the Work of Iranians, Officials Say," *New York Times* (January 8, 2013).

[10] Joseph Menn, Jim Finkle, and Dustin Volz, "Cyber Attacks Disrupt PayPal, Twitter, Other Sites," *Reuters* (October 21, 2016).

Regulators in the United States and abroad have recently accelerated efforts to protect the global financial system from the risks posed by cybercrime. The challenge of managing systemic risk in the financial system is complicated by the degree with which bank and nonbank financial firms are interconnected along with various nonfinancial service providers. The number of entry points into any systemically important financial institution or central bank is practically unlimited, as demonstrated by the Bank of Bangladesh heist. In September 2016, the New York State Department of Financial Services (DFS) proposed new cybersecurity regulations for banks and other financial institutions.[11] The new DFS regulations, which became effective on March 1, 2017, require nonexempt New York-regulated banks, insurance companies, and financial service companies (including New York-based branches and agencies of foreign banking organizations) to enact new cybersecurity measures. These measures include periodic risk assessments of information security systems and cybersecurity policies, annual penetration testing and semiannual vulnerability assessments, cybersecurity personnel who can keep up-to-date with the threat environment, and third-party service providers to adhere to minimum information security standards.[12]

In October 2016, the Board of Governors of the Federal Reserve System, the Federal Deposit Insurance Corporation, and the Office of the Comptroller of the Currency issued a joint advance notice of proposed rulemaking (ANPR). The agencies requested comments and suggestions on how best to enhance cyber-risk management standards for large financial institutions and third-party service providers with access to information systems or private customer data at regulated entities.[13] The agencies also expressed concern that, although the

[11] 23 NYCRR 500, "Cybersecurity Requirements for Financial Services Companies," New York State Department of Financial Services (last accessed on August 22, 2018).

[12] Refer to 23 NYCRR 500, Section 500.01, for the definition of *covered entity*. Exemptions are listed in Section 500.19.

[13] 81 FR 74315, "Enhanced Cyber Risk Management Standards," *Proposed Rules*, *Federal Register*, 81:207, pp. 74315–74326 (October 26, 2016).

financial services industry was particularly vulnerable to systemic cyber risk, they lacked a means of measuring overall risk exposure and the risk posed by any single institution. Recognizing that the first step toward quantifying the industry's exposure to systemic risk was to measure risk at the institution level, the ANPR requested comments on methodologies that could be used to quantify institutional cyber risk to make it comparable across the industry.

To understand some of the challenges facing regulators and industry participants seeking to quantify cyber risk, on February 23, 2017, the Consumer Finance Institute at the Federal Reserve Bank of Philadelphia hosted a workshop on measuring cyber risk with James Fox, partner and principal at PricewaterhouseCoopers (PwC) and a leading authority on cybersecurity in the financial services industry.[14] Fox discussed the importance of measuring cyber risk, highlighted some challenges that businesses face in measuring cyber risk, and assessed several leading cyber-risk management methodologies. Fox also provided some recommendations for bank exams and insights into how federal agencies might begin to quantify systemic cyber risk.

This paper is organized as follows. Section II introduces the concept of cyber risk and discusses the value and challenges of measuring it at the institution level. Section III summarizes the highlights of Fox's presentation. The paper concludes in Section IV.

II. What Is Cyber Risk and How Is It Measured?

Businesses are targeted by a host of different cybercriminals for a variety of reasons. Malicious insiders, transnational organized crime rings, foreign intelligence services, competitors, and hacktivists have all attempted to access corporate networks.[15] These groups

[14] This event was organized by the Payment Cards Center, which was later renamed the Consumer Finance Institute.

[15] "Who Is Stealing Your Trade Secrets? An Overview of Key Threats," Center for Responsible Enterprise and Trade, Create.org (September 29, 2015).

often disrupt business operations and steal trade secrets to obtain power, influence, or profit.[16] As the number and sophistication of cyberthreat actors evolves, so too does the need for a systematic approach to managing *cyber risk*. The Institute of Risk Management defines *cyber risk* as any risk of financial loss, disruption, or damage to the reputation of an organization from a failure of its information technology systems.[17] Cyber risk may arise unintentionally (i.e., an accidental destruction of data or intellectual property) or intentionally and maliciously (i.e., a cyber attack).[18] According to the Federal Financial Institutions Examination Council (FFIEC), a *cyber attack* is an attempt to damage, disrupt, or gain unauthorized access to a computer, system, or network.[19] Unauthorized access can lead to the destruction or theft of confidential or sensitive information or loss of control over internal computing systems or customer-facing websites. It can also provide cyberattackers with a gateway to other businesses via shared networks or common entry points.[20] The possibility that an isolated cyberattack could have consequences for the entire financial system is referred to as *systemic risk*.[21]

The increasing prevalence of cyberattacks has catapulted cyber risk to prominence among business risks, prompting companies to carefully consider how current business operations may be inadvertently creating vulnerabilities that could expose the company to attack. For example, it may be necessary to reconsider the processes and protocol for engaging with

[16] "Strategy to Combat Transnational Organized Crime," The White House (July 19, 2011).

[17] "Cyber Risk and Risk Management," Institute of Risk Management (last accessed October 17, 2018).

[18] Saffet G. Ozdemir, "Non-Malicious Destruction of Data," SANS Institute, SANS Security Essentials Version 1.2f (2001).

[19] "Information Security," FFIEC Information Technology Examination Handbook (September 2016).

[20] "Understanding Systemic Cyber Risk," World Economic Forum, *White Paper*, REF 181016 (October 2016).

[21] "Understanding Systemic Cyber Risk," World Economic Forum, *White Paper*, REF 181016 (October 2016).

third-party service providers, vetting potential employees, and managing internal systems and information, any of which could create a vulnerability.

A. The Tools of Cyber-Risk Measurement

When quantifying cyber-risk exposure at the enterprise level, businesses often start by enumerating their *threats* — anything capable of damaging a business asset or causing a loss to occur — and *vulnerabilities* — conditions in which a threat capability is greater than a firm's ability to resist it.[22] With that information, a business can populate the cells of a cyber-risk threat matrix — a simple cross-tabulation of the likelihood and severity of potential cybersecurity events. Rudimentary as they are, such tools can provide insights into how the business should allocate defensive information technology (IT) investments. Figure 1 presents a typical risk matrix, with rows corresponding to the likelihood of a cyber event occurring and the columns to the event's expected impact. Each threat and vulnerability the business identifies can be placed into one of the boxes according to a qualitative likelihood and impact assessment.

[22] Jack A. Jones, "An Introduction to Factor Analysis of Information Risk (FAIR)," Risk Management Insight (2005).

Figure 1. Risk Matrix (example)

		Business Impact		
		Very costly	Moderately costly	Minor or negligible cost
Event Likelihood	Very likely	HIGH	HIGH	MEDIUM
	Likely	HIGH	MEDIUM	LOW
	Not likely	MEDIUM	LOW	LOW

Source: Author's illustration

The cyber-risk threat matrix provides company executives with a qualitative vision of enterprise-level cyber risk, which is a good starting point for many businesses. However, as the conversation evolves from educating and informing business leaders to empowering them to make business decisions, information security personnel may not be able to translate the matrix into actionable insights. For example, since no business has unlimited resources, accepting incremental risk in one area in exchange for lower risk in another. Such decisions are complicated and require the ability to assess existing risks relative to the level of risk the business can tolerate and still function smoothly. This is precisely the impetus for drafting a risk appetite statement — a company-wide statement of the amount of acceptable risk for day-to-day business affairs.[23] Risk appetite is discussed in greater detail in Section III.

Despite its shortcomings, a cyber threat risk matrix — and reviewing it periodically — is an important step to a more sophisticated risk management system. If the matrix is not meeting business needs, leaders may consider refining their qualitative likelihood and impact

[23] Alan Gemes, Peter Golder, Yogesh Patel, and Hussein Sefian, "What Is Your Risk Appetite? A Disciplined Approach to Risk Taking," *strategy&*, PwC, (originally published by Booz & Company; October 20, 2009).

assessments. To do so, each threat and vulnerability should be assigned both a statistical likelihood (probability of an event occurring) and an expected dollar loss consistent with its position on the risk matrix. Consider threat ABC, which is considered to be a medium risk (or likely to occur, with moderate business impact). An experienced information security specialist might use her industry knowledge and subjective judgment to estimate a 10 to 15 percent chance of ABC occurring within the next 12 months and be 95 percent certain that the business impact would fall between $1 million and $5 million. The chief information security officer (CISO) can then communicate to other executives that the expected loss from threat ABC is between $100,000 and $750,000.[24]

III. Highlights from the Cyber-Risk Workshop

During his talk, Fox, who leads PwC's cybersecurity and privacy assurance practice in the New York metropolitan area, discussed the importance of initiating an internal discussion of cyber risk and working through the process of quantifying threats and vulnerabilities. He noted that some organizations may only have a single layer of information security protection and little understanding of the specific company assets that should be safeguarded and what precautions are required for each. In his experience, many organizations are inefficiently allocating investment budgets, often spending too much money on perimeter controls (network boundaries that are used to isolate zones with different security policies) and data loss prevention tools (software that monitors outgoing information to prevent accidental leaks or exposure of sensitive information), and too little on tools to protect the organization's true assets, as discussed next.[25] A 2017 report by the Ponemon Institute confirms Fox's experience. As shown in Figure 2, five

[24] Calculations: 10%* $1 million = $100,000; 15%* $5 million = $750,000.

[25] Axel Buecker, Per Andreas, and Scott Paisley, "Understanding IT Perimeter Security," IBM Redpaper, IBM Corporation (2008); Prasida Menon, "How Data Loss Prevention (DLP) Technology Works," McAfee (2018).

of the nine security technologies reviewed for the study had a negative value gap in which the percentage spending level is higher than the relative value to the business, suggesting that many organizations may be misallocating information security spend.

Figure 2. Value Gaps Associated with Security Investments

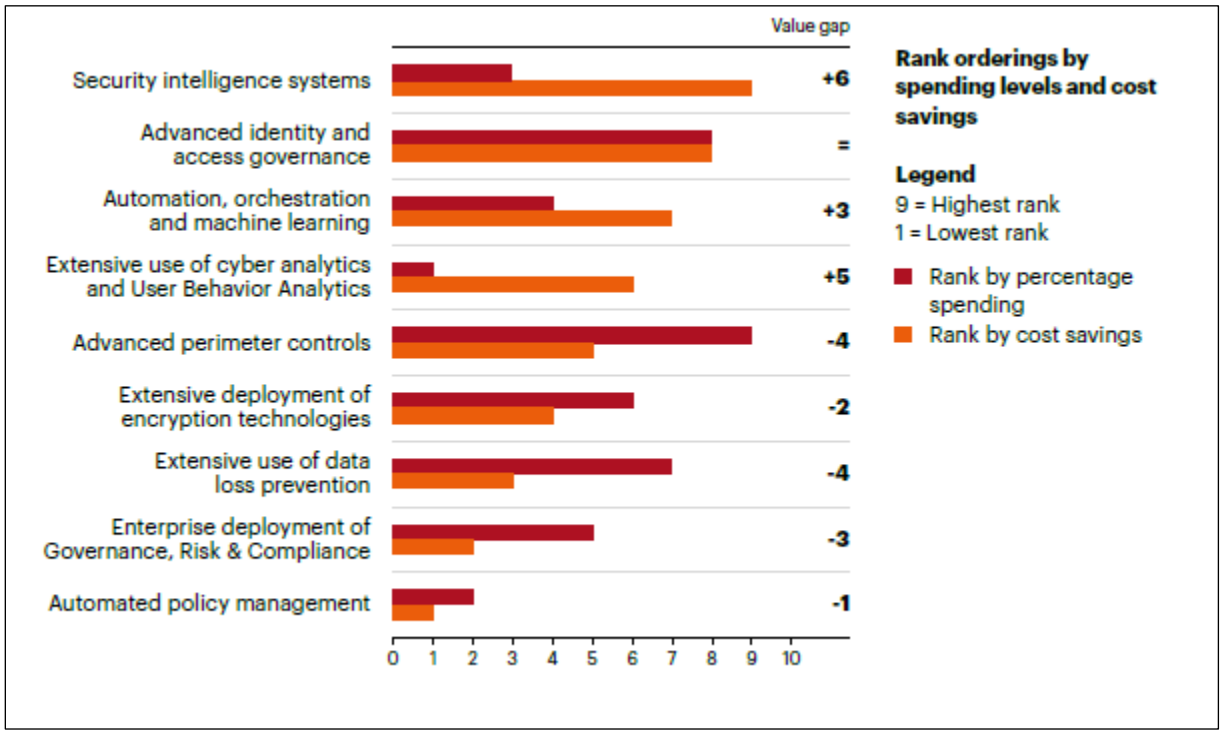

Source: *2017 Cost of Cyber Crime Study: Insights on the Security Investments That Make a Difference*, Ponemon Institute LLC

A. The Benefits of Measuring Cyber Risk

In their book *How to Measure Anything*, Douglas Hubbard and Richard Seiersen write that, if a business feels the need to measure something, it must be because it has tangible consequences for the business, even though the thing itself is intangible.[26] Thus, measuring even the most obscure or intangible cyberthreat begins with understanding its consequences. What are

[26] Douglas W. Hubbard and Richard Seiersen, *How to Measure Anything: Finding the Value of Intangibles in Business* (Wiley, Hoboken, NJ, 3rd edition, 2014), pp. 6–7.

the consequences of permitting contractors to access confidential data? Do underwriters working in a less secure off-campus workspace pose a heightened security risk? What is the risk associated with switching to a payment processor that employs a new, cloud-based technology? Answering these questions are at the heart of measuring cyber risk. During the workshop, Fox discussed the importance of initiating this often painstaking process, noting that measuring cyber risk enables business leaders to:

1) *Make Informed Decisions*

    - Risk decisions involve trade-offs. Protecting one asset or system will come at the expense of protecting another. It is impossible for a business to make informed risk decisions without having quantified its cyber risk since only then can decision makers understand what they must forego in pursuit of a particular direction.

2) *Allocate Resources Efficiently*

    - Likewise, spending an additional dollar on protecting a valuable but already well-protected system may not generate as much benefit as it would if the dollar were used to protect another, less secure system. Quantifying cyber risk helps leaders to make informed spending decisions that can mitigate risk while generating higher returns on investment.

3) *Recognize the Value of Systems and Information*

    - All businesses have "crown jewels" — data, systems, or intellectual property that, in the hands of cybercriminals, could bring the business to a halt. In some cases, the value of a particular data set or system may be apparent (e.g., internally developed software), while others may not (e.g., minutes from board meetings or discussions between members; information

on planned mergers, acquisitions, and divestitures). Quantifying cyber risk requires decision makers to know precisely which assets they need to protect.

4) *Resist the Urge to Boil the Ocean*

- Decision makers should focus on what is important — the crown jewels — and resist the urge to protect every data point or solve every problem at once. They need to start with something simple and allow it to evolve to meet their needs.

5) *Greater Visibility into IT*

- The IT infrastructure may not be well understood by executives and other leaders outside the IT department. Going through the process of assessing threats, identifying assets, and analyzing cyber risk provides the entire leadership structure with greater visibility into the firm's key risks, impacts, and threat likelihoods.

6) *Improve Incident Response*

- Many U.S. companies, including those in the financial services industry, are not adequately prepared to respond to a cyberincident. A 2016 PwC survey indicated that about 54 percent of surveyed companies had an active cyberincident response plan, while 17 percent had no plan at all.[27] When it comes to incident response, boards of directors and executive leadership should maintain a high level of engagement throughout the fiscal year and ensure that first responders are fully trained and that crisis management

[27] "Global Economic Crime Survey 2016: US Results," PwC (2016).

teams include IT professionals as well as legal, human resources, and digital forensic experts.

B. Challenges to Measuring Cyber Risk

Financial services firms may find it difficult to improve their cyber-risk measurement abilities. By all accounts, it is a painstaking process made worse by a lack of agreement on standards and a common language. Fox laid out the top five challenges faced by financial services companies seeking to enhance their cyber-risk measurement and cybersecurity.

1. The Business Lacks a Formally Defined Risk Appetite.

A successful cyber-risk management scheme requires a firm understanding of the company's risk appetite.[28] As discussed previously, a well-defined risk appetite statement translates risk metrics into business decisions, dynamically linking business strategy, performance targets, and corporate risk management.[29] Defining corporate risk appetite begins with establishing a risk baseline — a catalog of the firm's many types of risk exposures in financial terms. Once the risks are cataloged, leadership can determine whether their current operating practices and risk profiles are aligned with their desired risk appetite and tolerances. Figure 3 demonstrates how the pieces can be integrated into a qualitative assessment of corporate risk appetite and tolerances. In the example, only one out of the five risk categories has a risk profile within the firm's stated tolerance range. Credit, financial, and reputational risks are out of tolerance (unfilled circle in column 8), and operational risk is slightly out of tolerance (half-filled circle). With so many categories out of tolerance and 60 percent of the firm's

[28] In 2013, the Financial Stability Board provided a set of eight goals that an effective risk appetite statement should achieve. For ease of reference, the goals are reprinted in the Appendix. "Principles for an Effective Risk Appetite Framework," Financial Stability Board (November 18, 2013).

[29] Richard Barfield, "Risk Appetite — How Hungry Are You?" *the journal*, Special risk management edition, PricewaterhouseCoopers, pp. 8–13 (2005).

economic capital already allocated to credit risk, the hypothetical firm may need to allocate more assets to financial capital.

Figure 3. Corporate-Level Risk Appetite and Tolerances (Sample)

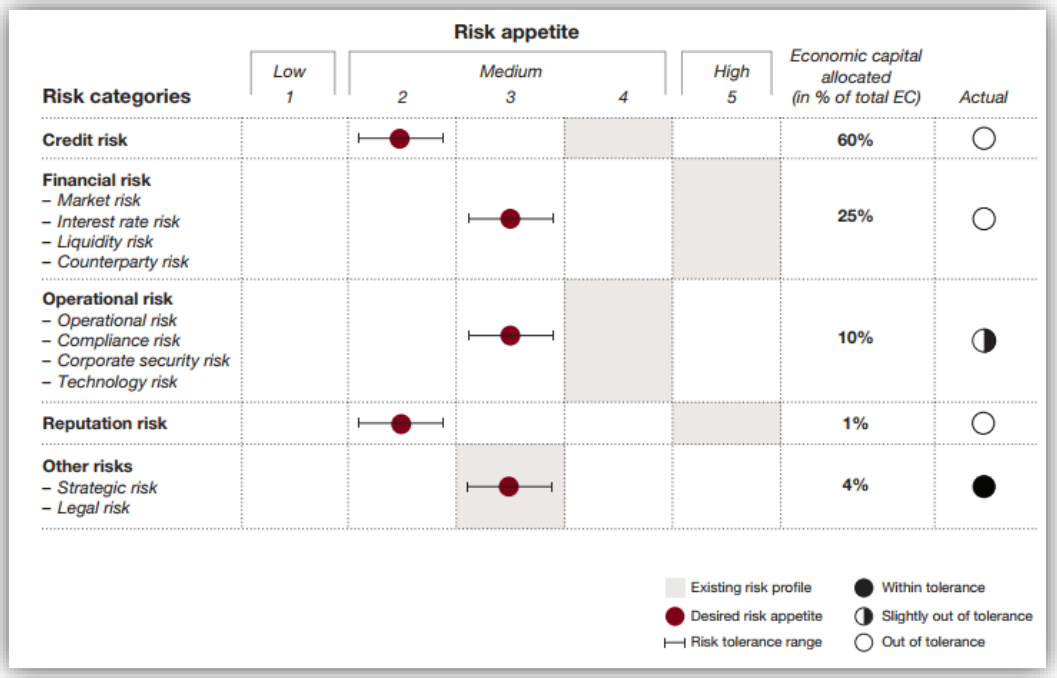

Source: Analysis from PwC Strategy&

2. *Cyber-Risk Management Is Not Fully Integrated into Enterprise Risk Management.*

Fox noted that, despite cyber risk being the biggest risk facing the financial services industry and a leading topic at board meetings, many financial institutions have cybersecurity risk models that are outside the enterprise risk management structure. Traditionally, cybersecurity has been managed within IT, separate from the operational risk management and compliance functions. Excluding cybersecurity risk from enterprise risk, intentionally or unintentionally, perpetuates communication barriers among chief risk officers (CROs), chief information officers (CIOs), and chief information security officers (CISOs) and can lead to a misplaced focus on prioritizing the protection of IT assets over business assets.

Fox stressed the value of integrating cyber-risk management functions into the firm's wider enterprise risk management strategy. Indeed, cyberattacks from external threats such as hacktivists or theft of intellectual property from within an organization are increasingly being recognized as forms of *operational risk*, the risk of incurring a loss due to external events or internal processes, people, and systems.[30] One of the main benefits of including cyber risk under the operational risk umbrella is that it provides the opportunity for the CISO to work alongside the CRO to develop contingency plans for cyberattacks. Such plans typically involve the deployment of backup systems until affected IT systems are fixed and brought back online and may include customer outreach, communication with regulators, and identifying losses incurred by the company or its customers.

3. The Business Lacks Consistent, Organization-Wide Risk Nomenclature and Measurement Standards.

Just as IT risk management should be fully integrated into enterprise risk management, the language of IT risk measurement should translate across the organization. As discussed in Section II.A, IT experts suggest that every cyber risk be described as the (a) likelihood and (b) dollar impact of a particular and well-defined threat exploiting a similarly well-defined vulnerability that has one or more describable consequences for the business.[31] Doing so can mitigate the possibility of miscommunicating with other business units and can facilitate analysis, aggregation, and comparison across risks of any kind.

It's also important for business units to agree upon and adhere to a set of clear measurement standards — mathematical methods and algorithms for calculating likelihood and

[30] Steve Kulp, "Banks Face Challenge of Integrating Cyber and Operational Risk," *Forbes* (April 26, 2017); "Consultative Document on Operational Risk," Basel Committee on Banking Supervision (January 2001).

[31] Stephen Bailey et al., "Understanding Cyber Risk Management vs Uncertainty with Confidence in 2017," NCC Group Whitepaper, 2017.

impact estimates. Standards ensure that business leaders are able to compare likelihood and impact estimates across business units and across time as the business evolves through mergers, acquisitions, and divestitures.

4. The Business Is Unable to Quantify the True Cost of a Cyberincident.

Organizations tend to respond to a successful cyberattack by attempting to isolate and remediate the threat to resume normal business operations as soon as possible. As a result, the total cost of a cyberincident is typically calculated as the sum of the incident response cost and lost business during downtime. According to Fox, the total cost of a cyberincident can run much higher than that and often involves accounting for costs that are difficult to measure. For example, a cyberattack may result in the degradation of a brand's reputation to customers and the loss of future business. Since employees who are working on threat remediation are unable to pursue other business priorities, the total cost calculation should also include the opportunity cost of employee salaries and remediation expenditures. Because of such omissions, organizations affected by a cyberattack tend to report lower losses than actually experienced. It can also cause them to underestimate the expected costs of potential cyberevents, leading to a misallocation of investment spend.

5. The Business Lacks Qualified Cybersecurity Employees.

A growing shortfall of qualified information security professionals means that businesses may not have enough employees to develop, operate, and monitor a robust information security program. According to Fox, demand for cybersecurity personnel continues to outstrip the supply of qualified labor, with 39 percent of cybersecurity jobs in the U.S. unfilled. By 2021, the global shortage of cybersecurity talent is expected to reach 3.5 million.[32]

[32] Steve Morgan, "Cybersecurity Labor Crunch to Hit 3.5 Million Unfilled Jobs by 2021," *Cybersecurity Business Report*, CSO (June 6, 2017).

One reason for the shortage is the rapid growth of the industry. The U.S. Bureau of Labor Statistics projects that employment of information security analysts will grow 28 percent from 2016 to 2026, compared with 13 percent for all computer occupations and 7 percent for all occupations.[33] Another challenge is that the job requirements continue to evolve, making it difficult for employers to define career paths for employees. Fox finds that employers would prefer to hire people with a mixture of business, cybersecurity, and risk management knowledge; however, few candidates have the required qualifications. The third reason for the shortage — lack of education opportunities at the high school, undergraduate, and graduate levels — may be a symptom of the first two reasons. Colleges and universities are often slow to build new degree programs or adapt existing programs in response to a rapidly growing industry with rapidly evolving needs. One survey of eight countries including the United States found that only 7 percent of top universities offer an undergraduate major or minor in cybersecurity.[34] A 2016 study of 121 U.S. universities found that none of the top 10 computer science programs required a cybersecurity course for graduation and that three of the top 10 programs offered no cybersecurity courses at all.[35]

The education gap is one area in which Fox expects improvement in the coming years. He noted that universities such as Carnegie Mellon, Syracuse, Baylor, and the University of Maryland have built multidisciplinary cybersecurity education programs both at the undergraduate and graduate level. In addition, the U.S. National Security Agency (NSA), through its National Centers of Academic Excellence, offers colleges and universities the

[33] "Information Security Analysts," *Occupational Outlook Handbook*, Bureau of Labor Statistics (last accessed on July 24, 2018).

[34] "Hacking the Skills Shortage," McAfee (July 2016).

[35] "CloudPassage Study Finds U.S. Universities Failing in Cybersecurity Education," CloudPassage (April 7, 2016).

opportunity to become accredited in cyberdefense and cyberoperations. There are currently about 249 accredited cyberdefense programs and 20 accredited cyberoperations programs.[36]

C. Cyber-Risk Management Methodologies

Fox introduced the audience to some of the leading IT risk assessment and analysis tools. He noted that the financial services industry to date has yet to coalesce around a single cyber-risk management methodology. Cybersecurity groups and standards organizations continue to advance various methodologies intended to quantify, mitigate, and monitor cyber risk at both the institution and industry level; however, no clear frontrunner exists. Organizational barriers and inadequate communication between risk and information technology functions make it difficult to integrate cyber risk into enterprise-level risk management. Many firms are unable to measure their own exposure to cyber risk, citing a lack of data, expertise, or guidance.

Of the approximately 20 different cyber-risk management methodologies in use today, Fox noted that five have the most traction in the financial services industry, although overall penetration is still limited. In his experience, most financial services organizations are either not using a cyber-risk management methodology or are not using it consistently when making cyber-risk management decisions. The primary constraint to implementation is a lack of available, verifiable data. In what follows, we introduce each of the five methodologies covered in Fox's workshop as well as some of his thoughts about each. This section is not intended to be a comprehensive review of the methodologies; please refer to the footnotes for additional resources and more detailed information.

[36] "NSA/DHS National CAE IN Cyber Defense Designated Institutions," Information Assurance Department, National Security Agency (2018); "Centers of Academic Excellence in Cyber Operations," National Security Agency (last modified on July 3, 2018).

1. *Information Risk Assessment Methodology, version 2 (IRAM 2)*

IRAM 2 is a qualitative information risk assessment **methodology** created by the Information Security Forum (ISF), an independent information security group with membership that includes corporations, public sector groups, and government agencies.[37] Fox considers IRAM 2 one of the better documented and easier to follow methodologies. Its six phases guide users from scoping information risk assessments across the business and technology to risk evaluation and treatment (see Figure 4). Among IRAM 2's highlights are its spreadsheet template and built-in reporting functions. It employs a comprehensive and somewhat rigorous methodology that ties in with other leading standards. Fox sees many organizations using ISF IRAM to augment their current IT risk management practices rather than as their primary methodology.

The prevalence of IRAM 2 may be constrained by its availability to ISF members only. Thus, it is not considered a public standard. Fox also noted that the methodology is overly detailed in certain areas, suffers from weakness in residual risk calculations, and requires the organization to build its own attack vectors and threats.

Figure 4. The Six Phases of IRAM 2

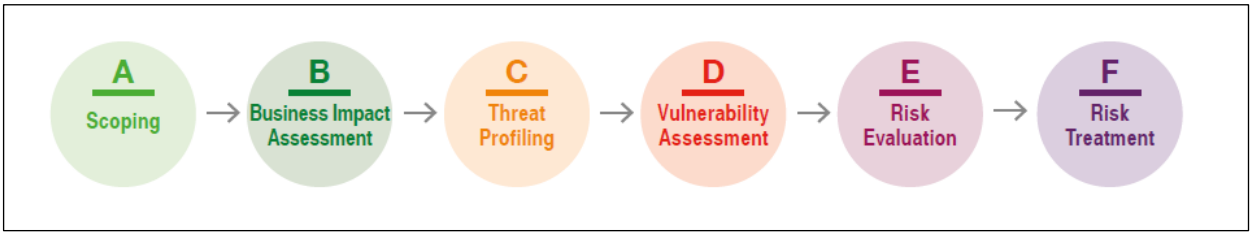

Source: Information Security Forum (requires registration)

[37] "About Us," Information Security Forum (last accessed on August 10, 2018).

2. *Risk IT*

Risk IT is a **framework** developed by ISACA (formerly the Information Systems Audit and Control Association) that is designed to help organizations manage all IT-related risk.[38] It was originally intended to supplement COBIT (formerly Control Objectives for Information and Related Technologies), a comprehensive framework for governing and managing enterprise IT risk. However, with the introduction of COBIT version 5 in 2012, it has been integrated into the broader framework.[39] As a stand-alone product, one of Risk IT's key features is its ability to work in harmony with enterprise risk management and other common enterprise risk categories.

Risk IT has some drawbacks that may make it difficult for less sophisticated organizations to adopt. Fox noted that both Risk IT and COBIT require a significant time investment compared with other frameworks, both in their initial ramp-up time and in ongoing maintenance. Relative to other frameworks, its risk appetite guidance may lack sufficient granularity for some organizations. In addition, where cyber-risk frameworks have recently tried to strike a balance between risk and operations, Risk IT favors a risk-centric approach. Lastly, within IT risk, cybersecurity risks may not receive adequate focus.

3. *What's FAIR?*

Factor Analysis of Information Risk (FAIR) is a quantitative risk analysis **standard** that views enterprise risk in totality. FAIR starts by defining a risk taxonomy that decomposes risk into a set of well-defined factors that provide business leaders with a straightforward and

---

[38] "Risk IT Framework for Management of IT Related Business Risks," ISACA (last accessed August 10, 2018). IT-related risk includes both information security, which is the practice of protecting the confidentiality, integrity, and availability of data and systems, and IT risk, which is the practice of identifying, monitoring, and mitigating information risk.

[39] "COBIT 5 A Business Framework for the Governance and Management of Enterprise IT," ISACA (2012). Note: Requires registration.

repeatable means of assessing risk. FAIR also provides a 10-step process for assessing the risk to any business asset from a particular threat community (e.g., employees, cleaning crew, hacktivist group). The final steps of the process enlist a computational engine to derive risk estimates and a mathematical simulation model that enables users to analyze various risk scenarios and to challenge and defend risk decisions.[40]

According to Fox, FAIR is helping to transition cyber-risk management from an art to a science. Some of its main advantages are that it is rooted in information security but translates to more general risk applications, puts information risk into financial terms comparable with other risks, and understands how time and money will affect a firm's cybersecurity risk profile.[41] In his experience, Fox noted that one area in which the FAIR standard is lacking is in its ability to help business leaders make decisions involving risk trade-offs, such as how much confidential information to share with another business to maintain customer confidentiality while also creating a verifiable transaction.

Fox noted that FAIR is gaining popularity within the financial services industry, particularly with the largest financial institutions. This is due, in part, to the recent growth of software, workflows, and support tools that have been developed within the FAIR ecosystem. Institutions have the opportunity to leverage either the free open source tool or one of the commercialized versions now available. Recently, The Open Group, a global consortium of more than 600 organizations, adopted the Open FAIR Standard as the international standard

[40] Jack A. Jones, "An Introduction to Factor Analysis of Information Risk (FAIR)," Risk Management Insight (2005).

[41] More generally, a risk profile is a decision support tool that identifies the business' known risks, their potential effect on business operations, and the controls in place to mitigate those risks. Sources: "Understanding Your Company's Risk Profile," Aon, (last accessed on August 21, 2018); ERM Initiative Faculty and Chris Cox, "Understanding Risk Appetite," NC State University (May 1, 2014).

Frameworks, Methodologies, and Standards: What's the Difference?

The various approaches to cyber-risk management are generally described as frameworks, methodologies, and standards. To make matters more confusing, the same approach is often referred to as a framework in some instances and a methodology or standard in others.

The term *framework* describes a flexible approach to managing cyber risk that provides guidance without requiring objectives to be accomplished in a specific manner. Rather, a framework allows the user to determine the best way to achieve the desired results. An integrated cyber-risk framework is one that fits into a broader enterprise risk management framework.

Methodologies tend to be more prescriptive and less flexible. A methodology may include a set of rules, prescribe the use of one or more methods, or spell out a specific process in detail.

A *standard* is a document, established by consensus and approved by a recognized body that provides principles, rules, or guidelines for some activity or result. Standards help facilitate order and repeatability and ensure uniformity across organizations.

So how do they all fit together? A framework can support many methodologies, and can satisfy any number of public standards. A methodology can coexist with other methodologies and may satisfy a public standard. Standards themselves are generally neither frameworks nor methodologies.

Sources: "International Standards," International Electrotechnical Commission (last accessed July 25, 2018); Michael Wood, "Why You're Confusing Frameworks with Methodologies," ProjectManagement.com (May 6, 2013).

information risk management model.[42] Fox believes that FAIR is well positioned to be adopted as a cyber-risk measurement standard across multiple organizations.

4. *NIST Cybersecurity Framework*

The National Institute of Standards and Technology (NIST) is responsible for developing measurement standards for the U.S. federal government.[43] The NIST Cybersecurity Framework (CSF) was mandated by executive order in February 2013 and builds on existing guidelines and standards, including NIST SP 800-30 and FIPS 199.[44] Its intended purpose is to

[42] "About Us," The Open Group (last accessed on August 8, 2018). "What Is FAIR?" FAIR Institute (last accessed on August 8, 2018).

[43] "About NIST," National Institute of Standards and Technology (last accessed on August 8, 2018).

[44] "Executive Order — Improving Critical Infrastructure Cybersecurity," White House (February 12, 2013).

help strengthen cybersecurity-risk management at organizations that manage critical national infrastructure, such as financial services, energy, and telecommunications, although it is flexible enough to be deployed at any company.[45] The CSF is a comprehensive risk-based **framework**. Its five core concepts — identify, protect, detect, respond, and recover — provide business leaders with guidelines for evaluating and managing their company's cyber risk as well as how best to respond if a cyberattack should occur.

Risk Analysis or Risk Assessment?

Cyber-risk management methodologies are typically distinguished by certain criteria: whether they focus on risk analysis, risk assessment, or both. Generally speaking, *risk assessment* refers to the overall process of identifying, evaluating, and analyzing risks. *Risk analysis* is a subprocess of risk assessment and refers to the process of comprehending the nature of risk and determining the level of risk.

Source: ISO/Guide 73:2009 https://www.iso.org/obp/ui/#iso:std:iso:guide:73:ed-1:v1:en.

As a government-sponsored public standard, CSF is the preferred framework of many regulatory agencies, according to Fox. Benefits of adopting the framework include access to its threat assessment templates and suggested risk metrics (e.g., threat, impact, likelihood of initiation). He also noted that it is important to distinguish the NIST framework from more analytical approaches such as FAIR. Rather than prescribe specific metrics and tools, the CSF allows users to leverage any number of analytical tools. Fox noted that some financial services companies run complex Monte Carlo analyses, while others perform a simple Pareto analysis, all within the NIST framework.

The CSF is not without shortcomings. Fox explained that many users find it difficult to understand its terminology. The lack of an accompanying toolset and limited guidance on how to determine the business impact of a cyberevent and how to establish a risk appetite means that businesses must pull in additional resources to complete their cyber-risk management build. He

[45] "Cybersecurity 'Rosetta Stone' Celebrates Two Years of Success," National Institute of Standards and Technology (February 18, 2016).

also expressed concern that, because the CSF is broadly applicable to businesses of all sizes across all industries, it may not be the best fit for any one particular company or industry.

5. *CyberVaR*

First envisioned by the World Economic Forum in 2015, the term *CyberVaR* describes an approach to quantifying cyber risk that borrows a statistical method called value-at-risk (VaR) from the financial services industry.[46] For a given financial portfolio (a group of bonds, equities, or other investment vehicles), a VaR analysis estimates an upper bound for the dollar amount the portfolio might lose under typical business conditions.[47] For example, a VaR might estimate a $5 million upper bound for what the portfolio would lose within the next day, with 99 percent certainty.[48] In his comments, Fox noted that VaR's ability to enumerate cybersecurity risk in a single number makes it appealing to many financial services executives, particularly those who are familiar with VaR but unfamiliar with cyber risk. However, he also noted that CyberVaR requires both the availability of sufficient clean data for model estimation as well as employees trained in financial statistics. Once a frontrunner in the financial services industry, CyberVaR has lost traction to other methodologies.

D. Adoption and Consensus

To date, no consensus has emerged around a particular cyber-risk management standard, framework, or methodology. Each of the leading approaches to quantifying risk has its own advantages and disadvantages, many of which were discussed in the previous section. In addition, some institutions may not be in a position to implement one of the more sophisticated

[46] "Partnering for Cyber Resilience Towards the Quantification of Cyber Threats," World Economic Forum, *Industry Agenda* (January 2015).

[47] Krzysztof Ostaszewski, "Value at Risk," Illinois State University (2007).

[48] A common criticism of VaR is that it places no upper bound on the losses that occur the other 1 percent of the time.

and data-intensive approaches such as FAIR and CyberVAR. In deciding which one to implement, Fox recommends that financial institutions focus less on determining which approach is subjectively better, but instead, select the approach that is most implementable given the institution's current operational and information constraints. Once an approach is selected, the financial institution should ensure the framework can be maintained and that it is integrated into the executive decision-making process. Once these conditions are met, Fox recommends that institutions proceed according to the Pareto Principle, with a focus on identifying and managing the 20 percent of threats and vulnerabilities that are expected to cause 80 percent of the total cyber risk.[49]

Down the road, Fox sees the industry adopting an amalgamation of the five approaches. One possibility is that all financial institutions agree (or are required) to adhere to the NIST cybersecurity framework but are permitted to construct their inputs from one or more of the other approaches. Propelling the industry toward this outcome is the increasing prevalence of cyberattestations. The American Institute of Certified Public Accountants' cyberattestation requires an independent auditor to examine the financial institution's program and attest to its adequacy. The attestation must then be signed by the institution's board of directors and chief executive officer. Fox noted that a financial institution that is not employing any one of the five approaches might find it difficult to pass an audit.

[49] In business, the Pareto Principle is a decision-making tool used to promote efficiency. It is based on the economic concept of diminishing marginal returns, whereby each additional unit of work generates a smaller quantity of results than the previous unit. Decision makers adhering to the Pareto Principle prioritize their efforts and resources by identifying the 20 percent of the total workload that will generate an outsized proportion (80 percent) of the results. The remaining 20 percent of results, then, are achieved with much less efficiency, requiring the business to perform the remaining 80 percent of the work. For an example, see Daniel Moody and Peter Walsh, "Measuring the Value of Information: An Asset Valuation Approach," European Conference on Information Systems, submission (1999).

E. Recommendations for Bank Exams

In order for bank examiners to be able to compare cyber risk by institution, Fox recommends that, as part of the regular examination process, examiners ensure that supervised banks are taking the following four steps:

1. Maintaining an awareness of their threat landscape;
2. Identifying their crown jewels;
3. Creating a well-defined and well-understood risk appetite statement; and
4. Implementing the risk appetite statement to decide where to allocate investment dollars.

A financial institution that is aware of the cyber-threat landscape knows where its most vulnerable cyberthreats are, the types of cyberthreats that may arise in the near future, and is thinking about ways to respond to both. Firms can gain awareness by constantly monitoring in-house systems, threats, and vulnerabilities and by maintaining a productive dialogue with peers via industry trade groups and cybersecurity experts. To determine the veracity of an institution's threat awareness, Fox recommends that examiners use penetration testing and other means to validate the bank's threat data and review the institution's strategy for maintaining awareness over time.

As discussed in Section III.A, business leaders often find it difficult to identify the firm's true crown jewels. For example, IT executives may consider configuration management databases (CMDBs) and other repositories of important IT configuration data to be among the company's key assets. While information contained in a CMDB does have target value to cybercriminals, it is often less valuable than information regarding the firm's strategic plans. For example, confidential communications between board members and C-level executives, secret

merger and acquisition plans, or sensitive customer information are all high-value targets to cybercriminals intent on disrupting business operations.

Once examiners are comfortable with the bank's responses to the first two items, Fox recommends they review the bank's risk appetite statement, if one exists, and verify that bank leadership is using the risk appetite statement to determine where to make cybersecurity investments and where to "let it ride" with existing protections. According to Fox, a bank that hasn't defined its risk appetite is probably not thinking of cybersecurity investment in the right terms and may instead be dividing a limited cybersecurity budget across a range of threats without regard for the likelihood and impact of any particular one.

F. Quantifying Systemic Cyber Risk

In the October 2016 joint ANPR, federal banking regulators expressed concern about the financial sector's vulnerability to systemic cyber risk. The regulators recognized that measuring institutional cyber risk in a standardized and comprehensive way, with a focus on the largest and most interconnected financial institutions, was the first step in beginning to understand the financial sector's overall vulnerability to cyberthreats. The regulators proposed establishing enhanced standards that would contain binding requirements for the largest, most interconnected U.S. financial entities and recommendations for other regulated financial entities.

Fox expressed concern about the agencies taking an overly prescriptive approach in any new regulations. He shared his thoughts on how banking and securities regulators could help promote widespread adoption of cyber-risk management tools. In particular, he noted that regulators could accelerate the process by writing rules that require financial institutions to do the following:

1. Have a credible cyber-risk management plan in place;

2. Collect the data necessary to support the plan;

3. Complete a cyberattestation stating that the board of directors is aware of the cyber-risk management plan and is actively involved in cyber-risk management decisions; and

4. Provide evidence that the cyber-risk management plan is being used to consistently make investment decisions.

If regulators focus on enforcing those four requirements across all organizations and required companies to secure independent verification, Fox believes it will significantly lower risk across the entire financial system. He also suggested that federal financial regulators use their power to help industry participants collaborate on common standards, such as FAIR. Only when these requirements are fulfilled by regulated financial entities and common standards have been agreed upon can federal regulators and other cybersecurity experts take up the challenge of determining how to measure and monitor systemic cyber risk.

When that time arrives, Fox recommends regulators approach the problem by identifying the inherent risk of different pieces of the ecosystem (e.g., payment, trade, and settlement networks), establishing a set of recommendations to address those risks, and developing a plan to monitor the deployment and ongoing operation of the recommended actions. Once those steps have been taken, Fox believes that federal regulators will be in a much better position to gauge industry-wide residual risk and subsequently deploy one of the leading cyber-risk management frameworks at the industry level.

IV. Conclusion

The 2016 ANPR issued by several U.S. financial regulatory agencies proposed enhanced cybersecurity standards for many financial institutions to better understand and manage the

sources of systemic vulnerability in the financial system. Section VII of the ANPR asked the public to provide feedback on their experience with existing cyber-risk management methodologies and share their recommendations on how best to leverage the methodologies to compare the levels of cyber risk posed by entities across the financial sector. To gain a better understanding of how the various methodologies are used to quantify cyber risk, the Consumer Finance Institute hosted a workshop in 2017 with James Fox, head of PwC's cybersecurity and privacy assurance practice in the New York metropolitan area. Fox discussed the importance of measuring cyber risk, highlighted some of the challenges business face when trying to measure cyber risk, and laid out the pros and cons of several leading cyber-risk management methodologies. He also provided recommendations for bank examinations and his thoughts on how federal agencies might go about beginning to quantify systemic cyber risk.

Both the ANPR and new regulations established in 2017 by the New York Department of Financial Services demonstrate the heightened awareness that cyber risk is receiving in the financial services industry and among its community of regulators. The Consumer Finance Institute will continue to monitor developments and share industry progress.

Appendix. Principles for an Effective Risk Appetite Framework

An effective risk appetite statement should follow these basic guidelines:

1. Include key background information and assumptions that informed the financial institution's strategic and business plans at the time they were approved;
2. Be linked to the institution's short- and long-term strategic, capital and financial plans, as well as compensation programs;
3. Establish the amount of risk the financial institution is prepared to accept in pursuit of its strategic objectives and business plan, taking into account the interests of its customers (e.g., depositors, policyholders) and the fiduciary duty to shareholders as well as capital and other regulatory requirements;
4. Determine for each material risk and overall the maximum level of risk that the financial institution is willing to operate within, based on its overall risk appetite, risk capacity, and risk profile;
5. Include quantitative measures that can be translated into risk limits applicable to business lines and legal entities as relevant, and at group level, which in turn can be aggregated and disaggregated to enable measurement of the risk profile against risk appetite and risk capacity;
6. Include qualitative statements that clearly articulate the motivations for taking on or avoiding certain types of risk, including for reputational and other conduct risks across retail and wholesale markets and establish boundaries or indicators (e.g., nonquantitative measures) to enable the monitoring of these risks;
7. Ensure that the strategy and risk limits of each business line and legal entity, as relevant, align with the institution-wide risk appetite statement as appropriate; and
8. Be forward looking and, where applicable, subject to scenario and stress testing to ensure that the financial institution understands what events might push the financial institution outside its risk appetite and/or risk capacity.

Source: "Principles for an Effective Risk Appetite Framework," Financial Stability Board (November 18, 2013).

FEDERAL RESERVE BANK OF PHILADELPHIA

Consumer Finance Institute Discussion Paper Series

http://www.philadelphiafed.org/consumer-finance-institute