CLTC
Center for Long-Term Cybersecurity
UC Berkeley

RESILIENT GOVERNANCE FOR BOARDS OF DIRECTORS

CONSIDERATIONS FOR EFFECTIVE OVERSIGHT OF CYBER RISK

In partnership with

Booz | Allen | Hamilton®

Abstract

How should boards of directors oversee cybersecurity risk for large global companies? This has become an urgent question for directors, for firms, and for the economy and society as a whole.

Over the summer of 2019, a team from Booz Allen Hamilton and the Center for Long-Term Cybersecurity at UC Berkeley interviewed 20 board members (mainly but not exclusively from U.S. companies, representing nine out of 11 industry sectors and with over 130 years of combined board service) to assess their beliefs, practices, and aspirations on cybersecurity governance.[1] Several important insights emerged that shed light on this high-stakes area, with which many boards have only started to grapple in the last four or five years. We found there is no single governance playbook for cyber that can be applied across sectors and risk profiles. We identify instead four "dynamic tensions" that characterize a set of trade-offs that board governance processes reflect in practice, although not always with clear intention. We explain the risks associated with different positions boards take on the dynamic tensions, as well as synergies and inconsistencies among them, and propose practices that directors should incorporate in their governance activities.

Cybersecurity risk requires a different and more dynamic governance model than is common among boards for handling other risks, a mindset we define as "resilient governance." We propose that boards should become fully intentional and self-consciously aware of how they are positioned, and rigorously test the actions they have taken to extend the upsides and de-risk the downsides of those choices around the four dynamic tensions. Boards should re-assess those decisions on a regular basis to take account of changes in the internal and external risk environment. Discussions about the role of the chief information security officer (CISO) and the CISO-board relationship are dependent upon and should be framed within this broader picture of board governance and oversight.

We will conduct future studies to follow the evolution of board thinking and practice in this area and continue to integrate the results into guidance and recommended practices as a means of driving and measuring progress toward resilient governance of cybersecurity risk.

JUST GETTING STARTED

Until four or five years ago, it was uncommon for boards of directors to address cybersecurity and related risks in a regular and disciplined fashion. "Cyber" (for shorthand) was mostly treated as an operational issue in the hands of information technology (IT) management and business units, not sufficiently critical to rise to the threshold of the board for strategic, enterprise-level oversight and

FOUR DYNAMIC TENSIONS

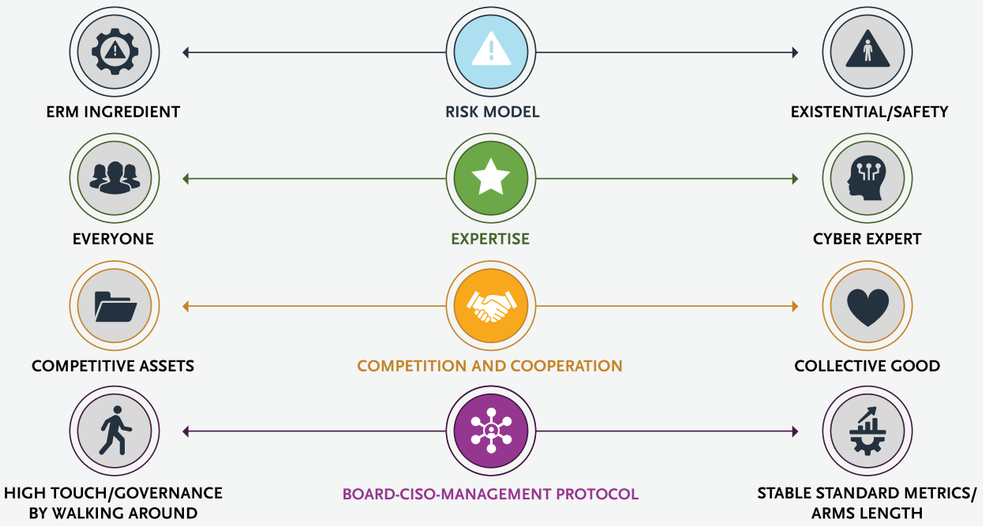

[1] See call out box on page 3 for more details on the interviews. Appropriate caveats apply for sample size.

governance. Now nearly the opposite is true. Boards feel a deep sense of urgency to exercise a central role in improving cybersecurity postures and outcomes for the firm. This attitude is appropriate, because by most common measures cybersecurity problems are morphing and mounting in importance faster than they are being solved or managed.[ii] In interviews with 20 board members mainly from U.S. companies and representing 9 of 11 industry sectors, a large majority said that cyber is one of the biggest risks that organizations face. Cyber will without question remain a board-level issue for the foreseeable future.

One of the practical consequences of this rapid shift is that, for many boards, cyber is a relatively new area of work. Board members are aware that they are in the early stages of thinking about how to best carry out oversight and governance responsibilities. They are also aware that the stakes are high. The risks associated with digital technology are mounting as increasingly complex regulatory and legal obligations make cyber governance more challenging. Put differently, almost no one expresses confidence that their board has "gotten it right" on important issues related to cyber, or that they know how to get to a place where they feel that way. The more typical experience is a mix of uncertainty and anxiety.

When we asked a straightforward question —"What information and processes do you feel you need to provide effective oversight and governance of cyber within your organization?"— the most common answer was, "we don't yet really know." There is no clear consensus on what "good" looks like in this domain, but there is a deep sense that the question is pressing and urgently needs to be answered.

The 20 BOARD MEMBERS interviewed for this paper represent 9 OF 11 INDUSTRY SECTORS as defined by the Global Industry Classification Standard (GICS), including: Communication Services, Consumer Discretionary, Consumer Staples, Financials, Healthcare, Industrials, Real Estate, Information Technology, and Utilities.

They represent 52 BOARD ROLES (many interviewees serve on multiple boards) and have 130 YEARS of combined board experience (in their currently sitting positions).

They have an average of 7 YEARS of service on their current board(s), ranging between 1 AND 22 YEARS.

Much of the emerging literature and debate around governance of cyber within the enterprise focuses on chief information security officers, particularly their role within the management structure and relationship with the board.[iii] These issues are becoming more salient as the seniority and compensation of CISOs continue to expand. But we believe it is impossible to get clarity on the role of the CISO outside of a broader model that addresses the prerequisite question of how boards should in practice oversee cyber for the enterprise. Without that model, CISO-board relationships tend to devolve into compliance and box-checking exercises that both boards and management find unsatisfying and ineffective. What is needed is a larger frame that contextualizes the CISO-board relationship (and the relevant responsibilities of other senior managers) within a model of governance that is appropriate for cyber-risk. We began our study thinking the central question would be what strategic role the board needs the CISO to play, but quickly discovered we needed to start earlier in the story and understand how boards' mindsets on cyber-risk are evolving in the context of rapidly-changing threats and evolving competitive and regulatory environments.

Currently, there is no stable and consensual playbook for board oversight of cyber. There are in fact significant differences in what directors mean when they assert that cyber has become a board

[ii] See Malwarebytes' Q1 2019 report, Cybercrime tactics and techniques, which illustrated that since Q1 2018, detections of threats to businesses have increased 235 percent. https://resources.malwarebytes.com/files/2019/04/MWB-CTNT-2019-state-of-malware_FINAL.pdf page 4.; Wired's extensive and growing list of reports on cyberattacks. https://www.wired.co.uk/article/hacks-data-breaches-in-2018; McAfee's Feb 2018 Report, Economic Impact of Cybercrime, estimating that cybercrime in 2018 cost the world almost $600 billion, or 0.8% of global GDP, a massive increase from the $445 billion estimated in 2014. https://csis-prod.s3.amazonaws.com/s3fs-public/publication/economic-impact-cybercrime.pdf?kab1HywrewRzH17N9wuE-.

[iii] See for example National Association of Corporation Directors' Cyber-Risk Oversight Director's Handbook Series Appendix I, page 38, about the CISO's role and mandate. This report also points to board composition and culture/mindset issues as important considerations, overlapping with and reinforcing some of our interview findings that we report below.

issue. For example, should cyber-risk be addressed as a central part of overall business strategy? Should it figure prominently in investment decisions—for example, about new product development or new lines of business—that rise to the board level? Should cyber-risk take prominence in board discussions about mergers and acquisitions? Board members have very different views on these questions, but they often express low confidence in their answers and a lack of clarity about what their colleagues believe and why.

A surprising proportion of our interviewees began by asserting that cybersecurity is now an "existential risk," which we took to mean a fundamental hazard to the continued existence of the enterprise. This is surprising because, while large firms have endured tens to hundreds of million dollars in direct losses from cyberattacks (consider Maersk's business losses, Capital One's equity market impacts, and Starwood's GDPR fines as recent examples), it is hard to identify a major firm or government organization that has ceased to exist as a result of a cyberattack. (Small enterprises have proven more vulnerable to existential risk from cyber, but of course they are more vulnerable to existential risk in many respects, and cyber isn't special in that respect).[iv] Reinforcing this point, relatively few CEOs of large firms have lost their job as a result of cyberattacks, and those who did were arguably punished more for their perceived mistakes in responding to the consequences than for failing to defend against the attack itself.[v]

When pressed, many board members emphasized the deep reputation risks associated with cyberattacks, while acknowledging that these risks are difficult to quantify, particularly as the time horizon gets longer. A common refrain is that a significant breach "damages how people view you," but how much and for how long is highly uncertain. The risk quantification challenge in cyber—a combination of insufficient historical data and the rapid evolution of the threat—is felt intensely by boards. Risk management techniques and models from other domains that might be applicable to cyber-risk management, such as market risk or credit risk, are typically not considered adequate. From evolving technology and attack surfaces to the risks of legal and regulatory exposure, many board members feel the goalposts in cyber are in constant and rapid motion.

Most boards feel they are just getting started with oversight of cybersecurity; that they got started later than they should have; and that they have to upgrade quickly. Our research identified key areas of agreement that are shaping perspectives and decisions about where to go:

- **As a board issue, cyber-risk is no longer confined to a set of operational decisions to be left solely in the hands of IT management. This is true regardless of firms' dependency on digital networks and assets, suggesting that board attention to cyber is driven by a change in the perceived need to engage, not only in the escalation of threats. The level of risk varies, but no organization is immune.**

- **Standard board governance frameworks are applicable at the highest level of abstraction, but are not specific enough to create a best-practice model for lower-level implementation and action. Boards can still think in terms of assessing first, second, and third lines of defense. They can focus, as is common, on process, risks embedded in process, controls, and effectiveness of controls. The trio of detection, prevention, and remediation is useful as general guidance. But filling in the details required to make these high-level oversight frameworks operational is a different challenge rendered unique by the dynamic nature of the threat.**

- **Sectors differ in their overall exposure and in their relative sophistication around cyber-risk. Firms vary in their capacity to understand and manage all kinds of risk, and cyber is no exception. But within the variation is a common theme about the crucial importance of corporate culture. Culture shapes to a considerable degree the manner in which mindsets about and models of governance play out in cyber. This is particularly salient for thinking about operational cyber-risks that flow down through the organization to lower-level employees, such as phishing attacks, where the weakest link is often at the individual behavioral level.**

- **Boards need CISOs to translate complex technical and engineering concepts into relatively simple language, just as is needed for other specialized areas of risk. There is a deeply felt need for metrics that can be compared over time (if not necessarily across enterprises or sectors). But there is also suspicion about the quality of metrics, their persistent value, and their potential to become obsolete or be manipulated. How a CISO talks to the board is an area of legitimate concern and attention, but this is a symptom of a larger systemic challenge, not merely an issue of better communication.**

- **There is a "black swan" aspect to the cyber threat that contributes to the difficulty of integrating cyber into broader risk management processes. A large number of minor cyber events are mitigated every day at most large organizations, but cyber professionals and boards are mindful that a major cyber event (including the most impactful and catastrophic ones) could come from a surprising or unanticipated direction. That makes failures of cyber defense in some cases — possibly the most important ones — not necessarily a failure of operational rigor but equally or more so a failure of imagination.**

[iv] See 2019 Verizon Data Breach Investigations Report findings: 43% of breaches involved small business victims. https://enterprise.verizon.com/resources/reports/2019-data-breach-investigations-report.pdf page 5; Malwarebytes' Cybercrime tactics and techniques Q1 2019 stating that the most vulnerable business targets are those of small and medium size (SMBs), which the 2018 study on "The State of Ransomware Among SMBs" demonstrated are battling the same number of threats but with the fraction of the security budget of a large enterprise corporation. https://resources.malwarebytes.com/files/2019/04/MWB-CTNT-2019-state-of-malware_FINAL.pdf page 6; Mark Smith's article in The Guardian stating that the 2015 UK Government Security Breaches Survey found that nearly three-quarters (74%) of small organizations reported a security breach in the last year, a significant increase from previous years' surveys. https://www.theguardian.com/small-business-network/2016/feb/08/huge-rise-hack-attacks-cyber-criminals-target-small-businesses Feb 8 2016; Gary Miller's article in The Denver Post stating that the U.S.' National Cyber Security Alliance found that 60% of small companies are unable to sustain their businesses over six months following a cyber-attack; according to the Ponemon Institute, the average price for small businesses to recover from attack stands at $690,000; for middle market companies, it's over $1 million. https://www.denverpost.com/2016/10/23/small-companies-cyber-attack-out-of-business/ Oct 23 2016.

[v] Examples include the 2017 Equifax data breach, where 147.9 million accounts were compromised, with records containing social security numbers, birth dates, addresses, and in some cases driver's license numbers; the CEO stepped down within two months following the breach; and the December 2013 Target data breach in which the debit/credit/contact information of 110 million people was compromised, with an ultimate cost estimated at $162 million; Target's CIO and CEO resigned by the first half of 2014. https://www.csoonline.com/article/2130877/the-biggest-data-breaches-of-the-21st-century.html Dec 20 2018.

These areas of conceptual agreement are not nearly enough to define a governance playbook that can be applied across organizations. And this could be a benefit as much as a problem. Established best-practice disciplines are attractive and seductive, but they can also be quite dysfunctional when they are premature and brittle to a changing risk environment. No board member we spoke with believes that the cyber-risk environment is stabilizing or that it is likely to do so in a predictable way over the next few years. So what kind of framework can boards use to oversee and govern cybersecurity in the enterprise *right now*, while *evolving* over time alongside the technology-business-regulatory complex at the heart of cybersecurity risk?

We propose a model of four dynamic tensions that board governance of cybersecurity will necessarily engage and manage. In the simplest terms, these represent four fundamental questions that boards grapple with around cyber:

1. What is our overall risk model for cyber, and how does it relate to other risks that we try to govern and oversee?
2. Where, how, and when do we access the expertise we need to carry out our work on cyber?
3. How does cybersecurity fit into our competitive strategy and sectoral relationships?
4. How do we want to share and exchange information and perspectives on cyber with management, and particularly with the CISO?

These are straightforward questions with complicated and interdependent answers. We present the choices that emerge as *dynamic tensions*, depicted visually as continua with a number of possible landing spots along the axes. Dynamic tensions are different than explicit "trade-offs," in the sense that there are no optimal landing spots that can be calculated given a known set of parameters. Dynamic tensions are, in fact, dynamic, as the terms of the relevant trade-offs are in motion. We articulate the most salient strengths and weaknesses associated with particular choices along each of the tensions. Note that the four tensions are not entirely independent; there are clusters of landing spots across the four that are more likely to appear in board practice than others, and there are logically inconsistent clusters that hardly appear at all.

Our guiding hypothesis about what good looks like thus becomes more dynamic and evolutionary. We propose that a healthy board process for cybersecurity governance:

- **Locates *self-consciously* and *explicitly* on each of the dynamic tensions. Board members need to know where the board is and why it has chosen to be there, without locking in to the idea that any particular choice is always best and stable over time.**

- **Understands and is relatively comfortable with the balance of pros and cons that characterize a chosen location on a continua. Most important, boards should be actively working with management to multiply the upsides and de-risk the downsides of their chosen landing spot.**

- **Re-evaluates the landing spot on each dynamic tension on a regular basis (possibly annually) to test for possible upgrades in the context of changes in the threat landscape or regulatory and business environment. This kind of meta-process is taxing and onerous; many boards will resist such frequent re-assessment. But we believe that cybersecurity risk management demands it.**

- **Grades for effectiveness *and* adaptability. Boards need a measure not only of how the enterprise is performing with regard to cybersecurity, but also of how their own oversight process is doing in contributing to better overall performance. While metrics can help satisfy regulatory and compliance demands, static measures of performance are insufficient, and adaptability is at least as important. Rather than "survival of the fittest," boards should think in terms of "survival of the most adaptable" to achieve a model of resilient governance.**

FOUR DYNAMIC TENSIONS

Dynamic Tensions 1. RISK MODEL

The most important dynamic tension that boards have to manage is a special case of the most familiar — what is the overall risk model or mindset within which cybersecurity fits. If boards believe it makes good sense to treat cyber as a business risk conceptually like any other, then the objective becomes fitting cyber into the existing enterprise risk management system so that it can be evaluated as part of an overall risk portfolio. But few board members believe that their firms have successfully done this. And about half of our interview subjects maintain that this objective is misplaced because cyber is conceptually distinct as a category of risk. This tension manifests most prominently in situations where cyber is seen as a truly existential issue or a baseline "safety" requirement that must be in place before any other risk can be appropriately addressed. It creates a *dynamic* tension because these beliefs change over time and can be highly sensitive to external events, such as a high-profile cyber incident at a competitor (or any other major organization) that demonstrates unexpected vulnerabilities and costs.

Directors that fall further toward the Enterprise Risk Management, or ERM, side of the continuum aspire to fully integrate and compare cyber risk to other risks, so as to achieve balance of

investment and attention against all of them. The common-sense intuition here rests in a tenacious awareness of the marginal value of a dollar invested in cybersecurity (rather than in the reduction of some other risk). Directors who share this mindset tend to reject claims like "cyber is different" or "we've never seen anything like this before" as self-defeating and logically suspect.

There's a recognition of roadblocks in practice to achieving ERM integration for cyber. The simplest and most important might simply be figuring out where in the organization that integration and balancing should actually happen. Many current CISOs aren't really equipped to play this role, and those that are find their structural position something of an obstacle. More likely this role would belong to a Chief Risk Officer, but many directors don't believe that their CROs (if they have one) are sufficiently well-resourced in cyber.

A second roadblock appears around the issue of how centralized or decentralized responsibility should be for the risk reduction actions the ERM model would prescribe. Some directors assert that cyber risk assessment should be integrated into product risk assessments *from the start* of a development process, in the same manner that revenue projections typically are. This is a reasonable view of how to put the concept of "security-by-design" into practice, but almost none of our interviewees (most notably those who would like it to be true) indicated that their organizations are close to that point. It is still much more common to rely on cybersecurity professionals to come in later in the process to manage the risk that other parts of the organization create, rather than partnering proactively to design for risk reduction from the start. It is also common for product design teams to push cybersecurity into the background as a secondary concern. Boards could choose to engineer a change in the overall process mindset or push the CEO and top management to do so, but there is a keen awareness of just how profound a cultural shift this would be for most organizations—and for many security professionals as well.

On the opposite side of the continuum, many directors believe that cybersecurity is a fundamentally different category of risk, above and beyond conventionally understood challenges like insufficient history. This sense of difference starts with but goes far beyond discomfort with complicated and unfamiliar technologies, as challenging as these can be for non-specialists to parse. The more fundamental difference seems to lie in the perception of third-party risk that fans out in the shape of a massive network rather than a discernible supply chain. Board members are deeply aware that cyber risk extends to this external network and acknowledge that some aspects of it will be nearly invisible, even as they ask for more elaborate mapping of the enterprise digital ecosystem. It is notable that, for some organizations, the invisibility lies much closer to the center of the network — for example, when firms discover that they do not know just how many devices are connected to their internal information systems, and/or what those devices are doing at any given time.[vi]

The digital interdependencies that affect enterprise cyber risk seem to many directors to be like an organizational chart that exists on a separate plane; no one can quite see it, and it cuts through the legal boundaries that define what is "inside" and "outside" the firm's jurisdiction as if those boundaries didn't exist. Organizations may focus on protecting their most valuable assets, but the risks to the "crown jewels" may sometimes lurk deeply in mundane places (like HVAC contractors in the Target case or tax filing systems in the NotPetya case).[vii] Of course interdependencies in enterprise risk are not unprecedented, but the distinctive complexities of digital networks make it for some directors distinctively hard or impossible to visualize.

[vi] See for example Tim Mintner's tech blog post on Tanium, "IT Operations Starts with Visibility to All Devices" revealing that he once asked a CIO how many endpoints were on the network and the CIO stated somewhere between 250,000 and 400,000 endpoints – a potential gap of 150,000 endpoints. https://www.tanium.com/blog/it-operations-starts-with-visibility-to-all-devices/ Dec 19 2018.

What should boards *do* if they regard cybersecurity as a different kind of risk? We heard several arguments. Some directors believe that boards need to prepare for worst-case scenarios and rehearse what would happen if cyber risk mitigation were to catastrophically fail, including by developing a playbook for survival. Others default to incrementalism aimed at buying time for the situation to clarify itself (though that was not generally paired with high confidence that the threat environment would get better over time). Some directors suggest searching for areas in the organization where security practices seem to be evolving more quickly — for example, in core data-related initiatives — and then scaling and extending those practices to other parts of the enterprise.

On balance, board members had low to moderate confidence in the efficacy of these ideas. Directors who see cyber as an existential risk are not quite throwing up their hands to surrender, but neither do most feel that they have a clear path forward. This point will come back later in discussions about oversight strategies and trust in management.

2. EXPERTISE

A second dynamic tension that boards must address lies in the distribution of expertise on the board and how directors access the knowledge they need to inform judgments that underpin governance and oversight. An oversimplified version of this tension would be to pose it as a blunt trade-off — should boards reserve a seat for a single "cyber expert" who brings a high level of technical knowledge and to whom other board members would turn for help. This framing is oversimplified, as there are increasing numbers of board candidates who possess both excellent technical credentials and meaningful business experience. Indeed, every director we spoke with agreed that all board members need baseline knowledge about cyber-relevant technology in order to do their work competently. But the dynamic tension is still present, as boards must weigh how much expertise is necessary and how much authority should be delegated, as well as whether it makes sense to create a board technology committee or even a cybersecurity committee.

Many board members know they are starting with some handicap when it comes to cyber. The average age of a S+P 500 director is

[vii] See Brian Krebs' "Target Hackers Broke in Via HVAC Company" explaining the Target systems intrusion: attackers first broke into the retailer's network on Nov. 15, 2013 using network credentials stolen from Fazio Mechanical Systems, a provider of refrigeration and HVAC systems; then pushed their malware to a majority of Target's point-of-sale devices, actively collecting card records from live customer transactions. https://krebsonsecurity.com/2014/02/target-hackers-broke-in-via-hvac-company/ Feb 5 2014. See also Ellen Nakashima's "Russian military was behind 'NotPetya' cyberattack in Ukraine, CIA concludes" explaining the NotPetya attack: hackers used what is known as a "watering hole" attack, infecting a website they knew their targets would navigate, in this case, a Ukrainian site that delivered updates for tax and accounting software programs; additionally, the attackers used malware that appeared to be ransomware, thus it took a few days before people realized the malware's actual objectives of permanently wiping data https://www.washingtonpost.com/world/national-security/russian-military-was-behind-notpetya-cyberattack-in-ukraine-cia-concludes/2018/01/12/048d8506-f7ca-11e7-b34a-b85626af34ef_story.html Jan 12 2018.

now about 63.[viii] That is not to say that people in their 60s are necessarily hobbled when it comes to new technology, but board members of that demographic are less likely to have had as much experience dealing with cyber risk while they were in top management positions as they would have had with other classes of risk. Directors generally agree that you do not need a PhD in computer science to govern enterprise cyber risk, any more than you need a PhD in economics to oversee financial risk; a little de-mystification goes a long way. But a lot of de-mystification happens at the intersection of experience and common language, and some directors compare cyber to the early days of learning "CFO language," which two decades ago seemed to some boards to be mysterious and impenetrable.

Given their backgrounds and experience, many directors naturally think like "good people"; they are adept at figuring out how to fix things and make things work. Yet security professionals and the cybersecurity world generally often benefit from a different mindset, as they learn to think like a "bad actor" and find ways to break things and make systems fail, which is a different mindset and culture. Scenario thinking and simulation/war-gaming can compensate for this gap and help directors grasp the shape of the risk, but board members disagree about whether such exercises are truly valuable for planning and insight, particularly if they focus too heavily on cataclysmic scenarios.

A majority of our interviewees believe that concentrating expertise, for example by reserving a board seat for a designated cyber expert, has more downsides than upsides. The arguments against such a decision are familiar. Cyber is an important issue, but it is too narrow in scope to devote a scarce and valuable board seat. (This concern was raised even by those who fully understand that cyber expertise and broader business expertise are not mutually exclusive.) There is concern that other board members might be overly deferential (consciously or subconsciously) and relinquish more decision rights than they should to a single colleague with such deep expertise. A few directors expressed concern that a board cyber-expert might not be able to avoid her knowledge-base corroding and becoming dangerously obsolete, more quickly than she or others on the board would expect or comprehend.

But there are advantages to concentrating expertise, and a vocal minority of directors pointed to these as important enough to tilt the scales. Some directors believe the issue of deference is not so problematic in practice and is outweighed by the value of having a cyber-expert colleague to whom others can turn as a sounding board. One director used a baseball analogy: a team can perform well with a mix of highly specialized and semi-specialized players, not all of whom can play every position but some of whom can play more than one; there is value at times in having a player who is world-class at only one position and cannot be moved. Some directors were skeptical about being multi-role players, worrying whether they and their colleagues can really be educated sufficiently to keep up with the rapid evolution of relevant technology and threat intelligence. Frontier knowledge about the evolving threat landscape is not principally a technical competency, but board members who prioritize the sense that attacks and penetration of their firm are essentially inevitable feel particularly strongly about the value of having direct access to threat intelligence at the board level. There is some sense that the Securities and Exchange Commission could create additional requirements for demonstrated cyber-expertise on boards, and a few directors expressed hope that this would happen as a way of standardizing expectations.

The majority of directors we interviewed believe that all board members need significant cybersecurity knowledge to do their jobs. Directors are seeking continuing education opportunities, including formal programs offered by third parties, and some report investing a substantial amount of time and effort in these. But they also express concern about whether these programs are sufficiently technical, or too broad and "high level." Are third party programs impaired by the embedded complexities and peculiarities of how cybersecurity risks present in specific organizations, so that a program which appeals to a broad range of directors turns out too general to be truly valuable to any of them? How frequently do directors need to refresh and update their education? These are natural worries for a relatively immature field where there is no clear consensus about what knowledge and understanding of risk are required (unlike in the core of finance, for example), but waiting for a consensus to emerge is not really an option.

While a significant majority lean toward distributed expertise, most board members we spoke with recognize that their beliefs and positioning around the distribution of expertise could change as the field evolves. Thus, this second dynamic tension is understood to be potentially quite dynamic despite the strong majority that leans at this moment toward distributed expertise.

3. COMPETITION AND COOPERATION

The third dynamic tension that directors confront is finding the right balance between cooperation and competition with other enterprises when it comes to cybersecurity. The publicly stated conventional wisdom on this (at least in the U.S.) almost always portrays cybersecurity as a collective good, generally for an economic sector or for a country, and sometimes for the world as a whole. Accordingly, there constantly arise new initiatives for cooperation (often in the form of information and threat-intelligence sharing) and similar kinds of joint actions that are meant to 'level up' a group of firms so that security becomes a table-stakes ingredient of license-to-operate rather than a quality that differentiates one firm from another in a competitive manner.

But why should it necessarily be that way? No law of nature declares that firms can't compete over security. An alternative logic sometimes shows up in more private conversations, captured in the aphorism, "I don't have to outrun the bear, I only have to outrun you." It is certainly possible to imagine a different kind of environment, where firms compete both privately and publicly to surpass their

[viii] See the 2018 United States Spencer Stuart Board Index https://www.spencerstuart.com/-/media/2018/october/ssbi_2018.pdf

competitors on cybersecurity just as they compete on other aspects of performance and risk management. For most of recorded history, banks have competed on the basis of how securely they protect money, so why should they naturally cooperate now when it comes to protecting data?

Most directors we interviewed see the answer to that provocation as nearly obvious, but not necessarily definitive. The majority view contends that cybersecurity should be treated as a collective good, and that competition would not be an appropriate way to enhance the performance of individual firms or, more importantly, the ecosystem of firms that depend to a certain degree on each others' security. Trying to outrun your competitor rather than working together to outrun the bear is seen by these directors as perilous: shunting off risk to someone else might come back to haunt you later as counter-party risk in finance, and other kinds of interdependencies in sectors like health care. To publicly claim a competitive advantage in cybersecurity paints a target on your back for determined attackers; that requires a level of confidence that does not exist outside a very small number of firms. A few board members expressed concern that antitrust and competition law and policy are actually creating barriers to more active collaboration and provision of collective cybersecurity goods. They see cybersecurity not as an opportunity space where a firm can excel, but as a basic license-to-operate issue for the firm, the sector in which it sits, and possibly for the economy as a whole. The aphorism that captures this point of view is "security isn't a feature, the absence of security is a flaw."

This view is not unanimous however, and a minority of directors point out that increasing the level of competition in cybersecurity might not be all bad. This perspective starts with an honest reckoning that there is not now a level playing field among organizations in cyber, and there almost certainly will not be a level playing field anytime soon. In that context, it would be desirable to imagine leveling everyone up to match the best performers, but there is also an awareness that a best-practice mindset can sometimes lead to "averaging" or converging behaviors that bring the top performers down (or least slow their forward progress). Directors who articulate this perspective generally believe that even the top performers have a long way to go on cyber, and thus see value in competitive pressure to innovate, accelerate, and improve on cyber as in other domains. They certainly do not want to see that competitive pressure diminish, even as they recognize some of the risks associated with it.

The dynamic tension between cooperation and competition is not "either/or." In practice, the mix is more subtle and fluid than it appears. For example, many directors express a desire for their firm's CISO to be held accountable as a value creator, not just a defender, and that implies some degree of competitive differentiation. A few board members (not only from cybersecurity and/or professional services firms) are considering leveraging internal corporate cybersecurity capabilities into their external service offerings, within their sector or perhaps even more broadly. On the "cooperation" side of the continuum, data- and threat-intelligence sharing initiatives have made uneven progress in part because some elements of what might be shared touch closely on the competitive assets firms seek to protect and control. Thus, the modal position for this dynamic tension leans toward the cooperative end of the spectrum, but not definitively, and even that position might not be stable in evolving threat and market environments.

4. BOARD-CISO-MANAGEMENT PROTOCOL

The fourth dynamic tension that boards navigate relates to how they choose to structure their oversight relationships on cyber within a higher-level model of information flows between management and the board. Another way to describe this is to dig one level beneath the conventional question about "how should the CISO talk to the board", to focus on what the objectives of that communication really are, and what communications protocols are needed to make it happen. Plain-speak translations of technical concepts by CISOs and clear discussion of what is (and is not) being done within the organization are only the starting point.[ix]

Beyond that baseline lie the dynamic tensions that most often manifest in the debate around what kinds of metrics the CISO should share with the board. Several of our interviewees posed the question as, "what does good really look like here," as they hope someone will be able to definitively answer that question, if not now then soon. Metrics are essential for tracking progress, and many board members understandably yearn for a standardized set of metrics that reflect a best practice discipline so they can be compared not only within their own firm over time, but also across firms within a sector or more broadly. Many of the directors express a looming discomfort that the metrics used by their organizations are deeply problematic in some respects and not sufficient for true insight.

Directors who serve on multiple boards confirm that reporting on cyber is highly variable, and there has been less convergence than they have hoped (and in some cases expected). Developing metrics that make sense and are useful at a moment in time is one challenge; ensuring those metrics remain relevant as the risk and threat landscapes change is an even greater challenge, one that many directors feel only the largest and most sophisticated organizations are in a position to do.

There is a clear desire for stable and standardized metrics, but also a lurking realization that rushing toward convergence too soon could be a net negative for effective governance and oversight. As in many other risk domains, the demand for "dashboards" that can easily illustrate progress (or lack thereof) has potential to lead to oversimplification and, in some cases, concerns about "gaming" the index. No board member we spoke with alleged that they were being intentionally misled; they simply acknowledged that managing to an index could be both effective and dangerous at the same time.

[ix]It's worth pointing out that not all board members feel their firms have yet achieved that goal.

Directors who see more potential risk tend to focus on speed of change and adaptive adversaries as core challenges. Concrete and stable metrics are "backward looking" and implicitly provide an attack roadmap for adversaries, or at the very least a landscape for arbitrage. Most board members place trust in CISOs to manage that dilemma, but some express a desire to have their CISO help them understand how the future risk and threat environments could be distinctly different from the past. Specifically, they desire more storytelling in the form of approaches like scenario and table-top exercises, but these are built out of a very different kind of communications protocol and CISO relationship than with standard and stable metrics.

Sometimes these issues manifest in questions about the key characteristics that organizations should look for in an outstanding CISO. Should we select a CISO for business savvy, skilled communicators who deeply understand the business model of the firm and can naturally integrate the security organization into that strategic story? Or should we select for deep technical skills, combined with security paranoia and single-minded obsession with protecting the organization's digital layers? The easy answer is to say "both," but individuals who possess a combination of these qualities are rare.

A subset of directors expressed concern that existing metrics lean too heavily toward showing how much progress has been made against yesterday's threats and unintentionally support the tendencies of some CISOs to talk "happy talk" to the board about how well the security organization is performing. Many directors want regular third-party input to keep the performance bar set at the right point, for fear that what was good enough this year may be below average next year. Some directors want to address this by looking for a much more granular and intimate relationship not only with the CISO, but also with employees further down the organization. They want to hear what keeps the CISO up at night and how things could go wrong. Building these relationships requires a "high-touch" approach, above and beyond informal dinners and other ways of talking with the CISO outside of formal board meetings. Some directors want more "governance by walking around," i.e. they want to go down to the security group on a normal day, sit on the corners of peoples' desks, and hear their stories about what they are doing and what they are worried about.

That kind of communication does not scale readily, and is subject to its own possible misinterpretations and distortions. It can be hard to translate such interactions into the kinds of high-level governance and oversight decisions that boards ultimately need. But a surprising number of directors feel right now that they might not be able to do their jobs well on cyber without more of this kind of input.

Next Steps

Board governance and oversight of cybersecurity risk are part of a new and rapidly evolving discipline. There is no reason to believe that the rate of change in the risk landscape is going to slow down anytime soon. It is more likely to accelerate, with on-going digital transformation programs, the introduction of artificial intelligence and machine learning technologies along with cloud and 5G networks that will enable a new generation of IoT architectures and more, and of course the evolving capabilities and strategies of adversaries, both criminal network and state-based.

This is why we were struck by a few areas of silence—concerns we raised in our conversations that we thought we would hear more about from board members, that were surprisingly rare. Like the "dog that didn't bark" in the Sherlock Holmes story, these silences are clues about where directors are focusing right now—and where directors may collectively have blind spots.

For example, we did not hear about foreseeable technological discontinuities, such as quantum computing or machine-learning breakthroughs, that might reshape the game. We did not hear about "security by design" as an aspirational practice for unlocking the value-creation potential of digital security programs. We did not hear much about basic culture and cultural changes — within the firm and more broadly— that could re-orient how people (employees, customers, and others) engage with digital technologies and how they practice (or do not practice) security behaviors. We also did not hear much about the current evolution of attack modalities from data breaches, ransomware, and information theft toward business disruption by data manipulation, deep fake videos (e.g., fake CEO speeches), and other emerging technologies.

This last point highlights that, just as the attack surface is growing in *size*, it's also changing *shape*, as the attack taxonomy and strategies for doing harm to the enterprise morph. Are these new attack modalities the responsibility of the CISO, and does the CISO have sufficient authority to manage these risks? For example, should the CISO be responsible for monitoring and protecting the firm against fake media and other increasingly sophisticated disinformation campaigns? This was an important "dog that didn't bark." Also strikingly absent from our conversations was the word "innovation," which was probably the loudest silence. We are confident that if we had interviewed a group of prominent cyber-criminals and nation-state agencies with offensive cyber roles, discussions about rapid and ambitious innovation would have been much more prominent.

These reflections on the "dogs that didn't bark" are an observation, not a criticism of where directors are focusing now. They point to the profound sense of urgency felt by board members to play better defense in cyber, and that urgency is clearly justified. But we also believe that boards need to spend more time and effort at the innovation horizons inhabited by adversaries. At some point, and possibly quite soon, it will be reasonable for boards to pivot toward a more aggressive mindset, and to embark on a more ambitious path towards cybersecurity oversight and governance. To that end, we put forward some aspirational principles that boards might consider going forward.

PRINCIPLES TO AMPLIFY THE UPSIDE

UNDERSTAND CYBERSECURITY AS AN ENABLER AND ACCELERATOR FOR DIGITAL TRANSFORMATION:

It is a problem when cybersecurity does not receive sufficient attention, but it is equally a problem when unmanaged security risks prevent firms from doing things that have potential to amplify the upside of digital transformation. The opportunity costs of actions not taken, products not created, and markets not entered because of unmanaged security risk are mounting, even though they are hard to quantify. Boards need to hold management to a standard in which cybersecurity and innovation are fully consistent with each other, if not synergistic. This means increasing boards' focus on imagining positive cybersecurity futures, not just fear-provoking, worst-case scenarios.

PROACTIVELY SHAPE THE REGULATORY ENVIRONMENT:

Regulation of cybersecurity- and privacy-related processes and performance is set to expand, and to simply react and comply is an overly passive position for most boards to take. Regulation often responds to dramatic events, but it is the underlying day-to-day "grey war" aspect of cybersecurity practice that is a more important determinant of long-term trajectories. In today's political environment, it is likely that regulation will take shape around the overreach of a few specific digital sectors — in particular, social media and the major internet platform businesses — rather than around the needs of the rest of the economy. Boards need to commit to an ongoing engagement with regulators in which both sides explain what they are seeing and what they need to improve the overall cybersecurity environment, rather than play whack-a-mole to stave off the latest form of attack. They also need to develop and promote internally and externally a pragmatic distinction between privacy and cybersecurity, establishing clear roles and responsibilities for CISOs and Chief Privacy Officers when their domains intersect.

INVEST IN DIGITAL TRUST:

The broad sense that customer and societal trust in digital systems is corroding poses a significant risk to firms' license to operate. Cybersecurity seen as table-stakes is a foundation for digital trust, but shifting perceptions even further could have a greater impact. Boards should be pressing for higher levels of ambition here. Incident response can be an opportunity to enhance trust; for example, Tylenol was a stronger brand following its response to the 1982 poisoning crisis than it was before. Positive externalities of security investments in the supply chain and the overall ecosystem (including investments in customer security) can be an opportunity to enhance trust. Cultural change that pushes cybersecurity-thinking down into product-level processes and treats people who touch the firm's ecosystem from multiple directions as responsible partners (rather than compliance machines or uncontrollable risks) is an opportunity to enhance trust.

The landscape for the board's responsibilities in cyber ultimately looks more like a "wicked problem" than a well-understood "optimization across trade-offs" problem.[x] Wicked problems are difficult to solve because the definition of the problem itself sits inside incomplete, contradictory, and changing requirements. In this context, the understandable yearning for a consensual scorecard, or a list of best practices that ground cyber governance in stable certainties, is likely to remain unfulfilled for some time to come. It would be premature to assert that one end of the dynamic tensions spectrum is "right" or better for all organizations, or even for a defined subset.

Although there is no single right answer, there are certainly right actions to take that can multiply the upsides and de-risk the downsides of choices a board will make. We present here a framework of suggested actions tied to the four dynamic tensions that can shape board governance processes. Keep in mind the overarching rationale, which is that boards can legitimately choose to lean in either direction on each of the dynamic tensions at a given time. The appropriate test is to make sure that the choices are *intentional, consensual* (as much as possible), and *adaptive* (re-assessed, perhaps annually).

RECOMMENDED PRACTICES FOR BOARDS TO CONSIDER

DYNAMIC TENSION 1: RISK MODEL

Leans "Enterprise Risk Management (ERM) Ingredient"

- Interrogate the balance between managing risk to *amplify the business upside* and managing to de-risk the downside for loss
- Promote cyber risk to a high level within the hierarchy or taxonomy for enterprise risk
- Harmonize enterprise risk management by clearly defining the expected role of the chief information security officer vis-a-vis others, such as the chief risk officer

Leans "Existential / Safety"

- Integrate cyber risk management into *early phases* of product and service design and development processes
- Prioritize due diligence of cyber-risk in the supply chain and for M&A activity
- Develop a culture of preparedness and stress-testing, including *semi-worst case scenarios*

DYNAMIC TENSION 2: EXPERTISE

Leans "Everyone"

- Ensure adequate training and education is defined, used, and *kept up-to-date*
- Engage external third-party expertise for specialized knowledge, and most importantly to prevent *group-think* traps
- Amplify accountability for cyber oversight in subset groups (likely committees)

Leans "Cyber-Expert"

- Seek out specific board members who offer deep specialized knowledge of cyber (e.g., crisis management, technology, and threat landscape)
- Prioritize full board discussion of cyber oversight over committee delegation
- Engage external subject-matter experts to *test and enhance* internal expertise

[x] Horst W. J. Rittel, and Melvin M. Webber. "Dilemmas in a General Theory of Planning." Policy Sciences, vol. 4, no. 2, 1973, pp. 155–169. JSTOR, www.jstor.org/stable/4531523.

DYNAMIC TENSION 3: COMPETITION AND COOPERATION

Leans "Competitive Assets"

- Ensure that security is clearly differentiated, with clear boundaries between what can be shared and what is proprietary (critical competitive assets)
- Seek and assess *return on security investments*, above and beyond protection/insurance on other investments
- Integrate privacy and security-by-design in product development and deployment

Leans "Collective Good"

- Actively invest in cyber information sharing capabilities across public and private sectors
- Assess progress toward *ecosystem health* in addition to firm performance ('herd immunity')
- Invest to improve the security of your supply chain
- Engage *proactively* in policy and regulatory development

DYNAMIC TENSION 4: BOARD-CISO-MANAGEMENT PROTOCOL

Leans "High Touch"

- Engage on security oversight with business units and other enterprise functions *beyond the Security Operations Center*
- Integrate quantitative and qualitative inputs in a consistent manner
- Deepen the "trust but verify" relationship with the CISO through more frequent interaction outside the boardroom
- Create and defend protected spaces for management and employees to expose cybersecurity challenges

Leans "Arms Length/ Stable Standard Metrics"

- Allow metrics to evolve, but choose a consistent framework for how metrics can support oversight
- Re-assess on a regular basis what really needs to be measured (such as *impact*, not just efficiency)
- Supplement quantitative metrics with integrated qualitative aspects (where possible) in a *balanced scorecard* mindset or model

It has become common at the end of the 2010s to talk about seeking *resilience* in cybersecurity, above and beyond defense and protection. The meaning of resilience from ecology is that a system, when perturbed or attacked, does not just return to its original state (which is what a successful defense promises, and what the word "robust" implies). Resilience is actually a higher bar, as it means that the system evolves in response and emerges *stronger* than it was prior to the challenge.

Immune system analogies can be banal at times, but in this context the analogy is apt. The adaptive and evolving nature of the immune system is precisely what makes organisms resilient in an environment filled with evolving pathogens. Resilience of that kind is an ambitious aspiration for cybersecurity, but the ambition is absolutely needed in order to deal with a similarly adaptive threat.

The same principle of resilience should operate in cybersecurity governance and oversight. A robust defensive framework would be better than what many boards currently believe they have. But it is still inferior to resilient governance, where frameworks and processes get stronger in the wake of challenge. Given the urgency to catch up to the present threat landscape, some directors may feel this is a bridge too far. But a resilient cybersecurity organization that improves its performance to keep up with threats will ultimately depend on resilient governance from the top.

CLTC
Center for Long-Term Cybersecurity
UC Berkeley

BILL PHELPS

Executive Vice President - Commercial Business Lead, Booz Allen Hamilton
Phelps_Bill@bah.com
Office: 202.346.9816

ANN CLEAVELAND

Executive Director, Center for Long-Term Cybersecurity
ann.cleaveland@berkeley.edu
Office: 510.664.7506

STEVE WEBER

Faculty Director, Center for Long-Term Cybersecurity
steve_weber@berkeley.edu
Office: 510.664.7506

In partnership with
Booz | Allen | Hamilton®