McKinsey&Company

McKinsey on Risk

Highlights

Number 4, January 2018

McKinsey on Risk is written by risk experts and practitioners in McKinsey's Global Risk Practice. This publication offers readers insights into value-creating strategies and the translation of those strategies into company performance.

This issue is available online at McKinsey.com. Comments and requests for copies or for permissions to republish an article can be sent via email to McKinsey_Risk@McKinsey.com.

Cover image:
© Steven Hobbs/Stocktrek Images/Getty Images

Table of contents

Introduction

Welcome to *McKinsey on Risk*, the journal offering McKinsey's global perspectives and strategic thinking on risk topics. We focus on critical decision making in the key risk areas that bear upon the performance of the world's leading companies. From our experience and expertise across sectors, we seek to help institutions better support their business objectives by transforming the way they manage risk.

The threats and opportunities companies face as they develop their risk strategies are increasingly complex and interconnected. They recall the dangers and rewards of the tiny spaceship *Cassini* as it navigated the perilous rings of Saturn at speed. To protect business value and promote strategic growth, risk management must be based on a contextual understanding of the connected risks. Undeniably, digital innovation will underlie this understanding. The digital strategies that companies in all sectors are adopting to capture business value must also be exploited to enhance risk management. In this fourth issue of *McKinsey on Risk,* we offer a selection of articles that apply a risk lens to these new powerful approaches, highlighting advantages and pitfalls.

We begin with a consideration of risk detection and how financial institutions, in collaboration with regulators, can apply digital methods to improve the accuracy of controls while lowering costs. The application of these methods, especially advanced analytics and machine learning, is the subject of further articles addressing anti–money laundering and financial fraud—among risk detection's most important disciplines. The new approaches come with risks of their own, however, as discussed in a forward-thinking piece on algorithmic bias and how companies can combat its distorting effects on business decision making. The related concern of data protection and the regulatory attention it is attracting are considered in a piece focusing on a forthcoming comprehensive European data law. Another article recommends an integral approach to cybersecurity, a topic of universal business importance, by which companies build a culture of resilience to protect their critical digital assets. The concluding piece lays out the strategic, business, and technical adjustments banks should make to ensure improved capital management under the proposed Basel IV regulatory framework.

We hope you enjoy these articles and find in them ideas worthy of your consideration. Let us know what you think at McKinsey_Risk@McKinsey.com. You can also view these articles and others, as well as previous issues of *McKinsey on Risk*, at McKinsey.com and on the McKinsey Insights app.

Thomas Poppensieker
Chair, Global Risk Editorial Board

The neglected art of risk detection

At the core of risk management is risk detection, an art that can be skillfully improved if banks and regulators accept new analytical methods.

Piotr Kaminski and Jeff Schonert

The modern risk-management framework generally relies on the "three lines of defense" scheme, with the businesses, control functions, and audit as the first, second, and third line, respectively. The concept borrows from the language of military strategy, in which intelligence plays a key role. For risk management, intelligence means effective detection: to prevent the bank's reputation, liquidity, and capital position from being harmed, the lines of defense must detect risks early.

Detection is fundamental in risk management, embedded in its activities and processes. Credit scoring, for example, is a tool for detecting potential borrower-default risk at the application stage, while customer due diligence is designed to identify high-risk customers during the onboarding process, as part of the bank's know-your-customer (KYC) program. Risk managers are practicing the art of detection when they identify instances of fraud, spot a drifting investment strategy in an asset-management business, monitor their network's end points to locate cyberintrusions and data theft, or identify potential rogue traders.

Most executives and risk professionals will quickly acknowledge the basic importance of detection. Yet the efficacy of detection—and the levels of "detection risk"—vary widely among risk disciplines and from bank to bank. With poor detection, threats can rise to existential proportions, as some of the world's largest institutions have learned in recent years. Weak detection capabilities can be costly. Manual controls, for example, are not especially effective and yet they always cost more than automated controls. Poor detection can result in high levels of false positives and the needless diversion of valuable risk resources.

Assessing control effectiveness

Most banks manage operational and compliance risks through detection processes. Accordingly, inherent risks are classified by their likelihood and severity. The effectiveness of the controls is then evaluated, usually on a three- to five-point qualitative scale, with such ratings as "unsatisfactory," "satisfactory," and "strong." In more advanced approaches, the effectiveness of the control is subtracted from the level of the inherent risk, producing a measure of residual risk. For example, an inherently high risk of noncompliance with the Truth in Lending Act can result in a low residual risk if the controls are considered strong (as when customer disclosures and redisclosures are automated).

This type of assessment is frequently deployed as part of the bank's risk and control self-assessment and independent risk assessment for operational and compliance risks. Used judiciously by trained frontline and risk personnel, the approach can yield valuable insights into the control environment. If applied mechanically, results will be less helpful. Suboptimal outcomes are often caused by the inadequate assessment of control effectiveness, including imprecise testing for accuracy. Without knowing how well their controls are detecting true risks and differentiating them from false positives, banks will be unable to identify gaps in control effectiveness and may have no choice but to add costly layers of controls.

Probability theory

Highly illustrative of the problem of accurate detection are diagnostic tests for diseases, which must account for the potentially high number of false positives resulting from relatively sound tests for rare conditions. Two primary parameters determine the reliability of such tests:

1. *Accuracy* is the probability that a sick person tests positive for a given disease. It reflects the sensitivity of a test in measuring the percentage of people predicted to be sick who are actually sick. A test that is 99 percent accurate means that if it is performed on 1,000,000 sick people, 990,000 will test positive for the disease.

2. *Specificity* is the probability that a healthy person tests negative. A test with 97 percent specificity means that if it is performed on 1,000,000 healthy people, then 970,000 will test negative.

A rare disease might have a frequency of .01 percent, affecting 1 person in 10,000. If the test to detect it is 99 percent accurate and 97 percent specific, then for every 1,000,000 subjects tested, 99 of the 100 who actually have the disease will be correctly diagnosed. At 97 percent specificity, however, the test will also incorrectly diagnose 29,997 of the 999,900 healthy individual as having the disease. *For those who test positive, therefore, the chances that they actually have the disease are less than one-third of 1 percent*. That this probability should be so small is counterintuitive and even astonishing. The unmitigated consequences can be devastating: healthy people might believe they have deadly conditions; qualified job applicants might be rejected for assumed drug use. For banks, the consequences of poor risk detection can be seriously damaging as well.

False positives and risk management

Banking executives and risk practitioners seeking to detect and prevent low-frequency events will recognize the problem of false positives. In anti–money laundering (AML), for example, a monitoring system is usually deployed that produces alerts on

atypical transactions. These are referred to a financial investigations unit (FIU) comprising experts who often have a background in law enforcement. Based on certain criteria, they attempt to identify likely instances of fraud from among the alerts and accordingly file suspicious-activity reports (SARs) with the appropriate authorities.

Should a transaction-monitoring control detect suspicious activity with 95 percent accuracy and specificity, 5 percent of the activity it determines to be legitimate or suspicious will not actually conform to the established criteria. If 0.1 percent of transactions truly do meet the criteria for suspicious activity, then this particular control could produce a false-positive rate of over 98 percent. Fewer than 2 percent of alerts will correspond to activity that upon further examination will qualify as suspicious. The FIU investigators will have to spend a lot of time investigating cases that do not qualify as suspicious, leading to a low conversion of alerts into SARs.

In practice, controls may be even less specific. If the control in the example above were 75 percent specific, more than 99.5 percent of transaction alerts would be false positives. Increasing the accuracy of the control to 99.9 percent will not reduce the false-positive rate significantly (the false-positive rate would remain above 99.5 percent).

The implications of inadequate control specificity

Improving control performance demands increased focus on specificity. A control that detects only 50 percent of positives, for example, but is highly specific—incorrectly signaling a positive for only 0.1 percent of negatives—would have a false-positive rate of 67 percent. If the specificity were improved from 0.1 to 0.01 percent, the false-positive rate would drop to 17 percent.

Good detection is not simply about reducing false positives. Controls must also be highly accurate, detecting a large percentage of positives. But equal attention must be paid to control specificity for the control environment to perform optimally. In AML, the objective is the accurate identification of transaction patterns associated with illegitimate activity. This implies the ability to distinguish such patterns from those originating with legitimate clients. To reach this objective, control assessments are vitally important. Unfortunately, many control assessments are merely qualitative or unable to differentiate between control accuracy and specificity.

Inadequately specific controls cause valuable resources to be diverted from actual risks. Fraud investigators are taken away from vital work, such as identifying connections between cases—"connecting the dots"—to detect networks of criminal activity. The problem of false positives is more than a matter of cost control. While regulators at times take a favorable view of increased spending on controls, the addition of manual controls is not always the best way to resolve detection issues. Banks should be focusing on improving the effectiveness of the control environment in critical risk areas—an approach that can also lower spending on manual controls.

Making progress in key areas

Leading banks are making progress in risk detection in several areas.

Anti–money laundering

AML activities are triggered by alerts generated by rules-based binary criteria. The alerts, investigated manually, usually have very high false-positive rates. Banks have discovered that tighter segmentation of alerts, the use of KYC data, and the admission of additional variables can improve the specificity of AML controls. In one example, the false-positive rate was cut in half with the use of additional data on internal transactions (see sidebar, "Deploying AML resources where they are most needed").

Deploying AML resources where they are most needed

At one large US bank, the false-positive rate in anti–money laundering (AML) alerts was very high. The remedial process involved a two-stage investigation. One team would determine whether an alert was truly triggered by suspicious activity. It would eliminate clearly false positives and pass on the remainder to experts for further investigation. Very few suspicious-activity-report filings resulted.

The bank rightly felt that this elaborate procedure and meager result was overtaxing resources. To improve the specificity of its tests so that AML expertise could be better utilized, the bank looked at the underlying data and algorithms. It discovered that the databases incompletely identified customers and transactions. By adding more data elements and linking systems through machine-learning techniques, the bank achieved a more complete understanding of the transactions being monitored.

It turned out that more than half of the cases alerted for investigation were perfectly innocuous intracompany transactions. With their more sensitive database, the bank was able to keep the process from issuing alerts for these transactions, which substantially freed resources for allocation to more complex cases (exhibit).

Exhibit

One bank used enhanced data and analytics to dramatically reduce false positives in anti–money laundering activities.

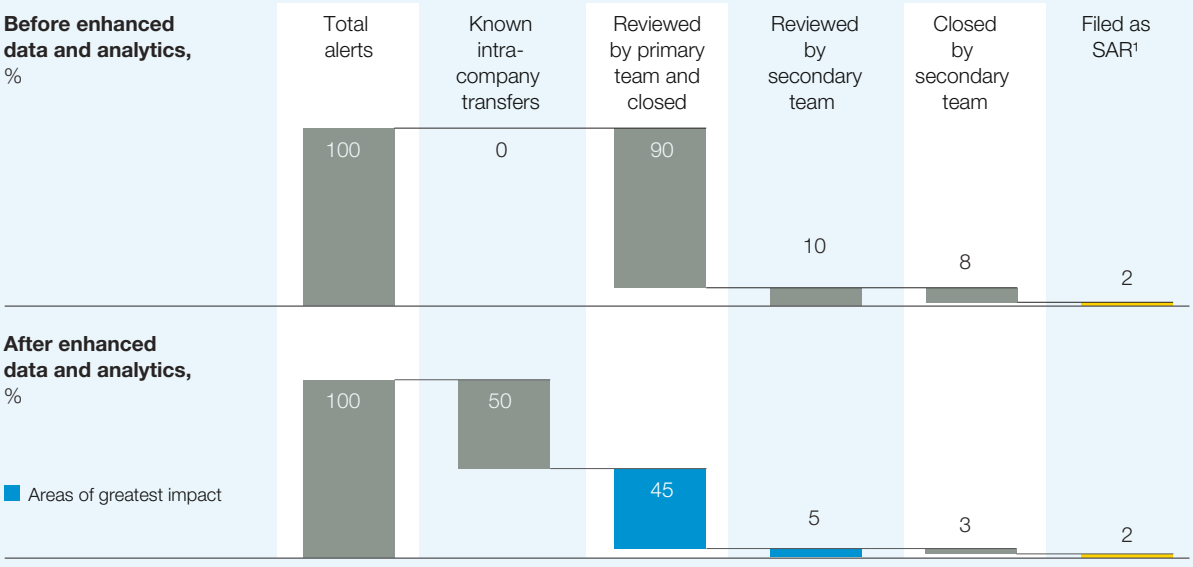

[1] Suspicious-activity report.

Source: McKinsey analysis

Compliance testing and monitoring

Banks frequently have difficulty determining whether compliance controls are truly effective at reducing noncompliant outcomes. One reason is that controls are often designed to ensure that internal procedures are followed, rather than to detect and prevent noncompliant outcomes. Having recognized the issue, some institutions are exploring ways to enhance the detection of compliance defects with system-generated process data—such as time stamps, queue status, work-flow history, and transaction attributes. In mortgage servicing, for example, institutions are increasingly using system data to achieve greater accuracy and specificity in their compliance controls.

Cybersecurity and fraud

Detection tools used to flag potential cyber-intrusions and unusual network and system activity generate alerts, which, like their AML equivalents, require further investigation. Many companies invest heavily in detection capabilities without giving sufficient consideration to how accurate and specific these alerts are in practice. The return on security investments should be the ability to separate "noise" (false alarms) from "signal" (real threats). One company is now pooling its security data into a data lake and exploring all potential correlations among the incidents and events that trigger alerts. The result is a material improvement in the signal-to-noise ratio and faster response times by company security analysts. This approach to reducing false positives and redirecting resources to real threats can be enhanced by adding external data sources, such as threat feeds or "dark web" scanning.

Consumer-credit underwriting

Most consumer-credit decision algorithms in the United States and Canada use a combination of credit scores and binary exclusion criteria. The credit scores correspond to the expected probability of nonpayment as derived from credit-bureau data, while the exclusion criteria might include such data points as "no bankruptcy in the last five years." The binary data are usually effective in filtering out high-risk borrowers but can produce a large number of false positives—credit applicants will be rejected who would have actually performed well. One credit-card executive referred to his firm's process as a "meat cleaver" that tended "to chop off" many potentially good customers. In response, the most sophisticated credit-card issuers are reducing the false-positive rate by improving the specificity of their proprietary credit scores and the binary exclusion criteria. This approach allows them to book profitable accounts that were deemed too risky by less advanced lending criteria.

Credit collections and portfolio management

The effectiveness of credit-collection strategies depends on the ability to identify high value-at-risk borrowers—those with significant credit balances that are also likely to be charged off. Often the best approach to these borrowers is to offer settlements and partial-payment programs early in the collections process. Obviously, such an approach cannot be deployed indiscriminately or it will result in unacceptable losses. Highly specific and reliable early detection of these borrowers is the goal, but it has been difficult to achieve in practice. Banks are beginning to explore advanced analytical techniques, such as random-forest algorithms, which can produce higher specificity than traditional logistic-regression models. To detect borrowers with high value at risk (VAR), one institution replaced a traditional VAR model with a machine-learning model and improved detection performance by 60 percent. In combination with contact-and-offer strategy models, the new analytics-based approach has significantly increased efficiency of the agents.

Credit portfolio management

In credit portfolio management, the early identification of borrowers at risk of imminent default is highly advantageous for retail and commercial lenders. This capability permits banks to respond by reducing credit lines, securing additional collateral, or modifying loans. Early-warning signals must be highly specific or unacceptable levels of false positives would result, outweighing the benefits of the approach. The costs of false positives include loss of revenues due to lower balances and customer attrition; in the long run, the institution's reputation can suffer. Lenders are already exploring ways of predicting defaults with greater accuracy, such as with machine learning applied to market data (such as credit-default-swap spreads) as well as unstructured data (such as sentiment analysis).

Five actions can improve risk detection significantly

Accurate detection is an essential capability for robust risk management. Institutions cannot improve detection effectiveness and efficiency overnight, but experience shows that meaningful progress can be achieved in 9 to 12 months. To be successful, near-term transformations of the control environment should include the following five actions.

1. Reviewing the control framework

This review improves the control environment by identifying gaps and inefficiencies and taking remedial action. Unnecessary or ineffective controls are culled, as are some manual controls or controls based on procedural adherence. The required accuracy and specificity of controls in high-risk areas, such as AML, fraud, and cybersecurity, should be determined by the frequency of the underlying risks. Throughout the review, testing based on residual risk or outcomes should be promoted over control testing, especially where results have been statistically unreliable.

Manual controls may be symptomatic of high false-positive rates in detection processes and therefore warrant close analysis. The control review should focus on replacing or augmenting manual controls with system-driven detection algorithms. Risk and control assessments, in particular those based on subjective or abstract criteria, may simply be unreliable, creating a false sense of security. Such programs should be evaluated for their efficacy relative to the nature and frequency of the risks. Assessments that do not enhance resilience of the bank will have to be dealt with critically. Resources freed by improvements (for example, the elimination of unnecessary manual reviews of false positives) should be deployed to critical areas.

The control review must be executed with good judgment to ensure that it increases control effectiveness. It cannot be approached as a template-driven, mechanical exercise: the risks are too high and regulators will be watching closely.

2. Changing the detection paradigm

Transformative change requires a fundamental shift in strategy. Banks should move beyond the detection of individual suspicious activities to detecting clusters of such activities. In AML and fraud, this means identifying the bad actors as opposed to focusing predominantly on potentially suspicious individual transactions. To do this effectively, banks will need to acquire more data sets. This will allow them to filter out more noise—the false positives—and to create risk scores that achieve better predictive power than binary detection criteria. Investigators of security threats may flag the purchase of fertilizer or the renting of a truck as warning signs of a potential terror attack, but they must also account for the fact that the vast majority of these transactions are completely legitimate. In our experience, traditional, rules-based detection methods in AML reach their potential for reducing false positives at around 90 percent. To go further,

banks and regulators alike must reframe the problem statement and apply advanced-analytics solutions to look for networks of events.

3. Applying advanced analytics and automation

Improving detection by replacing ineffective and expensive controls will require that banks develop significant capabilities in advanced analytics and automation. Where suitable, machine learning should also be integrated into existing analytics capabilities. Institutions can even set up a dedicated machine-learning factory as long as they guard against it becoming a "hammer looking for nails." Analytics efforts must follow practical necessity and not create problems to solve. The "decision science" model used by credit-card issuers to support credit, marketing, and collections strategies, for example, is a technique-neutral approach. The decision trees, logistic regression, and other modeling techniques it may employ are selected based on their applicability to a specific detection problem.

One misconception about advanced-analytics techniques is that they require vast quantities of high-quality data. While some techniques (such as neural nets) do require a lot of data, many do not. Furthermore, many very productive advanced-analytics methods thrive on unstructured and imperfect data. They can even help make the data accessible for more traditional techniques. Banks must always seek to improve data quality, but perfect data is a quixotic goal—unattainable in the near term and never cost-effective. With advanced analytics, data-quality issues are a fact of life, and banks must deploy measures to account for them.

Problems of model validation give pause to even the most committed proponents of advanced analytics. Current model risk management approaches have been honed on relatively well-understood regression and decision-tree techniques. For more complex machine-learning models, banks are still developing standards. Spurred by difficulties with validation revenue models for regulatory stress testing, some banks are starting to define universal principles of forecasting and modeling techniques that could be applied to nontraditional and advanced methods. This could open the way to a more flexible—but still policy-driven—model-validation approach.

4. Developing a portfolio of use cases and matching processes

Banks need to develop and manage a portfolio of use cases. The program should be designed to encourage expert creativity while ensuring a balanced portfolio. The use cases should address a diverse set of risks, with a range of probabilities and potential impact. Good governance is important, as the actual impact of the use cases will need to be validated against initial assumptions; furthermore, the feasibility of implementation must be assessed in light of regulatory requirements, system implications, and operational impact. A significant portion of the use cases developed for compliance and operational risk should contribute to simplifying and strengthening the control environment. The effort will generate demand for advanced analytics and automation in this area. Examples of potential use cases include monitoring employee conduct, contract compliance, and payment-fraud detection.

5. Engaging with the regulator

To improve detection, rationalize controls, and strengthen risk management, appropriate regulatory engagement is required. The conditions for such engagement may not yet be in place for banks addressing enforcement actions or major regulatory enhancements. Nevertheless, all regulatory issues need not be resolved before banks begin this program. Some banks are deploying advanced techniques in place of manual solutions as part of their major regulatory programs, including resolution planning and Comprehensive Capital Analysis and Review.

Banks might be surprised at regulator reaction to their plans to rationalize controls. Efforts to improve detection effectiveness in KYC and AML, for example, may be welcomed, given recent high-profile detection failures. In the approach we have been discussing, efficiency gains and greater effectiveness are closely linked. The business case for rationalizing expensive manual controls is based on better detection and risk management. Getting comfortable with efficiency gains as part of the business case for better detection is a requirement for success.

Nowhere is regulatory dialogue more important than in model risk management. Input from the regulator is required to meet the challenges posed in the validation of sophisticated models. Banks may therefore want to focus first on machine-learning models used to detect fraud and money-laundering activities before tackling models affecting consumer access to credit. With such checks in place as model performance controls and parallel runs, the models for detecting fraud and money laundering may be cleared for testing and deployment more readily.

■ ■ ■

Regulatory and competitive pressures as well as rising business costs are driving banks to improve the effectiveness of their risk controls. To reduce unmanageable levels of false positives—and all the added work they entail—leading banks are developing significant advanced-analytics capabilities and automating costly manual steps. Complex risk-detection problems, such as those involved in model validation, are inspiring innovative approaches to improve control standards and resource allocation. As the early movers are discovering, these investments lower operating costs while returning business, compliance, and reputational dividends. To make further progress, banks and regulators should consider the limitations of existing detection approaches and be open to applying methods that enhance results. ■

Piotr Kaminski is a senior partner in McKinsey's New York office, where **Jeff Schonert** is an associate partner.

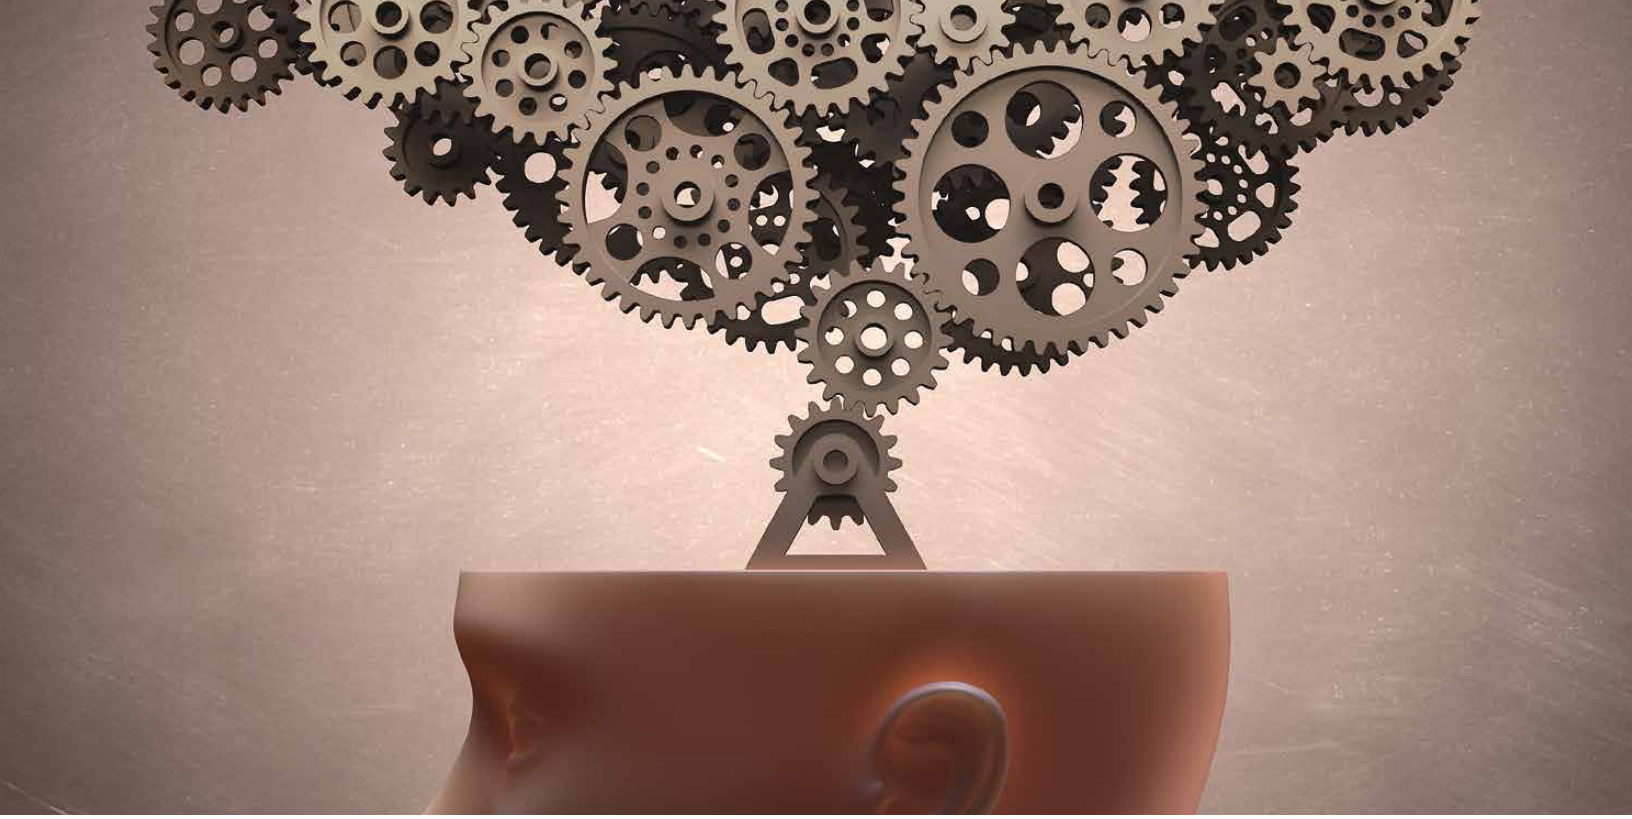

Controlling machine-learning algorithms and their biases

Myths aside, artificial intelligence is as prone to bias as the human kind. The good news is that the biases in algorithms can also be diagnosed and treated.

Tobias Baer and Vishnu Kamalnath

Companies are moving quickly to apply machine learning to business decision making. New programs are constantly being launched, setting complex algorithms to work on large, frequently refreshed data sets. The speed at which this is taking place attests to the attractiveness of the technology, but the lack of experience creates real risks. Algorithmic bias is one of the biggest risks because it compromises the very purpose of machine learning. This often-overlooked defect can trigger costly errors and, left unchecked, can pull projects and organizations in entirely wrong directions. Effective efforts to confront this problem at the outset will repay handsomely, allowing the true potential of machine learning to be realized most efficiently.

Machine learning has been in scientific use for more than half a century as a term describing programmable pattern recognition. The concept is even older, having been expressed by pioneering mathematicians in the early 19th century. It has come into its own in the past two decades, with the advent of powerful computers, the Internet, and mass-scale digitization of information. In the domain of artificial intelligence, machine learning increasingly refers to computer-aided decision making based on statistical algorithms generating data-driven insights (see sidebar, "Machine learning: The principal approach to realizing the promise of artificial intelligence").

Among its most visible uses is in predictive modeling. This has wide and familiar business applications, from automated customer recommendations to credit-approval processes. Machine learning magnifies the power of predictive models through great computational force. To create a functioning statistical algorithm by means of a logistic

Machine learning: The principal approach to realizing the promise of artificial intelligence

Artificial intelligence is the science and engineering of automated problem solving. The object is to generate solutions by using computers to mimic the cognitive functions associated with deliberative thought, including perception, reasoning, and learning.

Machine learning is the most prevalent means by which the potential of artificial intelligence is being exploited. The term refers to the ability of computers to detect patterns in large data sets through the application of algorithms. In addition to uncovering potentially powerful insights in the data, computers can be programmed to train themselves to make data-driven predictions.

Predictive modeling, also called supervised learning, is a machine-learning approach that builds pattern-recognition models using sample data with known attributes and outcomes (labeled "training data"). Working from the known patterns, the model can predict outcomes for new observations. The form of data used to predict outcomes can be structured or unstructured, whether or not supervised learning is applied. However, unstructured data can be processed directly only through machine learning; when more traditional techniques such as regression are used, the data scientist must first aggregate unstructured data into structured data based on business rules or independent analyses and procedures.

Deep learning is the most advanced technique for predictive modeling. It connects software-based calculators to form a complex artificial "neural network," often 50 or more layers deep. The simplest predictive-modeling techniques are regression modeling and simple decision trees. More advanced techniques include random forests (a more complex and sensitive decision-tree model) and support vector machines (for sophisticated data classification).

regression, for example, missing variables must be replaced by assumed numeric values (a process called imputation). Machine-learning algorithms are often constructed to interpret "missing" as a possible value and then proceed to develop the best prediction for cases where the value is missing. Machine learning is able to manage vast amounts of data and detect many more complex patterns within them, often attaining superior predictive power.

In credit scoring, for example, customers with a long history of maintaining loans without delinquency are generally determined to be of low risk. But what if the mortgages these customers have been maintaining were for years supported by substantial tax benefits that are set to expire? A spike in defaults may be in the offing, unaccounted for in the statistical risk model of the lending institution. With access to the right data and guidance by subject-matter experts, predictive machine-learning models could find the hidden patterns in the data and correct for such spikes.

The persistence of bias

In automated business processes, machine-learning algorithms make decisions faster than human decision makers and at a fraction of the cost. Machine learning also promises to improve decision quality, due to the purported absence of human biases. Human decision makers might, for example, be prone to giving extra weight to their personal experiences. This is a form of bias known

as anchoring, one of many that can affect business decisions. Availability bias is another. This is a mental shortcut (heuristic) by which people make familiar assumptions when faced with decisions. The assumptions will have served adequately in the past but could be unmerited in new situations. Confirmation bias is the tendency to select evidence that supports preconceived beliefs, while loss-aversion bias imposes undue conservatism on decision-making processes.

Machine learning is being used in many decisions with business implications, such as loan approvals in banking, and with personal implications, such as diagnostic decisions in hospital emergency rooms. The benefits of removing harmful biases from such decisions are obvious and highly desirable, whether they come in financial, medical, or some other form.

Some machine learning is designed to emulate the mechanics of the human brain, such as deep learning, with its artificial neural networks. If biases affect human intelligence, then what about artificial intelligence? Are the machines biased? The answer, of course, is yes, for some basic reasons. First, machine-learning algorithms are prone to incorporating the biases of their human creators. Algorithms can formalize biased parameters created by sales forces or loan officers, for example. Where machine learning predicts behavioral outcomes, the necessary reliance on historical criteria will reinforce past biases, including stability bias. This is the tendency to discount the possibility of significant change—for example, through substitution effects created by innovation. The severity of this bias can be magnified by machine-learning algorithms that must assume things will more or less continue as before in order to operate. Another basic bias-generating factor is incomplete data. Every machine-learning algorithm operates wholly within the world defined by the data that were used to calibrate it. Limitations in the data set will bias outcomes, sometimes severely.

Predicting behavior: 'Winner takes all'

Machine learning can perpetuate and even amplify behavioral biases. By design, a social-media site filtering news based on user preferences reinforces natural confirmation bias in readers. The site may even be systematically preventing perspectives from being challenged with contradictory evidence. The self-fulfilling prophecy is a related by-product of algorithms. Financially sound companies can run afoul of banks' scoring algorithms and find themselves without access to working capital. If they are unable to sway credit officers with factual logic, a liquidity crunch could wipe out an entire class of businesses. These examples reveal a certain "winner takes all" outcome that affects those machine-learning algorithms designed to replicate human decision making.

Data limitations

Machine learning can reveal valuable insights in complex data sets, but data anomalies and errors can lead algorithms astray. Just as a traumatic childhood accident can cause lasting behavioral distortion in adults, so can unrepresentative events cause machine-learning algorithms to go off course. Should a series of extraordinary weather events or fraudulent actions trigger spikes in default rates, for example, credit scorecards could brand a region as "high risk" despite the absence of a permanent structural cause. In such cases, inadequate algorithms will perpetuate bias unless corrective action is taken.

Companies seeking to overcome biases with statistical decision-making processes may find that the data scientists supervising their machine-learning algorithms are subject to these same biases. Stability biases, for example, may cause data scientists to prefer the same data that human decision makers have been using to predict outcomes. Cost and time pressures, meanwhile, could deter them from collecting other types of data that harbor the true drivers of the outcomes to be predicted.

The problem of stability bias

Stability bias—the tendency toward inertia in an uncertain environment—is actually a significant problem for machine-learning algorithms. Predictive models operate on patterns detected in historical data. If the same patterns cease to exist, then the model would be akin to an old railroad timetable—valuable for historians but not useful for traveling in the here and now. It is frustratingly difficult to shape machine-learning algorithms to recognize a pattern that is not present in the data, even one that human analysts know is likely to manifest at some point. To bridge the gap between available evidence and self-evident reality, synthetic data points can be created. Since machine-learning algorithms try to capture patterns at a very detailed level, however, every attribute of each synthetic data point would have to be crafted with utmost care.

In 2007, an economist with an inkling that credit-card defaults and home prices were linked would have been unable to build a predictive model showing this relationship, since it had not yet appeared in the data. The relationship was revealed, precipitously, only when the financial crisis hit and housing prices began to fall. If certain data limitations are permitted to govern modeling choices, seriously flawed algorithms can result. Models will be unable to recognize obviously real but unexpected changes. Some US mortgage models designed before the financial crisis could not mathematically accept negative changes in home prices. Until negative interest rates appeared in the real world, they were statistically unrecognized and no machine-learning algorithm in the world could have predicted their appearance.

Addressing bias in machine-learning algorithms

As described in a previous article in *McKinsey on Risk*,[1] companies can take measures to eliminate bias or protect against its damaging effects in human decision making. Similar countermeasures can protect against algorithmic bias. Three filters are of prime importance.

First, users of machine-learning algorithms need to understand an algorithm's shortcomings and refrain from asking questions whose answers will be invalidated by algorithmic bias. Using a machine-learning model is more like driving a car than riding an elevator. To get from point A to point B, users cannot simply push a button; they must first learn operating procedures, rules of the road, and safety practices.

Second, data scientists developing the algorithms must shape data samples in such a way that biases are minimized. This step is a vital and complex part of the process and worthy of much deeper consideration than can be provided in this short article. For the moment, let us remark that available historical data are often inadequate for this purpose, and fresh, unbiased data must be generated through a controlled experiment.

Finally, executives should know when to use and when not to use machine-learning algorithms. They must understand the true values involved in the trade-off: algorithms offer speed and convenience, while manually crafted models, such as decision trees or logistic regression—or for that matter even human decision making—are approaches that have more flexibility and transparency.

What's in your black box?

From a user's standpoint, machine-learning algorithms are black boxes. They offer quick and easy solutions to those who know little or nothing of their inner workings. They should be applied with discretion, but knowing enough to exercise discretion takes effort. Business users seeking to avoid harmful applications of algorithms are a little like consumers seeking to eat healthy food. Health-conscious consumers must study literature on nutrition and read labels in order to avoid excess calories, harmful additives, or dangerous allergens. Executives and practitioners will likewise have to study the algorithms at the core of their business and the problems they are designed to resolve.

They will then be able to understand monitoring reports on the algorithms, ask the right questions, and challenge assumptions.

In credit scoring, for example, built-in stability bias prevents machine-learning algorithms from accounting for certain rapid behavioral shifts in applicants. These can occur if applicants recognize the patterns that are being punished by models. Salespeople have been known to observe the decision patterns embedded in algorithms and then coach applicants by reverse-engineering the behaviors that will maximize the odds of approval.

A subject that frequently arises as a predictor of risk in this context is loan tenor. Riskier customers generally prefer longer loan tenors, in recognition of potential difficulties in repayment. Many low-risk customers, by contrast, aim to minimize interest expense by choosing shorter tenors. A machine-learning algorithm would jump on such a pattern, penalizing applications for longer tenors with a higher risk estimate. Soon salespeople would nudge risky applicants into the approval range of the credit score by advising them to choose the shortest possible tenor. Burdened by an exceptionally high monthly installment (due to the short tenor), many of these applicants will ultimately default, causing a spike in credit losses.

Astute observers can thus extract from the black box the variables with the greatest influence on an algorithm's predictions. Business users should recognize that in this case loan tenor was an influential predictor. They can either remove the variable from the algorithm or put in place a safeguard to prevent a behavioral shift. Should business users fail to recognize these shifts, banks might be able to identify them indirectly, by monitoring the distribution of monthly applications by loan tenor. The challenge here is to establish whether a marked shift is due to a deliberate change in behavior by applicants or to other factors, such as changes in economic conditions or a bank's promotional strategy. In one way or the other, sound business judgment therefore is indispensable.

Squeezing bias out of the development sample

Tests can ensure that unwanted biases of past human decision makers, such as gender biases, for example, have not been inadvertently baked into machine-learning algorithms. Here a challenge lies in adjusting the data such that the biases disappear.

One of the most dangerous myths about machine learning is that it needs no ongoing human intervention. Business users would do better to view the application of machine-learning algorithms like the creation and tending of a garden. Much human oversight is needed. Experts with deep machine-learning knowledge and good business judgment are like experienced gardeners, carefully nurturing the plants to encourage their organic growth. The data scientist knows that in machine learning the answers can be useful only if we ask the right questions.

In countering harmful biases, data scientists seek to strengthen machine-learning algorithms where it most matters. Training a machine-learning algorithm is a bit like building muscle mass. Fitness trainers take great pains in teaching their clients the proper form of each exercise so that only targeted muscles are worked. If the hips are engaged in a motion designed to build up biceps, for example, the effectiveness of the exercise will be much reduced. By using stratified sampling and optimized observation weights, data scientists ensure that the algorithm is most powerful for those decisions in which the business impact of a prediction error is the greatest. This cannot be done automatically, even by advanced machine-learning algorithms such as boosting (an algorithm designed to reduce algorithmic bias). Advanced algorithms can correct for a statistically defined concept of error, but they cannot distinguish errors with high business impact from those of negligible importance. Another example of the many statistical techniques data scientists can deploy to protect algorithms

from biases is the careful analysis of missing values. By determining whether the values are missing systematically, data scientists are introducing "hindsight bias." This use of bias to fight bias allows the algorithm to peek beyond its data-determined limitations to the correct answer. The data scientists can then decide whether and how to address the missing values or whether the sample structure needs to be adjusted.

Deciding when to use machine-learning algorithms

An organization considering using an algorithm on a business problem should be making an explicit choice based on the cost-benefit trade-off. A machine-learning algorithm will be fast and convenient, but more familiar, traditional decision-making processes will be easier to build for a particular purpose and will also be more transparent. Traditional approaches include human decision making or handcrafted models such as decision trees or logistic-regression models—the analytic workhorses used for decades in business and the public sector to assign probabilities to outcomes. The best data scientists can even use machine-learning algorithms to enhance the power of handcrafted models. They have been able to build advanced logistic-regression models with predictive power approaching that of a machine-learning algorithm.

Three questions can be considered when deciding to use machine-learning algorithms:

- *How soon do we need the solution?* The time factor is often of prime importance in solving business problems. The optimal statistical model may be obsolete by the time it is completed. When the business environment is changing fast, a machine-learning algorithm developed overnight could far outperform a superior traditional model that is months in the making. For this reason, machine-learning algorithms are preferred for combating fraud. Defrauders typically act quickly to circumvent the latest detection mechanisms they encounter. To defeat fraud, organizations need to deploy algorithms that adjust instantaneously, the moment the defrauders change their tactics.

- *What insights do we have?* The superiority of the handcrafted model depends on the business insights embedded in it by the data scientist. If an organization possesses no insights, then the problem solving will have to be guided by the data. At this point, a machine-learning algorithm might be preferred for its speed and convenience. However, rather than blindly trusting an algorithm, an organization in this situation could decide that it is better to bring in a consultant to help develop value-adding business insights.

- *Which problems are worth solving?* One of the promises of machine learning is that it can address problems that were once unrecognized or thought to be too costly to solve with a handcrafted model. Decision making on these problems has been heretofore random or unconscious. When reconsidering such problems, organizations should identify those with significant bottom-line business impact and then assign their best data scientists to work on them.

In addition to these considerations, companies implementing large-scale machine-learning programs should make appropriate organizational and cultural changes to support them. Everyone within the scope of the programs should understand and trust the machine-learning models—only then will maximum impact be achieved.

Implementation: Standards, validation, knowledge

How would a business go about implementing these recommendations? The practical application and debiasing of machine-learning algorithms should be governed by a conscious and eventually systematic process throughout the organization. While not as stringent and formal, the approach is related to

mature model development and validation processes by which large institutions are gaining strategic control of model proliferation and risk. Three building blocks are critically important for implementation:

- *Business-based standards for machine-learning approvals.* A template should be developed for model documentation, standardizing the process for the intake of modeling requests. It should include the business context and prompt requesters with specific questions on business impact, data, and cost-benefit trade-offs. The process should require active user participation in the drive to find the most suitable solution to the business problem (note that passive checklists or guidelines, by comparison, tend to be ignored). The model's key parameters should be defined, including a standard set of analyses to be run on the raw data inputs, the processed sample, and the modeling outputs. The model should be challenged in a discussion with business users.

- *Professional validation of machine-learning algorithms.* An explicit process is needed for validating and approving machine-learning algorithms. Depending on the industry and business context—especially the economic implication of errors—it may not have to be as stringent as the formal validation of banks' risk models by internal validation teams and regulators. However, the process should establish validation standards and an ongoing monitoring program for the new model. The standards should account for the characteristics of machine-learning models, such as automatic updates of the algorithm whenever fresh data are captured. This is an area where most banks still need to develop appropriate validation and monitoring standards. If algorithms are updated weekly, for example, validation routines must be completed in hours and days rather than weeks and months. Yet it is also extremely important to put in place controls that alert users to potential sudden or creeping bias in fresh data.

- *A culture for continuous knowledge development.* Institutions should invest in developing and disseminating knowledge on data science and business applications. Machine-learning applications should be continuously monitored for new insights and best practices, in order to create a culture of knowledge enhancement and to keep people informed of both the difficulties and successes that come with using such applications.

Creating a conscious, standards-based system for developing machine-learning algorithms will involve leaders in many judgment-based decisions. For this reason, debiasing techniques should be deployed to maximize outcomes. An effective technique in this context is a "premortem" exercise designed to pinpoint the limitations of a proposed model and help executives judge the business risks involved in a new algorithm.

Sometimes lost in the hype surrounding machine learning is the fact that artificial intelligence is as prone to bias as the real thing it emulates. The good news is that biases can be understood and managed—if we are honest about them. We cannot afford to believe in the myth of machine-perfected intelligence. Very real limitations to machine learning must be constantly addressed by humans. For businesses, this means the creation of incremental, insights-based value with the aid of well-monitored machines. That is a realistic algorithm for achieving machine-learning impact.

[1] Tobias Baer, Sven Heiligtag, and Hamid Samandari, "The business logic in debiasing," May 2017, McKinsey.com.

Tobias Baer is a partner in McKinsey's Taipei office, and **Vishnu Kamalnath** is an analytics specialist in the North American Knowledge Center in Waltham, Massachusetts.

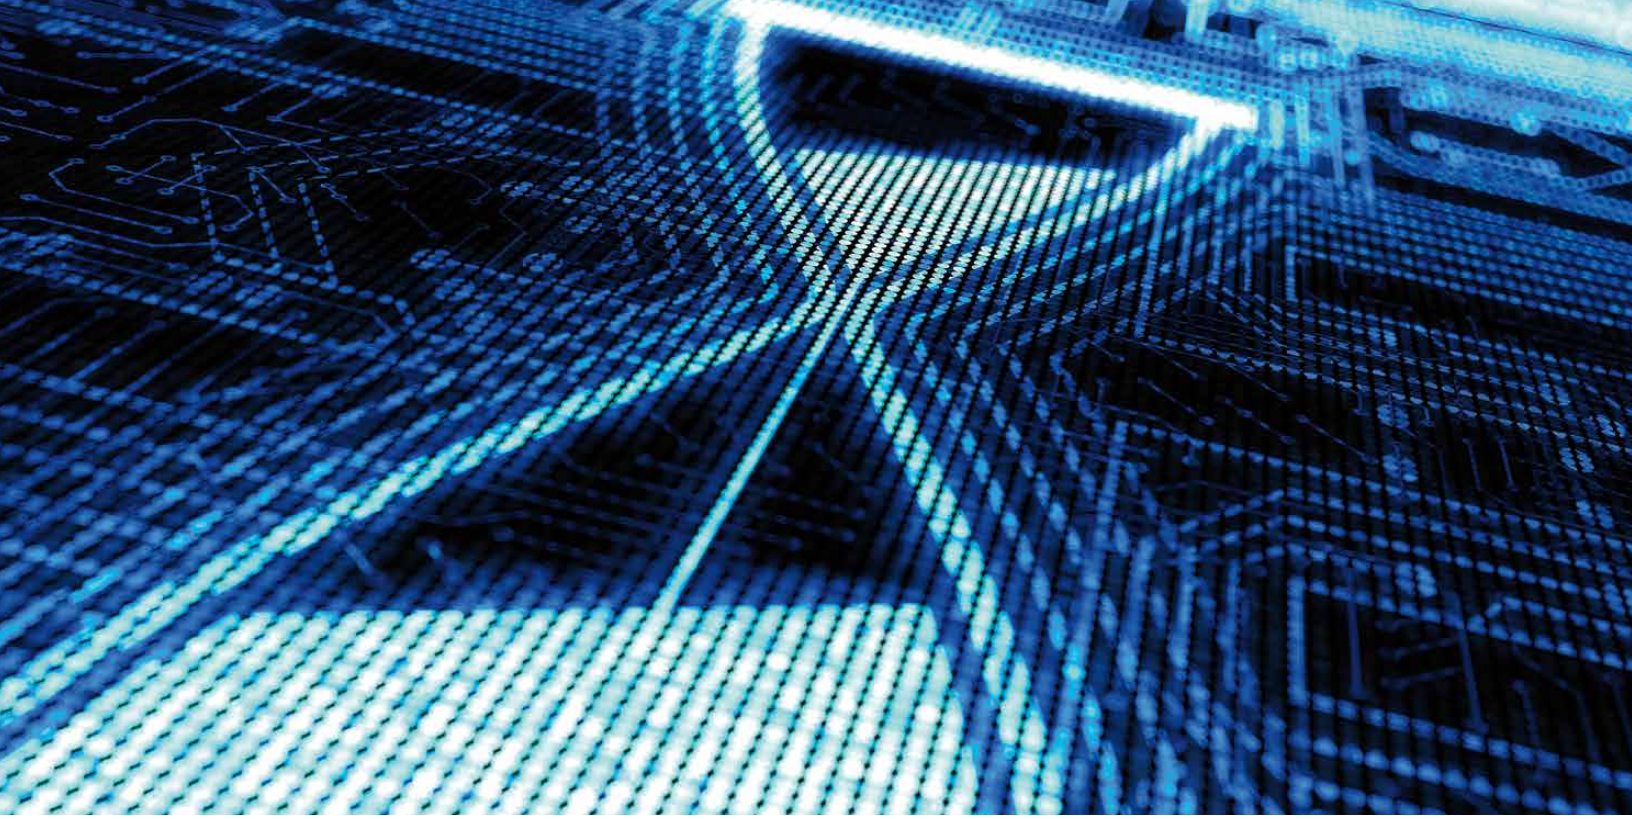

Tackling GDPR compliance before time runs out

Data protection has always been important. Now it's becoming urgent. Here's a primer on how companies can adapt to the new rules.

Daniel Mikkelsen, Kayvaun Rowshankish, Henning Soller, and Kalin Stamenov

Europe is on the brink of a sea change in its data-protection laws. In fact, when the General Data Protection Regulation (GDPR) takes effect on May 25, 2018, the effects will reverberate far beyond the continent itself. The GDPR goes further than harmonizing national data-protection laws across the European Union and simplifying compliance; it also expands the reach of EU data-protection regulation and introduces important new requirements. It seeks to ensure that personal data are protected against misuse and theft and to give EU residents control over how data relating to them are being used. Any entity that is established in the European Union or that processes the personal data of EU residents in order to offer them goods or services or to monitor their behavior—whether as customers, employees, or business partners—will be affected. Any failure to comply with the regulation could incur severe reputational damage as well as financial penalties of up to 4 percent of annual worldwide revenues (see sidebar, "The GDPR: Key facts," for a synopsis of the new rules).

After an initial wait-and-see approach, many companies in Europe and beyond—including those in Asia, the Middle East, and the United States—are starting to set up sizable compliance programs. Yet our recent surveys of major companies revealed that a third of the executives in the sample felt their organizations still had a long way to go on the road to compliance.[1] As the GDPR is based on principles rather than rules, the onus is on individual companies to determine implementation in their particular context (exhibit). This process

Exhibit

The General Data Protection Regulation sets out guiding principles for data protection.

Principle	Explanation
Lawfulness	Data should be processed only when there is a **lawful basis** for such processing (eg, consent, contract, legal obligation)
Fairness	The organization processing the data should provide data subjects with **sufficient information** about the processing and the **means to exercise their rights**
Transparency	The **information provided** to data subjects should be in a **concise and easy-to-understand format** (eg, the purpose of consent should not be buried in a lengthy document of terms and conditions)
Purpose limitation	Personal data may be collected only for a **specific, explicit, and legitimate purpose** and should not be further processed
Data minimization	The processing of personal data should be **adequate, relevant, and limited to what is necessary** in relation to the purposes for which those data are used
Accuracy	Data should be **accurate and kept up to date**
Storage limitation	**Data should not be held any longer than necessary** in a format that **permits personal identification**
Security	Data should be processed in a manner that **ensures security and protection against unlawful processing, accidental loss, damage, and destruction**
Accountability	The **data controller** is responsible for **demonstrating compliance**

Source: Regulation (EU) 2016/679 of the Council of the European Union, European Commission, and European Parliament

is fraught with uncertainty, and many companies are struggling to understand how they can best interpret, measure, and monitor compliance. Below we examine some of the main stumbling blocks and identify the steps that successful companies are taking to overcome them.

Why businesses are struggling with GDPR compliance

From our survey and conversations with executives, we have identified a number of ways that compliance efforts are falling short:

- *Underestimating the scope of the regulation.* Some of the executives who responded to our survey were not fully aware of the breadth of the GDPR, regarding it as merely an enhancement to existing regulations. Conversely, others felt that complying with the new provisions—especially the business and IT implementation of data-subject rights—would be onerous for their organization, and they were doubtful they would reach full compliance by May 2018. Indeed, only one of the 19 participants in our European survey believed his company would fully comply by the deadline.

The GDPR: Key facts

The scope of the General Data Protection Regulation (GDPR) is broad, covering any information that can be linked to an identifiable individual (such as search-engine entries, employee authentication, payment transactions, closed-circuit-television footage, and visitor logs) in any format (structured or unstructured) and in any medium (online, offline, or backup storage). The regulation introduces stringent consent requirements, data-subject rights, and obligations on organizations that gather, control, and process data. Its core requirements cover the following:

Documentation. Organizations should maintain a record of data-processing activities and be ready to present it to the regulator at any time.

Legal basis. All data processing should have a legal basis, such as the consent of the data subject or the need to fulfill a contract or legitimate business purpose.

Rights of data subjects. Organizations should implement rights such as the right to be forgotten (or, more accurately, to data erasure), the right to data portability, the right to object, the right to revoke consent, and the right to restrict processing.

Security. Organizations should protect data through means such as encryption or "pseudonymization" and have effective operational procedures and policies for handling them safely.

Third-party management. Vendors and suppliers, including outsourcing partners, should be required to protect personal data and should be monitored to ensure that they do so.

Privacy by design. Any organization planning a new technology, product, or service should consider data-protection requirements from the beginning of the development process.

Breach notification. Data breaches resulting in risk to individuals' rights and freedoms should be reported to the authorities within 72 hours, and subsequently to the data subjects as well in certain cases.

The new regulation will be enforced via national supervisory authorities within the European Union that are granted wide-ranging enforcement powers and sanctions, such as the power to ban data processing. The fines for failure to comply will be high, as much as 4 percent of annual worldwide revenues. The GDPR also allows individuals to seek civil actions (including class-action lawsuits) against organizations that violate their data-protection rights.

- *Uncertainty about how to interpret the requirements.* The GDPR sets out a number of principles that organizations should observe in processing personal data, but most companies have yet to decide how to put these principles into practice. For instance, under the principle of lawfulness, any organization processing personal data must have either the consent of the individuals concerned or some other lawful basis for that processing. Although the GDPR provides guidance on what might constitute a lawful basis—such as to carry out a contract, to comply with a legal obligation, or to serve the legitimate interest of the data controller or a third party—that guidance leaves a great deal of room for interpretation. In practice, we see organizations

taking very different views on issues such as the extent to which new consents are required from customers. In all these matters, companies will need to consult with lawyers. And lawfulness is not the only principle in the GDPR where there is uncertainty over interpretation. Take the accuracy principle, for example: it requires organizations to keep personal data up to date and take every reasonable step to rectify inaccuracies, but it is left to the organizations themselves to decide what steps they consider reasonable.

- *Slowness in identifying the additional security measures needed.* As the GDPR uses similar language to the current directive, many organizations are relying on their existing security measures, including protocols for particular customer segments, for compliance. However, as they build their records of processing activities, they will need to ensure that these measures are proportionate to the risks pertaining to different types of personal data. This calls for a structured approach to defining data risk and the measures necessary for mitigation—"pseudonymization," anonymization, encryption, deletion, and so on.

- *A struggle to build and maintain a comprehensive inventory of all their personal data–processing activities.* To satisfy this requirement, most of the banks we spoke with are relying initially on manual methods, typically using an internal survey to identify relevant data-processing activities within their organization. Such an approach may suffice for creating the inventory in the first place, but it is unlikely to be adequate to the task of keeping the inventory current and readily available to the regulator on demand. Sustainable processes and tools for maintaining detailed records have proved elusive so far for many organizations.

- *Lack of capabilities to fulfill their obligations.* Many companies are struggling to identify and develop the capabilities they will need to execute data subjects' rights in a timely manner. Consider, for example, the right to data portability. If a wealth-management firm receives a request from a customer to hand over all of her personal data to a different institution, what capabilities will it need to compile these data and transmit them to the new bank? Under the GDPR, the data covered by the portability requirement are not confined to the personal data an individual provides and the transactions they perform, but includes all observed data on that individual as well, such as the number and timing of their visits to the organization's website or calls to its call center. Building IT capabilities to fulfill these requirements may require banks to consolidate data from disparate systems, create new authentication methods, and introduce external application programming interfaces (APIs).

Steps to a successful GDPR effort

Drawing on our industry observations and regulatory experience, we have identified a number of actions that contribute to a successful GDPR effort and can help overcome some of the difficulties outlined above. Our advice is to check whether your institution is already taking these steps, and, if not, act now while there is still time.

- *Assign ownership of the program to a cross-functional task force.* A typical GDPR program does not have a natural owner in the organization; the challenge of ensuring compliance requires an approach that cuts across functions and businesses. All of the teams involved—legal, compliance, the business, IT, risk, and others—must commit to and share responsibility for a road map for change. Senior-leadership approval and buy-in is vital so that the program is securely anchored in a company's overall strategy.

- *Define the scope of your GDPR program and use a business lens to determine what should be ready for May 2018.* Most of the companies we surveyed believed they would not be fully compliant by the implementation date, so it is important to identify which aspects of the regulation and which data assets are critical to compliance and make them a priority. This means not only understanding legal requirements but also defining what risks the business is willing to accept, and what value it seeks to extract from the program.

- *Develop an in-house interpretation of GDPR requirements* that identifies the big strategic questions they pose and seeks to address them early on. The approach should reflect the most likely scenario, take the industry view into account, and neither downplay nor exaggerate the impact of the regulation. Adopting a black-or-white legalistic approach may not be helpful, so it will be important to stay close to peers as well as regulators and see what practical steps they are taking to comply. As your program progresses, take regular pulse checks to keep it on track. Given the heavy IT requirements, your program validation should be performed well before the second quarter of 2018 to allow time for course correction, if needed.

- *Assess your GDPR readiness to uncover any gaps* and plan measures to fill them, whether that involves modifying marketing processes to secure customer consent, developing new in-house data-protection measures, or carrying out vendor evaluations. Bear in mind that adopting manual solutions to satisfy requirements such as ensuring data portability can lead to high ongoing running costs. Building an automated solution at the outset—such as APIs for data transfer—could simplify compliance and reduce costs in the long run if you believe there will be sufficient demand (for instance, for data portability) to justify the investment involved.

- *Begin building a "golden record" of every personal data–processing activity* in the organization to ensure compliance and traceability. This goes beyond documenting the system inventory and involves maintaining a full record of where all personal data come from, what is done with them, what the lawful grounds for processing are, and whom the data are shared with. Map business or functional activities that use personal data and get the owners of these activities to complete a detailed questionnaire about the data processing involved. In parallel, work with vendors and internal IT experts to build tools and processes to maintain the inventory in steady state. This can be done as part of your software-development life cycle and data-protection impact assessments. Some companies adopt special data tools to discover personal-data assets and provide compliance reporting, but these tools have yet to be proved at scale in the marketplace.

- *Define your organizational setup for data protection.* Designating a data-protection officer (DPO) is not enough. Companies also need to weigh the pros and cons of different organizational setups to arrive at a reporting structure that enables the DPO to operate independently; to interact effectively with the chief information-security officer, chief privacy officer, and heads of legal, compliance, and risk; and to report to the highest level of management. Having decided on the new structure, companies then need to determine the resources required to support it and fulfill their data-protection responsibilities more broadly.

- *Define the uncertainties in interpreting the GDPR requirements, and identify unacceptable risks to your business and IT.* Many aspects of GDPR will be gradually resolved through industry practices and codes of conduct, regulatory guidance, or the court system. Interpretations of what is legally acceptable may also change over time. Frequent interactions with legal and business partners on compliance, legal issues, cybersecurity, application development, third-party vendor management, operations, marketing, and so on will help companies build a shared understanding of what they need to do. Beyond pure compliance, IT and the business should work together to define where the program should go the extra mile to minimize reputational risk, maintain customer trust, and avoid last-minute IT scrambles. This may involve implementing more stringent consent requirements, prominently announcing opt-out possibilities, implementing tougher-than-necessary security measures, and setting a high bar for sending personal data to third parties.

- *Consider strategic value.* Half the chief information-security officers in our sample regarded GDPR as primarily a hindrance to their business. Undoubtedly, the regulation will impose a burden on organizations, and with a matter of months to go before implementation, companies are racing to limit any negative impact it may have. However, what many leaders miss are the benefits that can be realized through a GDPR program. A well-conceived program can help an organization to build customer trust, improve customer relationships, establish better data controls, and improve internal data handling and availability. One company is taking advantage of its GDPR program to reengineer its master data-management platform so that all parts of the organization have a complete picture of all personal data on any given customer. Other companies are using GDPR-inspired reforms as an opportunity to build greater flexibility into their data platforms so that they can not only comply with the new provisions but also respond more readily to future regulatory changes. Seen in this light, a GDPR program can be an opportunity to embark on a wider data transformation that will benefit the whole business.

The steps above will help any institution get on the right track to meet next year's implementation date. GDPR should not be taken lightly. Organizations that fail to comply could face high fines, civil actions, and reputational damage, while those that use their GDPR program to spur a broader data transformation may be able to capture additional business flexibility and value. These are compelling reasons to treat the new regulation as a high priority for the whole organization, not just the risk, legal, and compliance functions. And with the implementation date imminent, companies need to act fast.

[1] We surveyed 19 executives at McKinsey's European General Data Protection Regulation (GDPR) Roundtable in February 2017; most were chief information-security officers. In May 2017, we conducted an informal online poll of eight US executives who were leading GDPR efforts.

Daniel Mikkelsen is a senior partner in McKinsey's London office. **Kayvaun Rowshankish** is a partner in the New York office, where **Kalin Stamenov** is a consultant. **Henning Soller** is an associate partner in the Frankfurt office.

The authors wish to thank Malin Strandell-Jansson for her contributions to this article.

The new frontier in anti–money laundering

New analytical tools and surgical automation can help banks take the fight to fraudsters.

Stuart Breslow, Mikael Hagstroem, Daniel Mikkelsen, and Kate Robu

In recent years, three factors have heightened the risk banks face when combating financial crimes. First, the growth in volume of cross-border transactions and greater integration of the world's economies have made banks inherently more vulnerable. Second, regulators are continually revising rules as their focus expands from organized crime to terrorism. Finally, governments have expanded their use of economic sanctions, targeting individual countries and even specific entities as part of their foreign policies.

Banks have responded to these trends by investing heavily in people, manual controls ("checkers checking the checkers"), and systems addressing point-in-time needs. For example, in the United States, anti–money laundering (AML) compliance staff have increased up to tenfold at major banks over the past five years or so. Banks have typically used a piecemeal approach, adding staff to areas with the weakest controls. Often this has resulted in compliance programs built for individual countries, product lines, and customer segments—with all the duplication that suggests. Banks have also hired thousands of investigators to manually review high-risk transactions and accounts identified through inefficient, exception-based rules. For example, one big US bank expanded the ranks of its compliance team by one-third in recent years, including many people who work on "know your customer" (KYC) and AML compliance. Banks are also spending hundreds of millions of dollars to maintain the processes and systems they built in response to remediation needs.

As a result, second-line AML compliance programs now look more like operational utilities or, as one executive put it, "factories," and less like the independent oversight functions that banks first envisioned. These factories are expensive yet might be acceptable if the huge teams and manual processes were working well. But many are not. Most financial institutions continue to face challenges that erode the effectiveness and efficiency of their AML programs, including the following:

- Poor-quality data, nonstandard data structures, and fragmented sources make data aggregation by legal entities, subsidiaries, and vendors difficult. For example, many banks are still making tens of thousands of costly customer calls every month to refresh KYC documents, updating information that is incorrect or missing in their databases.

- Analytical approaches for customer risk scoring and transaction monitoring suffer from high rates of false positives, resulting in significant resources focused on investigating low-risk accounts and transactions. Adding new calibration tools and thresholds often leads to another spike in the number of false alerts.

- Inconsistent standards in processes such as customer identification, enhanced due diligence, and account monitoring and screening mean that businesses do not agree on what constitutes risk and violation of compliance requirements.

- Similarly, inconsistency in the reporting of suspicious activities and currency transactions means banks sometimes produce too many reports, and sometimes too few, exposing them to the twin dangers of regulatory sanctions and excessive cost.

- Fragmented systems and platforms limit the ability to automate transaction monitoring and due diligence. Instead, compliance teams spend the bulk of their time collecting data, and then on "stare and compare" sessions, instead of investigative work.

- Reliable quantitative metrics to assess risk across products, geographies, and processes are often not available.

- Ever-faster launches of new products and services, as well as instant fund transfers and mobile payments, add complexity to real-time detection and prevention. For example, "intelligent" ATMs allowing customers to anonymously deposit and transfer cash even when banks are closed certainly offer convenience but lack adequate KYC and AML safeguards.

Leading banks are trying to crack these problems by turning to new technologies. Machine learning, real-time data-aggregation platforms using fuzzy logic, rapid automation, and text and voice analytics offer a fundamentally new approach to managing compliance. Even better, they also offer an opportunity to simultaneously cut structural costs and improve the customer experience. As they take up these new tools, banks are shifting financial-crime compliance toward a more forward-looking and sustainable approach.

Traditional improvements in operations, governance, and management information systems will continue to be important elements in financial-crime-prevention programs. But technology and advanced analytics can raise these programs to much higher levels of effectiveness and efficiency. While there are many opportunities, our experience shows that banks should invest in three areas: efficient *data-aggregation* platforms, *advanced statistical modeling* (such as machine learning–based risk scoring and alert-generation engines), and *automation of processes* (such as investigator visualization tools).

Banks that invest strategically in these three areas, rather than tactically reacting to market and regulatory changes, can over time substantially reduce their risk exposure and capture other substantial benefits. For example, compliance-error rates measured through sample-based testing can be reduced from more than 30 percent to less than 5 percent. At the same time, false-positive alerts can be brought down from more than 90 percent to below 50 percent. These steps reduce the risk of regulatory fines and other penalties related to noncompliance, as well as help banks avoid potential reputational issues. The following discussions review ideas and techniques in the three areas and suggest ways banks can apply them.

Data aggregation

Banks in all markets struggle with the quality of data they keep on their customers, creating a significant obstacle to data aggregation. Long-time clients may have signed up when information standards weren't as rigorous and manual forms were prone to error. Most banks have established modern data-entry processes for new customers—yet these might be followed inconsistently across countries or even branches. The challenge can be especially daunting in some countries like the United States or the United Kingdom that have only partial nationwide identification systems.

Banks are turning to new tools to aggregate poor-quality data that can help them avoid hundreds of thousands of dollars in cost for manual data structuring and cleansing as well as hundreds of millions of dollars in investments required to build central "data lakes." For example, intelligent data platforms use machine learning or "fuzzy logic" (an approach to computing based on degrees of truth, rather than the more conventional binary true/false logic) on unstructured account and transaction data, to create a 360-degree view of suspected cases of money laundering. In practice, these new tools allow banks to automatically validate more customer identities, identify beneficial owners faster, and map how specific customers are connected to other individuals and legal entities, especially those earmarked as higher risks.

This can have significant implications on the volume of accounts and transactions that get escalated for manual reviews. For example, our analysis at one global institution showed that about half of the transactions flagged as "suspicious" would not have needed investigation if the bank had been able to connect the data held by its various divisions, some of which had identified and previously cleared the parties involved.

As another example, in a typical bank, data infrastructure and systems are not well positioned to quickly spot the connections among small cash deposits made by many different customers and wire transfers sent by those customers to the same recipient. The exhibit illustrates how a typical "smurfing" scheme works, in which cash deposits are broken down into amounts below the reporting threshold of $10,000. Analytics-driven data aggregation can help overcome these challenges by instantly connecting these individuals to the same geographic location, same behavioral pattern (for example, transaction types, frequency, and sequence), same destination account, and even block the wires from leaving the bank early in the process, before the laundered amount gets big.

Advanced analytics

Intelligent data aggregation is not the only opportunity to apply advanced analytics in the AML space. Consider customer risk scoring and the tools used to generate alerts on suspicious transactions. Current tools are often not statistical models at all, but rather a series of linear rules based on an institution's experience, a typology of known money-laundering events, and explicit regulatory requirements (such as reporting any wire transfers of more than $10,000). Regrettably for banks, up to

90 percent of the alerts generated by these rules can be false positives, and should be quickly discarded by investigators (but often are not). Though rarer, false negatives (or criminal activity that goes unnoticed) also pose a significant risk to banks. It is relatively easy for criminals to understand the linear rules currently applied by many banks and then design approaches to circumvent them (like smurfing, including the use of dormant intermediary accounts before the funds converge into the target account).

Exhibit

How 'smurfs' tie together small deposits and wires, evading anti–money laundering detection.

Source: McKinsey analysis

Statistical models based on machine learning and other forms of artificial intelligence can help banks raise their game. Such models review verified events to identify the often obscure combinations of predictive variables most likely to help minimize losses. Learning algorithms take advantage of the large pools of data and heightened computing power now available to detect patterns that might go unnoticed by data scientists. Systems using artificial intelligence can discern, for example, whether a series of transactions represents possible money laundering or a more innocent activity, such as a sudden wave of overseas expenses. In our experience, machine-learning algorithms can help reduce the number of false reports by 20 to 30 percent. As a result, investigators can spend more time on high-risk cases, and the manual work required can be reduced by as much 50 percent.

The impact of advanced statistical modeling is further increased when it's applied across a network of financial institutions. For example, one major European payments processor implemented machine-learning algorithms to follow the money across many banks and various entities, accounts, and locations. The approach allowed investigators to identify the paths used by "mule" accounts that are notoriously difficult to detect. Such accounts, spread across several financial institutions, bounce and "clean" the funds as they move from an illegal source into the formal financial system. Besides identifying the at-risk accounts in their network, investigators were able to develop powerful predictive variables to flag suspicious transactions and accounts newly entering the legal payments system. Forewarned is forearmed: banks that are on the alert for markers of increased money-laundering risk—such as use of Bitcoin services, prepaid cards, accounts opened by foreign students overseas—are able to stop transfers in real time.

The financial industry has been slow to adopt advanced tools such as machine learning, partly because the models are difficult to explain and validate to satisfy regulatory requirements. However, the techniques are becoming commonplace in other parts of the bank. Machine-learning algorithms are being used to offer better products and advice to customers, as well as to manage customer retention more effectively. Regulators are becoming more comfortable with validation approaches involving random forest and other such algorithms, which produce models that are relatively easy to explain and test for stability.

Banks can start with simple uses of analytics, like those involved in smart triage and microsegmentation of accounts and transactions to reduce false positives. For example, instead of making binary "file/do not file" decisions, some banks score each account and transaction that did not immediately require filing of suspicious-activity reports (SARs). They "hibernate" them until the cumulative view of triggers over time surpasses a predetermined threshold. Some institutions achieved a threefold improvement in SAR conversion rates through tighter segmentation of accounts and transactions based on behavioral and demographic characteristics, allowing them to distinguish between suspicious and nonsuspicious transactions the same way experienced investigators do.

Down the road, other tools might accelerate progress, given AML's heavy reliance on human judgment and expertise. Deep learning is an advanced form of machine learning that is already being used in image analysis and human language processing. It attempts to mimic human thought processes like those used by financial-crimes investigators and requires large amounts of data and fine-tuned models. Deep learning will likely start being deployed at scale in the next three to five years for banks to combat money laundering, fraud, and other financial crimes.

Automated processes

Automation and standardization of critical portions of the due diligence and investigation processes can

make expert staff more effective and significantly reduce their caseload. Robots can be used to automate certain activities, including the population of case files for investigators, the closing of level-one alerts, and the population of SAR forms. These measures can reduce the investigation time for alerts and allow for workforce optimization.

For example, a leading North American wealth manager used many techniques to move from a largely manual process for customer identification and due diligence to a reengineered and tech-enabled process. The solution included case-management work flow to guide due-diligence analysts faster and more effectively through the process; an integrated interface to bring all the data and third-party applications that analysts typically need into a single screen; rules-driven pipeline management to ensure priority-based resolution of cases; and so on. Under the new processes, staff were able to make decisions on low- to medium-risk customers almost instantly, and within 24 hours on most high-risk clients. The initiative also enhanced the customer experience by speeding up decisions and eliminating unnecessary follow-ups for missing information. All told, the firm was able to improve operational efficiencies in KYC by up to 50 percent.

The integrated interface is particularly important for speeding up the alert-investigation process and can be quickly acquired and deployed from a number of third-party vendors. This type of tool automatically gathers information through online searches, internal data, and third-party databases and highlights concerns such as relevant sanctions, negative media, and political exposure. This information is visualized into a clear, on-screen report that helps an investigator quickly assess the case and make a decision.

Filing of SARs with regulators is another area that presents high potential for automation. Natural-language-processing software converts data into text and can replace most of the work that investigators are traditionally putting into writing the reports that support their decisions on a case when it's filed. Integrated with the case-management work flows, nonperforming-loan applications can be really powerful tools that automatically generate the SARs as soon as the investigator pushes the "generate report" button—all it takes is a quick review and edit, followed by pushing the "file report" button.

How to get there

Our experience suggests that analytics and technology are important, but they alone will not provide a silver-bullet solution to all AML challenges. The key to impact is being able to deploy analytics and technology in a business-specific way and to embed them organically into business processes, which in turn often have to be fundamentally reshaped to take advantage of new tools.

With this context in mind, leading institutions are focusing on four key initiatives to both generate substantial value in the near term and course-correct their in-flight efforts to achieve a more sustainable target operating model:

1. Develop a truly end-to-end view of an optimized, tech-enabled KYC and AML process, from new standards for customer-data intake to customer identification to risk-based due diligence to monitoring. The design of this "north star" process should cover complexity-based triage, rules-based routing of files to investigators, standards of work, quality tollgates, and so on. Currently no single third party supports the entire process; hence the bar is very high for up-front system-architecture design and integration of internal and third-party point-to-point solutions.

2. Define a strategy for data quality and aggregation, including linking KYC and AML data closer together. Consider application of analytical tools such as fuzzy logic and machine learning to connect the dots in the known KYC/AML reference data—such as customer investments

and associated entities and individuals—and sanctions data. This will address a significant share of customer due-diligence escalations. It will also end many of the sanctions and transaction-monitoring alerts that result from gaps and problems with matching reference data, including known intracompany and intercompany transfers and customer identification data from another business unit.

3. Fold simple analytical approaches like micro-segmentation into current systems and model-validation processes. In this way, added rules on top of existing rule engines for customer risk scoring and alert generation can be consolidated. In parallel, banks can invest in longer-term solutions such as training neural networks through supervised learning, which will further reduce false positives and false negatives. These tools can be brought up to high performance and be ready to go once regulatory approval is completed.

4. Implement a set of metrics and practices to measure effectiveness of the KYC/AML processes and assess impact from operational and system improvements. Potential metrics could include the following:

   - Establish the expected volume and quantity of alerts. For example, set targets in 90-day intervals to reduce false alerts as new controls are launched.
   - Set rate of conversion of alerts to cases: for example, aim to reduce the SAR conversion rate by 1 or 2 percent every 90 days.
   - Reduce time per case: for example, set a target to reduce the investigation time by case type.
   - Set targets to reduce false positives and negatives, rather than focusing on the number of SARs filed or overall transaction volumes.

The industry is at a turning point. Not only are many banks reconsidering their approach to KYC and AML, but many regulatory-technology start-ups are launching products to support and sometimes supplant their efforts. Every new technology reaches a point when the hurdles fall away, and the benefits become too numerous to ignore any longer. As pioneering banks are finding out, automation and analytics for AML are at that point. ■

Stuart Breslow is a partner in McKinsey's New York office, **Mikael Hagstroem** is a partner in the Charlotte office, **Daniel Mikkelsen** is a senior partner in the London office, and **Kate Robu** is a partner in the Chicago office.

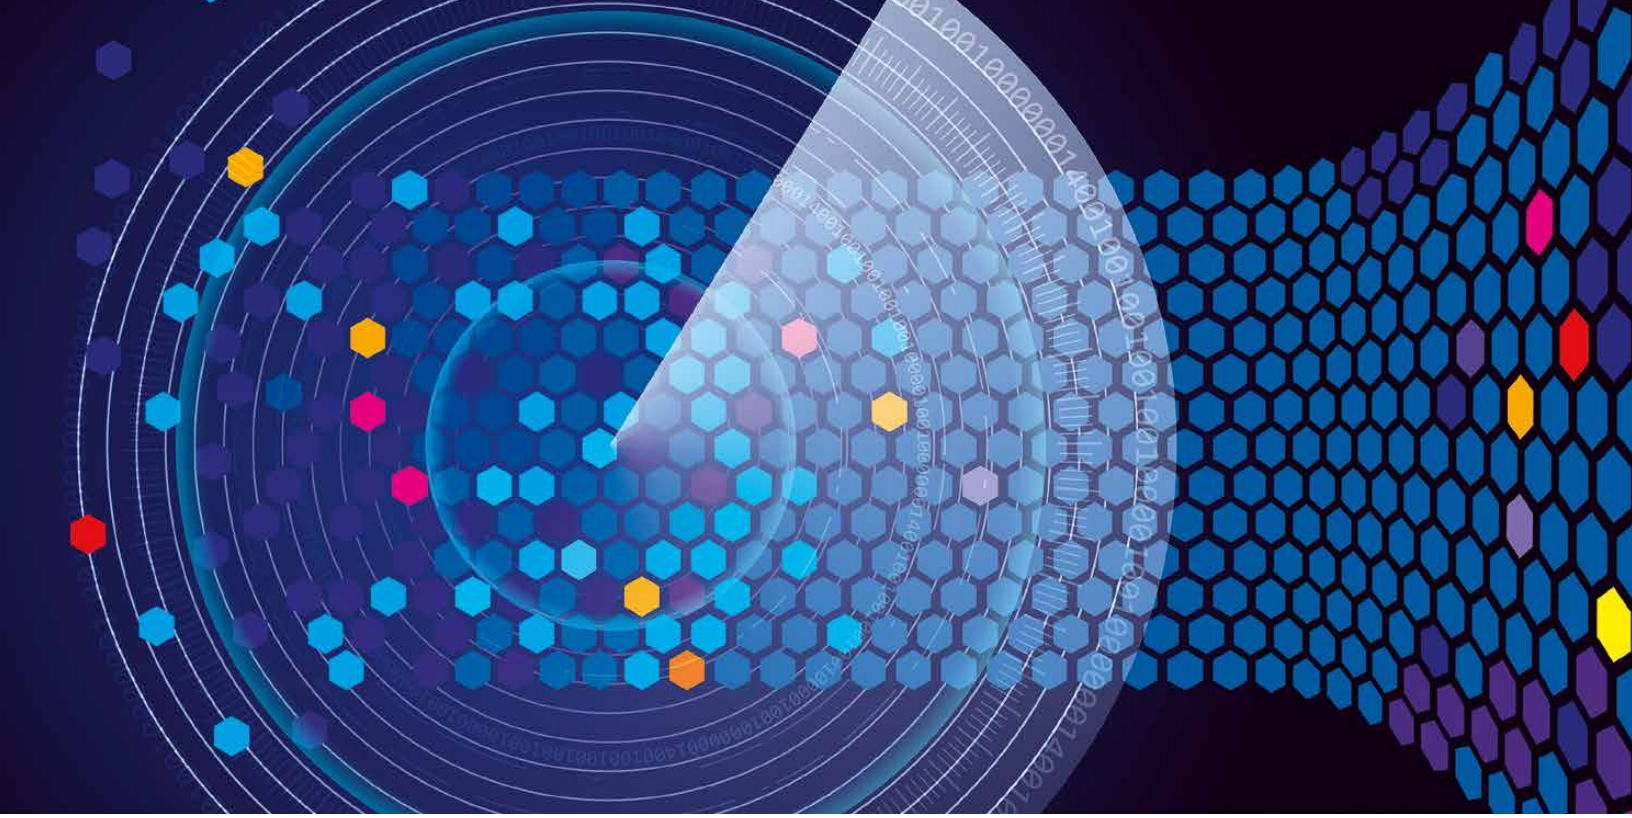

Using analytics to fight fraud

To fight financial fraud in a digital age, organizations need to apply advanced analytics and machine learning.

Philip Bruno, Jacomo Corbo, Carlo Giovine, and Chris Wigley

In decades past, check forgery was a typical form of banking fraud—an individual forger would copy the signature of an account holder on a check or withdrawal slip. Security depended on the alertness of bank employees, who were sometimes able to catch forgers in the act through signature comparisons. In the digital age, fraud has become more complex and organized—to say the least. A criminal gang is now able to target a student bank account, for example, which might suddenly receive a payment of £10,000. Within minutes, these funds can be cycled through many more accounts and then transferred out of the country. As investigators will eventually discover, the trail then runs cold. During the course of the crime, no alarms will have been sounded, no inquiries by affected customers and institutions will have been made to the original bank. By the time the fraud is recognized, no live trace of the perpetrators will remain.

Around the world, fraud is an ever-increasing risk for businesses in all sectors. In 2016, Kroll and the Economist Intelligence Unit reported that 75 percent of companies surveyed had been victims of fraud in the prior year, an increase of 14 percentage points from three years earlier.[1] An industry study within the financial sector found that 73 percent of finance professionals reported an attempted or actual payments fraud in 2015, a rise of 11 percentage points over 2014.[2]

In attempting to address fraud, financial institutions face a number of difficult challenges. Anti-fraud measures can be defeated by the sheer volume of transactions, all of which must be monitored in order to pursue comparatively few bad apples. In addition, fraud usually moves too fast for operative safeguards. Affected institutions might, for example, attempt to follow lightning-fast transfers with telephone calls to the sequence of banks involved. Further challenges are presented by institutional privacy practices and the built-in lack of information sharing among companies and across borders.

A digital problem requires a digital solution

In fighting fraud, institutions pass through decision cycles, in which potential fraud is detected, monitored, and acted upon. At the moment, these cycles do not turn fast enough to catch up. This is where data and analytics, including artificial intelligence and machine learning, come in. With a digital approach, financial institutions can accelerate the decision cycle. There are prerequisites, however. Appropriate processes have to be adopted and the necessary talent must be developed (or acquired). Institutions also need to create a culture in which data and analytics are fundamentally integrated into fraud-detection efforts.

Defining the role of analytics in addressing the challenges of financial fraud

The vast amount of data that flow through financial-services organizations provides a highly advantageous starting point for the application of analytics in fraud detection. Its effectiveness can be magnified by combining the proprietary data with other sources of information, since the most penetrating insights will tend to emerge where industries and data sets overlap. Likely sources are publicly available data sets, such as industry benchmarks and government data. Fraud markers—activity likely to indicate fraud—can also be added to enrich the data aggregation.

The object is to assemble the data that experts will need as they shape the appropriate analytical models to identify and prevent fraud most effectively. Success will produce financial savings but will also protect the institution's reputation and maintain public confidence. A recent example demonstrates how applying analytics to fraud detection can provide immediate and significant benefits.

A new model detects an unprecedented volume of invoice redirection

Imagine that your employees receive an email, sent seemingly from your account, requesting an update to the payment details of a key supplier. They might carry out this request—since it apparently comes from a trusted source—promptly and without question. In doing so they would also become unwitting accomplices to executive fraud. In this increasingly prevalent crime, imposters gain access to business email accounts and use them to convince unsuspecting employees to send funds to bogus accounts. Executive fraud nearly trebled in 2016 and has led to losses of more than $2.3 billion over the past few years.

Detecting this type of executive fraud, called invoice redirection, is especially challenging. To combat it, most banks use manual fraud-detection procedures or rules-based solutions. The effectiveness of these tools in spotting bogus accounts is limited. The accounts are designed to appear as genuine and automatically attract payments. How can banks detect the fraudulent look-alikes in time?

One large bank faced with this type of fraud created a tool that provides daily reports of suspicious transactions in time to intercept and evaluate them. The bank assembled the data needed—including one of the largest data sets in the country where it operates—to train an analytics model that could identify more than 80 percent of invoice redirection, in both value and occurrence.

To score every one of the millions of daily transactions for fraud risk, the bank built a supervised machine-learning model (Exhibit 1). To learn to detect fraud, the model needed a large data set. The number of potentially fraudulent transactions per day was very small, however, so the model would not have had access to the needed data for a long while. To get around this barrier, the data-analytics team decoupled the training process from the day-to-day operation and created a partially synthetic data set to train the model.

Exhibit 1 **A supervised machine-learning model helped monitor transactions for fraud.**

Supervised software training model flow pattern

Source: McKinsey analysis

The group worked with the data-engineering team to ensure computational performance, database best practices, and legal compliance. The curated data sets successfully trained the model to determine which transactions are safe and which are potentially fraudulent.

The model immediately and accurately categorizes most daily transactions as nonfraudulent. The remaining few thousand transactions are run through the machine-learning model, which uses analytics to combine the value and risk probability of each transaction. The model instantly ranks transactions by risk score, which is based on two different transaction patterns: one between the source and the destination account, and one that covers relationships established at the destination account. The risk score indicates which transactions are most suspicious and which are likely to be safe.

The model has significantly improved the bank's capability to detect high-value fraudulent transactions (Exhibit 2). It notifies the bank of 35 high-risk transactions on an average day, out of the several million processed. This allows the bank's fraud team to focus on the transactions that truly demand close investigation. These investigations identify new fraudulent cases and validate new safe relationships, results that are then fed back into the machine-learning model.

The predictive model identifies more than 85 percent of fraud cases in both value and incidents on the day the transaction is processed, allowing the bank to halt transactions before close of business and recover the funds. Within the first few weeks of live-scoring transactions, the model detected approximately £100,000 in fraudulent transactions. Other banks have expressed interest in the product, which is just the first step of applying analytics and modeling to the financial fraud-detection space.

Industry-wide applications for machine learning and analytics

Advanced-analytics techniques can be applied to organizations across the financial-services industry to anticipate and prevent fraud. As one auto insurer learned, moreover, the insights from machine learning often make significant contributions to the bottom line. A segment of this carrier's customer base had an accident rate in the first three months of its policies that was four times the company average. The insurer suspected that some customers were applying for insurance after being involved in an accident—an activity that its processes were unable to detect. To prevent these policies from being written in the first place, the insurer built a model using machine learning that could identify customers representing more than 80 percent of the risk. This enabled the insurer to improve its loss ratio and then reduce prices to attract a broader customer base. These measures increased profits by more than 10 percent.

Exhibit 2 **The model helped a bank improve its fraud-detection capability.**

A swindler uses a device to capture card information as client withdraws cash from an ATM. The information is then used fraudulently to withdraw cash or make purchases in a store or online. The model helped a bank spot the fraud.

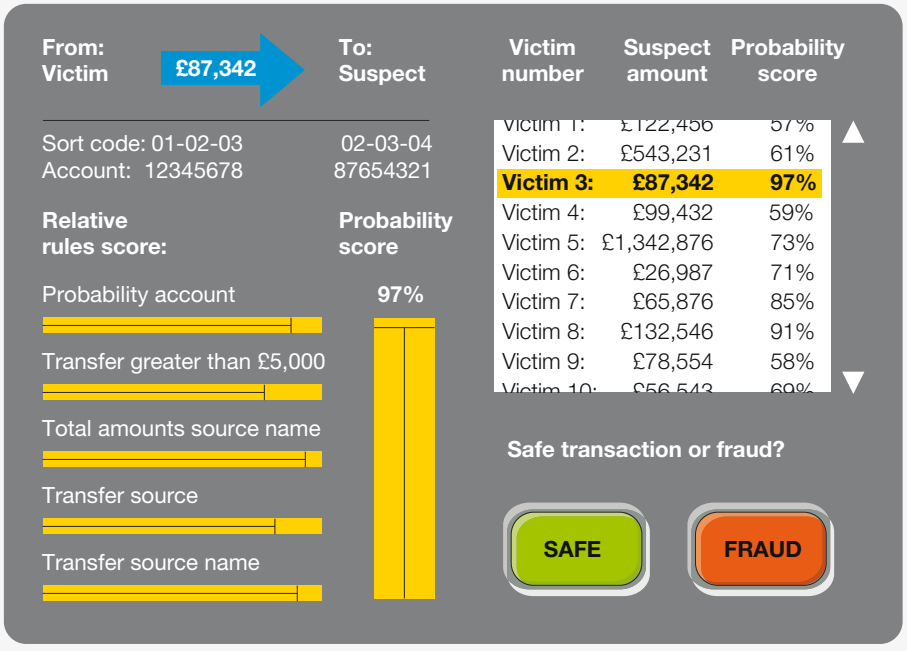

Source: McKinsey analysis

Working together to craft practical solutions

The foregoing examples illustrate some of the ways financial institutions have applied analytics to combat fraud. In the digital age, effective fraud detection and avoidance involves teamwork. Institutions that have been successful have brought together analysts and engineers to assemble the data needed to train the models that will identify fraud; those combined efforts are handsomely rewarded through a significant reduction in fraud losses and increased public confidence in financial institutions.

To benefit from the opportunities that data analytics presents to fight fraud, executives of financial institutions could implement a framework centered on four key areas: tools, processes, communication, and leadership.

- *Targeted tools and capabilities.* As the advanced-analytics solutions are developed, the institution has to provide training to ensure that people understand the fraud markers and achieve the desired results. A key element will be creating a culture of vigilance and data-driven decisions. In some areas, new talent will be needed.

- *Processes redesigned for speed and efficiency.* Determine how the organization will apply or alter its processes to improve fraud detection. This could involve changing the information that is reported or using new tools to obtain better information. An audit to identify data sources and measure data quality could be part of this phase.

- *Enterprise-wide communications.* Communicate the story of the fraud-detection effort and the new advanced-analytics capabilities. Explain how they will be deployed and their expected benefits. Ensure that the organization understands how the new capabilities will be used in day-to-day tasks. Use internal channels to share the story across the organization.

- *C-suite involvement.* Drive change from the top down. Executives should be involved in analytics initiatives and be vocal advocates for integrating data-driven decision making into all facets of the organization.

The talent component of this effort is also crucial. Institutions should determine whether to build their own internal data-science capability or work with an outside organization to close any gaps in analytics skills.

Analytics offers financial institutions the potential to identify fraud cases more quickly and frequently, sometimes even before the fraudulent act occurs. Since these organizations already collect a tremendous amount of data, the preconditions for successful detection models are already in place. The data sets do not have to be perfect to be useful, but most organizations will want to assess existing data and determine what other useful data might be collected.

To benefit from the fraud-fighting potential of data analytics, financial institutions must commit to developing the necessary skills and creating the appropriate culture. Given the potentially sizable rewards of reduced fraud losses and maintaining public trust, that commitment should be one that all organizations are willing to make. ■

[1] *Global fraud report 2015/16: Vulnerabilities on the rise*, Kroll, 2016, p. 7.

[2] *2016 AFP payments fraud and control survey: Report of survey results*, Association for Financial Professionals and J.P. Morgan, March 2016, p. 2.

Philip Bruno is a partner in McKinsey's New York office; **Jacomo Corbo** is a senior expert in the London office, where **Carlo Giovine** is an associate partner and **Chris Wigley** is a partner.

An earlier version of this article appeared on McKinsey.com and on the website of QuantumBlack, a McKinsey proprietary solution.

© fotomay/Getty Images

The true costs and impact of cybersecurity programs

Here's how business and technology leaders can ensure that important digital assets remain safe.

Jason Choi, James Kaplan, Chandru Krishnamurthy, and Harrison Lung

Companies are using all kinds of sophisticated technologies and techniques to protect critical business assets. But the most important factor in any cybersecurity program is trust. It undergirds all the decisions executives make about tools, talent, and processes. Based on our observations, however, trust is generally lacking in many organizations' cybersecurity initiatives—in part, because of competing agendas. Senior business leaders and the board may see cybersecurity as a priority only when an intrusion occurs, for instance, while the chief security officer and his team view security as an everyday priority, as even the most routine website transactions present potential holes to be exploited.

This lack of trust gives rise to common myths about cybersecurity—for example, about the types of threats that are most relevant, the amount of spending required to protect critical data, and even about which data sets are most at risk. Perceptions become facts, trust erodes further, and cybersecurity programs end up being less successful than they could be. If incidence of breaches has been light, for instance, business leaders may tighten the reins on the cybersecurity budget until the chief information officer (CIO) or other cybersecurity leaders prove the need for further investment in controls—perhaps opening themselves up to attack. Conversely, if threats have been documented frequently, business leaders may reflexively decide to overspend on new technologies

without understanding that there are other, nontechnical remedies to keep data and other corporate assets safe.

In our experience, when there is greater transparency about companies' cybersecurity programs and trust among the various stakeholders, companies reap significant benefits. Businesses can make better decisions about their security priorities and response plans, as well as the training and investments required to hold attackers at bay. In this article, we explore four common myths executives tend to believe about cybersecurity, and we suggest joint actions business and IT executives can take to create more transparency and understanding company-wide about the technologies and processes that are most effective for protecting critical business information.

Separating myths from facts

Based on our work with companies across industries and geographies, we've observed that business and cybersecurity leaders fall under the sway of four core myths when discussing or developing protection programs for corporate assets.

Myth 1: All assets in the organization must be protected the same way

Not all data are created with equal value. The customer data associated with a bank's credit-card program or a retailer's loyalty-card program are of greater value than the generic invoice numbers and policy documents that companies generate in-house. Companies don't have endless resources to protect all data at any cost, and yet most deploy one-size-fits-all cybersecurity strategies. When faced with a request from the IT organization for more funding for cybersecurity, C-suite leaders tend to approve it reflexively (particularly in the wake of a recent security breach) without a more detailed discussion of trade-offs—for instance, how much is too much to spend on protecting one set of critical data versus another? Or, if the company protects all external-facing systems, what kind of opportunities is it missing by not bringing suppliers into the fold (using appropriate policies and governance approaches)? Indeed, most business executives we've spoken with acknowledge a blind spot when it comes to understanding the return they are getting on their security investments and associated trade-offs.

In our experience, a strong cybersecurity strategy provides differentiated protection of the company's most important assets, utilizing a tiered collection of security measures. Business and cybersecurity leaders must work together to identify and protect the "crown jewels"—those corporate assets that generate the most value for a company. They can inventory and prioritize assets and then determine the strength of cybersecurity protection required at each level. By introducing more transparency into the process, the business value at risk and potential trade-offs to be made on cost would then be more obvious to all parties. A global mining company, for example, realized it was focusing a lot of resources on protecting production and exploration data, but it had failed to separate proprietary information from that which could be reconstructed from public sources. After recognizing the flaw, the company reallocated its resources accordingly.

Myth 2: The more we spend, the more secure we will be

According to our research, there is no direct correlation between spending on cybersecurity (as a proportion of total IT spending) and the success of a company's cybersecurity program. Some companies that spend quite a bit on cybersecurity are actually underperforming the rest of the market with respect to developing digital resilience[1] (Exhibit 1). In part, this is because those companies were not necessarily protecting the right assets. As we mentioned earlier, companies often default to a blanket approach (protecting all assets rather than the crown jewels).

Exhibit 1 **Companies' spending on cybersecurity does not necessarily correlate with level of protection.**

Note: Reflects responses from 45 companies in the Global 500 about their cybersecurity spending and capabilities.

[1]Companies' cybersecurity maturity is rated on a scale of 1 to 4, with 4 being the most mature (highest-level talent and capabilities).
[2]Spending is rated on a scale of 1 to 10; no companies allocated more than 10% of their budget to security.

Source: 2015 McKinsey Cyber Risk Maturity Survey

Throwing money at the problem may seem like a good idea in the short term—particularly when an intrusion occurs—but an ad hoc approach to funding likely will not be effective in the long term. Business and cybersecurity leaders instead must come to a shared understanding of costs and impact and develop a clear strategy for funding cybersecurity programs. The business and cybersecurity teams at a healthcare provider, for example, might agree that protecting patient data is the first priority

but that confidential financial data must also be secured so as not to compromise partner relationships and service negotiations. They could allocate resources accordingly. Without this shared understanding, business leaders may balk when a data breach occurs after they've funded significant changes in the security infrastructure. The lack of transparency and trust between the C-suite and the IT organization will only get worse.

Myth 3: External hackers are the only threat to corporate assets

It is true that threats from outside the company are a huge concern for cybersecurity teams, but there are significant threats inside corporate walls as well. The very people who are closest to the data or other corporate assets can often be a weak link in a company's cybersecurity program—particularly when they share passwords or files over unprotected networks, click on malicious hyperlinks sent from unknown email addresses, or otherwise act in ways that open up corporate networks to attack. Indeed, threats from inside the company account for about 43 percent of data breaches.[2]

Business and cybersecurity leaders must therefore collaborate on ways to improve internal risk culture. They must educate employees at all levels about the realities of cyberattacks and best practices for fending them off—for instance, holding town-hall meetings, mounting phishing campaigns, or staging war-game presentations to familiarize employees with potential threats and raise awareness. Many of these activities will need to be led by the CIO, the chief security officer, or other technology professionals charged with managing cybersecurity programs. But none will be fruitful if the company's business leaders are not fully engaged in a dialogue with the cybersecurity function and if companies don't build explicit mechanisms for ensuring that the dialogue continues over the long term. Business leaders at all levels must realize that they are the first line of defense against cyberthreats, and cybersecurity is never the sole responsibility of the IT department.

Myth 4: The more advanced our technology, the more secure we are

It is true that cybersecurity teams often use powerful, cutting-edge technologies to protect data and other corporate assets. But it is also true that many threats can be mitigated using less advanced methods. After all, most companies are not dealing with military-grade hackers. According to research, more than 70 percent of global cyberattacks come from financially motivated criminals who are using technically simple tactics, such as phishing emails.[3]

When companies invest in advanced technologies but do not understand how best to use them or cannot find properly skilled administrators to manage them, they end up creating significant inefficiencies within the cybersecurity team, thereby compromising the cybersecurity program overall.

Companies must, of course, explore the latest and greatest technologies, but it is also critical that companies establish and maintain good security protocols and practices to supplement emerging technologies—for instance, developing a robust patch-management program[4] and phasing out software for which vendors no longer provide security updates. This sort of foundation can help companies mitigate many of the biggest threats they may face. Consider the following example: a patch covering the vulnerabilities that could be exploited by the WannaCry cryptoworm was released March 14, 2017—some two months before the ransomware worked its way into more than 230,000 computers across more than 150 companies.

Building a culture of resilience

Rather than perpetuate myths, business and cybersecurity leaders should focus on bridging

the trust gaps that exist between them. We believe most companies can do that when technology and business leaders jointly train their attention on two main issues of control: how to manage trade-offs associated with cybersecurity, and how to discuss cybersecurity issues and protocols more effectively.

How do we manage trade-offs?

Technology professionals have a role to play in reeducating the C-suite about best practices in cybersecurity spending—specifically, illustrating for them why a tiered approach to cybersecurity may be more effective than blanket coverage for all. The budget cannot grow and shrink depending on whether the company recently suffered a system intrusion. Cybersecurity must be considered a permanent capital expenditure, and allocations should be prioritized based on a review of the entire portfolio of initiatives under way. Business and technology professionals must work together to manage the trade-offs associated with cybersecurity.

When discussing which initiatives to invest in and which to discontinue, business and cybersecurity professionals can use a risk-categorization model with four threat levels denoted, from minor to severe. The cybersecurity team can then engage the C-suite in discussions about the most important data assets associated with each part of the business value chain, the systems they reside in, the controls being applied, and the trade-offs associated with protecting higher-priority assets versus lower-priority ones.

At a broader level, technology professionals can help the C-suite create benchmarks for cross-company and multiyear expenditures on cybersecurity initiatives that can be reviewed regularly—for instance, cybersecurity spending as a percentage of overall IT expenditures. The CIO and his or her team could create a capital-expenditure index for security investments to help the C-suite justify cost per risk-adjusted losses or cost per percentage of infrastructure protected. Or, technology and business professionals could jointly develop a formula for quantifying the upside of making improvements to the cybersecurity program. In this way, they can make clear decisions about which tools to buy and add to the existing cybersecurity architecture, which systems to upgrade, and which to retire.

Regardless of the metrics used, it is important to have a comprehensive, formal approval process for planning and reviewing capital expenditures associated with cybersecurity. Priorities must be set from a business perspective rather than a systems perspective. CIOs and chief security officers must collaborate with the business to identify those assets with the potential to generate the greatest amount of value for the business and develop a cybersecurity road map accordingly. The road map would illustrate the distribution of crown jewels across the organization and the greatest surface areas of exposure. It would outline current controls and the sequence for launching new security initiatives, looking two to three years out. Of course, business and cybersecurity executives would need to revisit these plans quarterly or annually to ensure that they are still relevant given changes to the environment. The road map would also define roles and responsibilities, as well as mechanisms by which the C-suite and the leaders in the cybersecurity function could monitor progress made against the plan and revise it accordingly.

How do we talk about cybersecurity?

Weak communication accounts for much of the lack of trust between business leaders and members of the cybersecurity function. Our research indicates that in most companies, cybersecurity professionals are at least two layers from the CEO in the corporate hierarchy, with few opportunities for direct discussion about protection issues and priorities (Exhibit 2). What's more, in about half of the companies we studied, there was little to no formal

Exhibit 2 **Cybersecurity teams' access to the C-suite is limited.**

How many direct reports away is the senior-most cybersecurity executive from the CEO?,
% of survey respondents

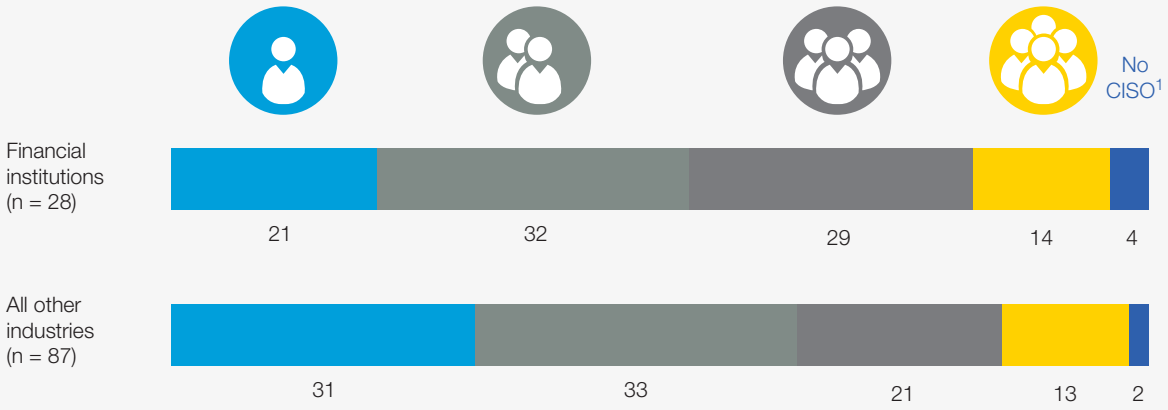

Note: Executives polled included chief information-security officers and other C-suite executives charged with making decisions about cybersecurity investments.

[1]Chief information-security officer.

Source: 2015 McKinsey Cyber Risk Maturity Survey

documentation shared by the cyber function with the C-suite about the status of their defense systems; many companies relied instead on occasional emails, memos, and notes (Exhibit 3).

Furthermore, when business and technology professionals do get in a room together, cybersecurity is usually discussed using highly technical language—for instance, "We already have measures to cover all CVE, however APT is something we need to watch out for. With our current SVM and SIEM infrastructure, there is no way we can defend these advanced attacks."[5] Jargon notwithstanding, the technology and business professionals in the room all understand how critical it is to build a robust cybersecurity program given the potential effects on the bottom line if corporate assets are compromised. But each side is typically only getting half the story.

Instead of reporting that "ten vulnerabilities were remediated," for example, technology professionals can use visual aids and outcomes-oriented language to help business leaders understand potential security threats and ways to address them. A status update might be better phrased in the following manner: "Our cybersecurity team has patched a security hole in our customer-relationship-management system that could have given hackers access to millions of packets of our retail customers' data, creating $100 million in financial damage." Cybersecurity professionals could also clearly delineate and communicate levels of systems access for intended and unintended users—a database administrator would have greater privileges than frontline employees, for instance.

Exhibit 3 **Many cybersecurity teams use informal means to communicate with business leaders.**

How do you summarize the status of defense systems to the chief information-security officer and business-level executives?, % of survey respondents

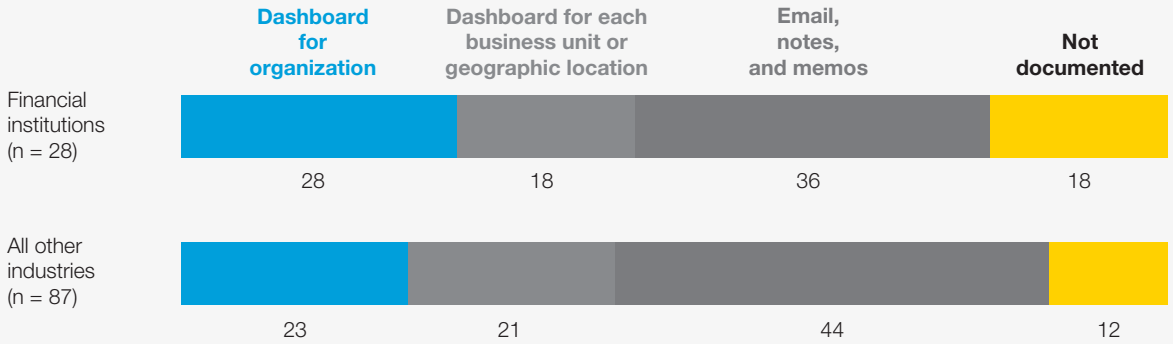

Note: Executives polled included chief information-security officers and other C-suite executives charged with making decisions about cybersecurity investments.

Source: 2015 McKinsey Cyber Risk Maturity Survey

Finding a common vocabulary is important not just for ensuring clear communication between the C-suite and the cybersecurity function but also for raising awareness about potential cyberthreats and risks among employees throughout the company. Members of the cybersecurity function should schedule frequent, regular check-ins with staff at all levels to educate them about relevant cybersecurity topics—how to recognize a phishing email, for example—and to showcase the company's security capabilities. The cybersecurity team at one technology firm conducts "road shows" to demonstrate which systems are being scanned and how they are being monitored. One online retailer, meanwhile, includes details about its cybersecurity efforts in existing financial reports—for instance, reporting on its development of an anti-malware scanner to protect the integrity of its recommendation engine, which helps drive advertising. It does this to illustrate that cybersecurity is part of the business process and can help drive revenue.

These discussions should take place regardless of whether the company is facing an imminent threat or not. The cybersecurity team at one company we observed shared with top leadership a simple breakdown of a typical security-event drill (Exhibit 4). The team wanted to give members of the board and the C-suite a step-by-step overview of what would happen in a typical attack—not just to prove the effectiveness of the company's security capabilities but also to familiarize individuals with potential threats so they might recognize them when they encounter deviations from the norm.

As we mentioned earlier, technology leaders may have to lead the charge in forging direct communications, creating cost transparency, and identifying business priorities. But the tasks suggested will require experience in C-suite-level communication, budgeting, and strategy planning—some of which may be beyond the core skill set of

Exhibit 4

Cybersecurity data theft has a pattern of event and response.

1 Insider takes sensitive data via flash drive

A disgruntled employee installs indexing malware in corporate systems and transfers files from servers to USB drive.

Visible hints

- Inquiry is made to senior executives about temp file being created and deleted.
- Slow laptops are reported to IT department and chief information officer.
- Help-desk ticket is sent to IT security lead.

Typical response

- Initially, the IT-security team does not realize that data are being threatened.
- Once the data are breached, the security team tries to determine best way to inform senior executives; the process is ad hoc, because protocols are not clear.

2 Insider gives or sells employee data to a cybercriminal

Cybercriminal uses old but valid credentials to access company servers and download employee records containing personally identifiable information (PII).

Visible hints

- Data-loss alerts are sent to the security lead in the IT organization.

Typical response

- Team focuses on the forensics of the alert but is not able to connect it to previous notifications.

3 Cybercriminal sells PII data to identity thieves on the black market

Identity thieves buy and use the employee data for fraudulent transactions.

Visible hints

- Based on individuals' and organization's complaints, the FBI detects the data breach and files a report with government affairs.

Typical response

- IT security reactively investigates employee data leak, trying to determine the scope of the breach.
- Team escalates event to privacy team.

4 Sensitive data are published on social media

Online bloggers publish video with references to the sensitive data stolen.

Visible hints

- An online video, found by employees, is sent to the head of communications.

Typical response

- The security team engages the communications group.

Source: 2015 McKinsey Cyber Risk Maturity Survey

those on the cybersecurity team. To come up to speed more quickly, cyber leaders may want to reach out to others with relevant expertise—for example, vendors and partners who can share best practices. In the spirit of agile development, cybersecurity teams may also want to take on these activities in "launch, review, adjust" mode. They could update threat and risk profiles in one- to six-month sprints, thereby ensuring they are responsive to the latest trends and technologies.

Make no mistake, the time to foster greater transparency about cybersecurity is now. The board must have trust in the C-suite and its ability to handle security breaches without dramatically affecting the company's value and brand. The C-suite needs to trust the chief information-security officer's claims that every penny spent on improving the security of IT infrastructure is worth it. The company needs to trust that vendors can properly protect shared data or ensure service stability if breaches occur. And, of course, customers need to trust that their personal data are being carefully safeguarded behind corporate walls.

The C-suite and the cybersecurity function can no longer talk past one another; security must be a shared responsibility across the business units. It must be embedded in various business processes, with the overarching goal of building a culture of resilience. The companies that take steps now to build greater trust between the business and the IT organization will find it easier to foster a resilient environment and withstand cyberthreats over the long term. ■

[1] Unless otherwise indicated, statistics relating to the composition and effectiveness of companies' cybersecurity programs are from the 2015 McKinsey Cyber Risk Maturity Survey.
[2] *Grand theft data*, Intel Security, 2015, mcafee.com.
[3] *2017 Data breach investigations report*, Verizon, 2017, verizonenterprise.com.
[4] Patch management is the structured process of acquiring, testing, and installing code changes to an administered computer system.
[5] CVE stands for common vulnerabilities and exposures, APT stands for advanced persistent threat, SVM stands for security and vulnerability management, and SIEM stands for security information and event management.

Jason Choi is a consultant in McKinsey's Hong Kong office, where **Harrison Lung** is an associate partner; **James Kaplan** is a partner in the New York office, and **Chandru Krishnamurthy** is a senior partner in the Atlanta office.

The authors wish to thank Suneet Pahwa and Chris Rezek for their contributions to this article.

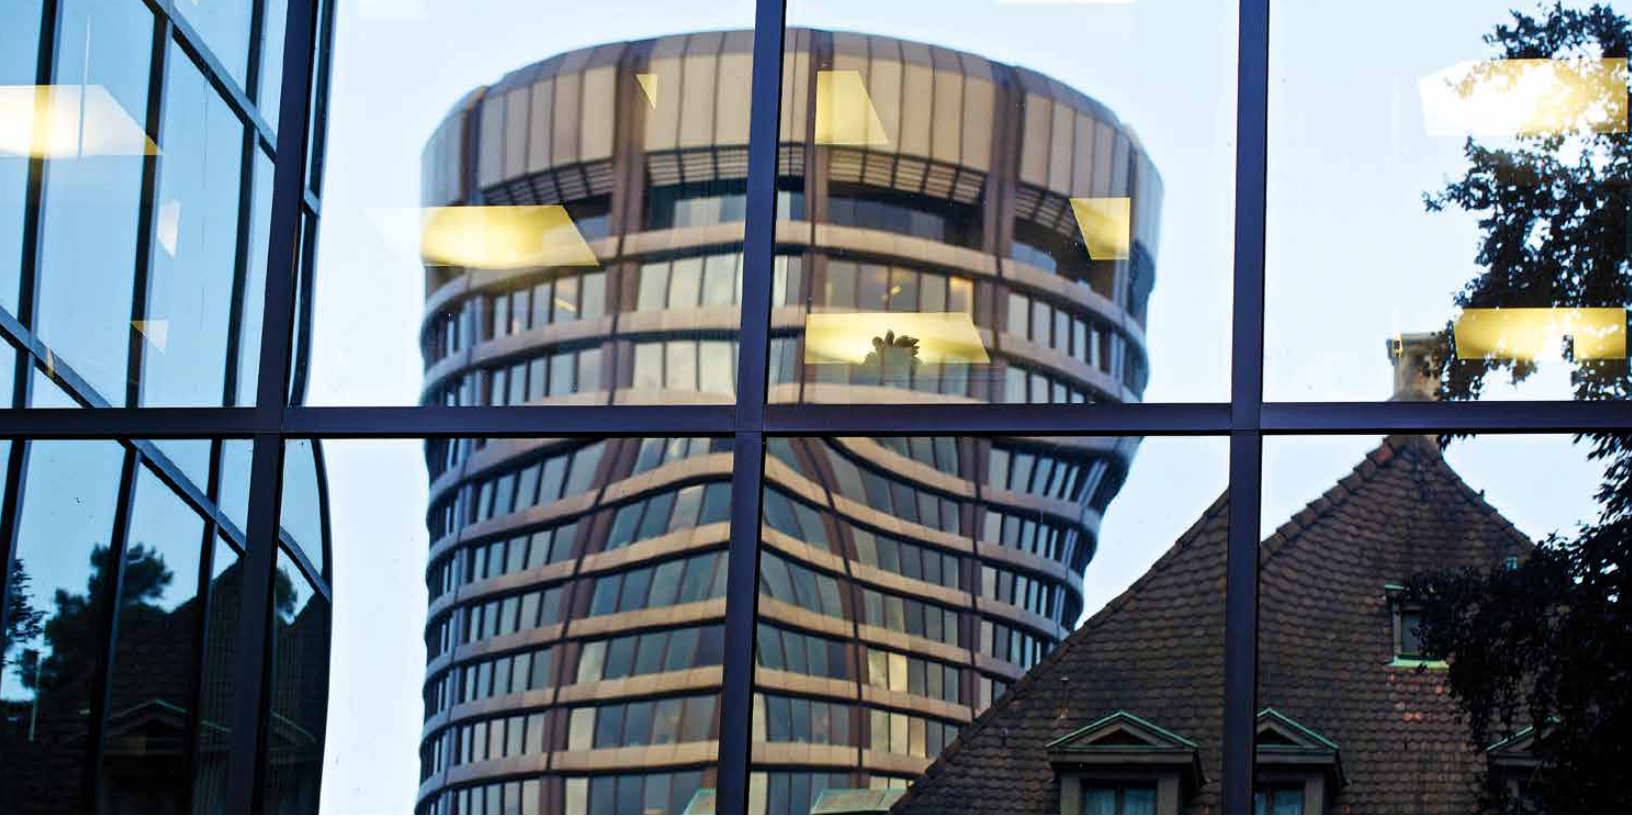

Headquarters, Bank for International Settlements, Basel, Switzerland
© Bloomberg/Getty Images

Bringing Basel IV into focus

How banks can mitigate €120 billion in capital requirements and avoid an ROE haircut.

Stefan Koch, Roland Schneider, Sebastian Schneider, Gerhard Schröck

The Basel III regulatory framework was developed to enhance the stability of the financial system by raising requirements on regulatory capital and liquidity. Basel III increased thresholds for capital quality and quantity, raising Tier 1 capital requirements, introducing buffers and leverage-ratio requirements, and adding the Common Equity Tier 1 requirement (CET1) (see sidebar, "Basel III, TLAC, MREL, and more").

Since Basel III was rolled out, the Basel Committee on Banking Supervision (BCBS) has been reviewing risk-measurement approaches internationally and among banks. One outcome of this review was the new standardized measurement approach (SMA) for operational risk, which was proposed in 2016. The committee also began a discussion on aggregated internal-rating model floors, concerned about the wide variation in the levels of risk-weighted assets (RWA) issuing from banks' internal models. The committee finalized standards for minimum capital requirements for market risk—the fundamental review of the trading book (FRTB)—in January 2016.[1] The committee plans to make technical revisions to this framework in 2017–18, however, and the country-level implementation is still under discussion.

Together, the changes are part of a Basel III amendment now more commonly referred to as Basel IV. The Group of Central Bank Governors and Heads of Supervision (GHOS) indicated that it does not intend to increase the total regulatory capital requirements in the industry as a whole. It did, however, acknowledge that the impact "may well be significant" for some banks. Recent

McKinsey research has gone further, suggesting that the impact of Basel IV will be significant throughout the banking industry.[2]

Banks will also have to deal with further regulatory adjustments and discussions that are indirectly affecting capital requirements under Pillar 1 and Pillar 2. These new mandates include risk-data aggregation and IT (BCBS 239), the revised standards for interest-rate risk in the banking book (IRRBB), and the introduction of IFRS 9 accounting standards. This new regulatory environment will require banks to run large-scale implementation programs and to ensure they have adequate resources to cover substantial one-time costs and provisioning needs. Moreover, additional capital requirements imposed by supervisors, such as during the EU Supervisory Review and Evaluation Process (SREP), will increase capital thresholds and loss-absorbency requirements (total loss-absorbing capacity, or TLAC, and minimum requirement for own funds and eligible liabilities, or MREL). The resulting higher funding costs from new issuance of eligible loss-absorbing liabilities could vary significantly from country to country if the European Union is unable to harmonize implementation throughout Europe. Coupled with the revised risk-measurement approaches, the new rules will no doubt entail expenses that affect banks' ability to build up organic capital.

According to our analysis, if banks do nothing to mitigate the cumulative impact, they will need about €120 billion in additional capital, while the banking sector's return on equity will be reduced by 0.6 percentage points. That is, current CET1 ratios of European banks would drop by 29 percent, according to our calculations, declining from a ratio of 13.4 percent now to 9.5 percent. The severest effect comes from internal-ratings-based (IRB) output floors, which would decrease CET1 ratios on average by about 1.3 percentage points. Other significant drivers are the new standardized measurement approach for operational risk (0.8 percentage points) and Basel III phase-in (about 0.5 percentage points). The average return on equity for European banks would drop to 7.4 percent from 8.0 percent, assuming that banks take no mitigating actions and keep Basel III capital requirements fully phased.[3]

A holistic program for improved capital management

The totality of the Basel IV adjustments has not yet emerged, and therefore the expected impact of the new regulations cannot yet be fully articulated. Some recent developments can be reported. BCBS now seems close to settling some Basel IV details, including an aggregated IRB floor of total RWA. Based on the standardized approaches, the floor is 70 to 75 percent, a level supported by most EU countries, including Germany (France and the Netherlands are exceptions). Likewise, the US Federal Reserve has lately signaled a willingness to accept floor levels below 80 percent of RWA. For another aspect of Basel IV, expectations have changed. The transitional period in which the rules are phased in may run through 2027, rather than 2025 as had been forecast earlier, due to the effects on mortgage portfolios.

The impact of Basel IV will vary by location, bank type, and business model, and no set of mitigating actions could uniformly address every situation. Each bank will have to work out an appropriate capital-management strategy to mitigate the impact of Basel IV based on its own position. Optimal responses will vary by bank: for example, banks with focused business models could face a significant IRB output-floor requirement.[4] In response, these banks will either have to adjust the composition of their business or move assets off the balance sheet. Banks with a more diversified portfolio will likely be able to respond with many smaller actions.

Each bank will likely need to adopt a package of changes big and small to improve capital management. Proposed *strategic shifts* in business

models will have to be tested for sustainability in the new regulatory environment. Most banks can make beneficial *business changes* that do not require a new strategic focus, including the application of methods to increase capital efficiency and profitability. Also of primary importance will be more rigorous *technical measures* to measure risk-weighted assets more accurately and improve regulatory capital—for example, by reducing capital deductions (exhibit).

Exhibit

To raise more capital without losing ROE under Basel IV, banks will need to introduce a holistic program to improve capital management.

Capital-management strategy	Typical impact, basis points[1] ROE potential	Typical impact, basis points[1] Capital ratio	Structural enablers to ensure strategy is sustainable
Strategic adjustments • Optimize balance sheet under all regulatory and business constraints • Location strategy for headquarters, booking models, and legal entities • Portfolio strategy, including exit of selected portfolios and regions, and identification of priority growth areas	50–100	To be determined	Embedded in strategic financial planning + Capital-conscious culture and action at front line + Effective performance management around capital usage + Ability to simulate and optimize balance sheets under normal and stressed conditions, and all regulatory requirements to control for business-model changes + Capital-steering model in place with consistent set of metrics and efficient capital-allocation process + Sustainable IT and process solutions for inaccuracies in regulatory risk-weighted-asset (RWA) reporting + Timely and accurate RWA reporting for management and front line
Business changes (no strategic change) • Tactical actions, including product and collateral optimization • Client review to promote profitable customers • Commercial actions: cross-selling, pricing, different product offerings	50–100	~100	
Technical measures • Address data quality and process issues (such as unrecognized collateral, ratings, or cash-flow-based effective maturity) • Reduce capital deductions and buffer requirements, including goodwill, intangibles, minorities, and G-SIB[2] and Pillar 2 buffers	50–100	~100	

[1] Not cumulative.
[2] Global systemically important bank.

Basel III, TLAC, MREL, and more

Basel III specifies capital target ratios of 7.0 percent for the core Tier 1 capital requirement, including a minimum of 4.5 percent of core Tier 1 capital and a required capital conservation buffer of 2.5 percent. A broader requirement for all Tier 1 capital is set at 8.5 percent; this includes the core Tier 1 minimum of 7.0 percent and an additional minimum for noncore Tier 1 capital of 1.5 percent. Additional hurdles are imposed by the Financial Stability Board's minimum standard for total loss-absorbing capacity (TLAC) and the minimum requirement for own funds and eligible liabilities (MREL). TLAC is intended to elevate capital and leverage ratios for the 30 institutions determined by the Basel Committee on Banking Supervision to be global systemically important banks (G-SIBs). A G-20 agreement applies TLAC to all G-SIBs, and the rule is scheduled to become codified in the European Union on January 1, 2019. MREL is an EU bail-in standard with a minimum requirement of 8 percent of liabilities; introduced in 2014, MREL is being revised to align with TLAC. As the regulatory bar rises, furthermore, the Swiss regulator has recently set capital requirements even higher than EU standards.

Strategic adjustments

The forthcoming requirements provide an opportunity for banks to rethink their portfolio of businesses, as well as individual business models. It will be important to review the business to identify activities that will become a drag on capital in a Basel IV environment. Given that internal models are restricted by applicable capital floors, banks whose business models are less sophisticated might suddenly become more competitive in terms of capital cost in certain product classes. The environmental change will increase competition and margin pressure for banks serving segments like specialized lending, where those using slotting or standardized models face significantly higher capital charges. Banks with less sophisticated models might suddenly become more competitive.

The top-down review of business activities should be based on a thorough understanding of how the new capital requirements affect each segment and product in both the current cycle and under stress scenarios. It will be crucially important to uncover the interdependencies and trade-offs among business segments and under different regulatory constraints.

First, banks will need to clarify the contributions made by each division to scarce regulatory resources (capital, funding, and liquidity) and their consumption. The complexity of this task should not be underestimated. Some of the metrics it requires are not consistently present and ready to use in typical IT systems. In addition to addressing the missing metrics, banks will need to focus on capital steering and allocation. This is because a number of diversification effects might arise once a bank is constrained by the IRB output floor. Banks could have to figure out how to allocate excess capital from operational-risk or market-risk standardized approaches to other business units or down to products, for example. This complicates capital steering for banks as they attempt to manage stressed and regulatory capital, capital buffers, and RWA-based and capital-requirements-based calculations. Only through a full understanding of the balance sheet at a group level will banks be able to quantify the total impact of division and product characteristics. Then banks can figure out how to adjust the balance sheet to optimize performance.[5]

Several leading banks have begun to use advanced modeling and optimization approaches to understand the evolving regulatory requirements. This process is typically interactive, in that strategic direction and business mix define the

parameters of the modeling, and the model can help quantify feasibility and implications of a chosen strategic direction. Once the review is complete, the businesses that remain in the portfolio must adjust their business models to the new capital realities. Some businesses may require only small adjustments, while others will be fundamentally changed. Two areas for strategic focus are banks' portfolio strategy and their legal-entity setup:

Updating the portfolio strategy. Banks should review their capital allocation to each client segment and region to ensure that capital is preferentially allocated to areas that generate higher returns (adjusting for risk, funding, and increased capital costs). Most banks have yet to institutionalize these capabilities. In trade approval, for example, value adjustments (xVA) are not often considered for changes in capital requirements, margin, or collateral requirements over the life cycle of a trade. Likewise, most banks need to adjust costs charged on the banking book for funding, liquidity, or capital to the new regulatory requirements. Client segments should be evaluated in growth and economics but also by capital requirements and capital efficiency—based on the current economic cycle and stress scenarios (to mitigate tail risk). With the evaluation for guidance, banks can then scale back business in segments and regions that do not add economic value—such as those that account for a big share of the bank's risk-weighted assets without returning the cost of capital.

Reviewing the legal-entity setup. Many banks are already questioning the number of legal entities in their structure in light of resolvability requirements. Reducing the number of subsidiaries typically leads to substantial capital and funding savings; it can also achieve some cost savings, better transparency, and improved governance. At the same time, an optimized legal-entity structure improves resolvability and may help in reducing MREL and TLAC requirements. While many local supervisors want subsidiaries to exercise better control over risk exposures and balance sheets, supervisors are also in favor of simpler legal structures. In deciding the appropriate legal-entity setup, banks will take into account client impact and strategy, regulatory and legal factors, the financial impact, and the effects on operations, as well as governance and implementation requirements.

Business changes

Certain adjustments to the business can increase capital efficiency, sometimes significantly. Some changes might slightly reduce revenues but also release capital demand such that overall profitability and capital efficiency increase.

Tactical moves include small adjustments to the current product offering or to the requirements or deals, making them more capital efficient for the bank. Collateral requirements could be revised, for example, so that more collateral is required or collateral allocation is improved. To optimize product offerings, banks could phase out under-utilized lines, adjust contract clauses (for example, committed versus uncommitted, or maturity clauses), or pursue product swaps—especially for limits such as overdrafts and revolvers.

A review of client relationships can result in improved profitability. After FRTB, banks will see significant shifts in profitability in both the trading and banking book. To respond, begin by identifying the essential relationships—high-revenue clients, those with international significance, or those that are sources of funding. For low-performing clients, relationships can be modified or ended. To refine and prioritize this list, leading banks have created thresholds by segment and then analyzed each relationship. The prospects for making relationships more profitable will then emerge. Actions can then be proposed to relationship managers and criteria developed for renegotiating or ending the relationship. The approach requires that the bank set top-down targets for risk-weighted assets,

capital efficiency, and cross-selling and work with relationship managers on client action plans. For trades, pricing and valuation will have to be revised to incorporate expected FRTB capital charges and value adjustments. By accounting for xVA, including for regulatory capital and marginal adjustments, banks would be able to capture the expected lifetime profitability of trades and understand whether to novate or break up trades. Where profitability targets are not met, an exit plan would take effect. The approach can then be expanded to the bank's entire client portfolio.

Commercial actions can be undertaken to ensure that banks continue to meet client needs while also increasing capital efficiency. Areas for consideration are products, collateral and guarantees, repricing, and cross-selling. Banks can help themselves in achieving these changes by aligning front-office incentives with the new realities.

- *Product offerings* can be adjusted, since over time some products by themselves become less attractive after capital costs. (Of course certain products will remain important for customer relationships.) For example, some institutions might need to stop offering certain mortgages and real-estate products that do not meet calculated tolerances to loan-to-value exposures or risk weights.

- Banks should review their policies on financial *collateral and guarantees*. While IRB models accept nonfinancial and physical collateral for credit-risk mitigation, the standardized measurement approach does not. Banks using the IRB approach might want to prefer financial collateral or guarantees, eligible under both approaches, once the floor framework based on the applicable SMA risk weights is introduced. Guarantees of highly rated counterparts will become more important compared with physical collaterals due to deviation of economic and regulatory credit-risk mitigation.

- *Repricing and cost-management actions* will be needed to support future profitability. Banks should assess the expected impact of the new regulatory requirements, especially in the product areas that are affected most, such as mortgages and commercial real-estate exposures. There will likely be opportunities to amend prices or reduce operating costs to make up for increased capital costs.

- Banks should also look for opportunities to increase their *cross-selling* of fee-based products that do not create any additional capital charge.

Finally, banks will want to model the impact of tactical and strategic changes on profitability, and optimize the balance sheet accordingly.

Technical measures

Industry evidence indicates that all banks can improve the accuracy of their RWA calculations.[6] The quality of the data used can be insufficient for producing accurate results. Data can be incomplete and data usage can also be imperfect—as, for example, when data on collateral do not make their way from front-office systems to the RWA calculation engine. By improving RWA accuracy and processes, banks can often reduce their RWA under both the SMA and the IRB approaches. They can also avoid RWA increases, which otherwise could force extreme measures, such as exiting business lines. While many banks have already conducted an accuracy-improvement program for RWA, significant opportunity to reduce RWAs and improve economic profit remains. Even for banks with solid data accuracy in their IRB portfolios, further improvements are often possible. To make them, banks can use technical tools that do not require significant investments. Their impact,

however, is often significant: an RWA reduction of €1 billion, for example, typically corresponds to an increase in economic profit of €10 million to €15 million.[7]

Many banks rightly focus on RWA accuracy, but considerable opportunity lies in reducing other capital needs, including capital deductions (such as minority interests, goodwill, intangibles, and nonconsolidated investments), capital buffers (for G-SIBs, under Pillar 2, and the countercyclical buffer), and trapped capital. For some banks, the Pillar 1 requirements constitute only half or less of their total capital requirements; the remainder is determined by these other demands. Depending on the situation, banks can take several actions to bolster capital. These include increasing RWA in entities with an excess of CET1, netting intangibles and goodwill deductions with linked deferred tax liabilities, and reviewing activation policies and amortization periods of expenses related to intangible assets.

Some banks have already begun to take a more holistic approach to improving capital. One global bank increased its capital ratio significantly by implementing a number of diverse measures. The bank correctly classified intangible assets, applied netting procedures in deferred tax assets and deferred tax liabilities to reflect goodwill and pension-fund deductions properly, and adjusted its legal-entity setup and asset-booking locations in line with minority-interest deductions. Another global bank significantly reduced its RWA by changing the regulatory treatment of one of its major participations. In close alignment with the national regulator, the bank managed to move from a look-through approach for calculating the RWA of the participation to the approach laid out in the most recent capital-requirements directive and regulation (CRD IV, CRR). This enabled the bank to consider the RWA and capital deductions of the participation itself. The overall CET1 ratio at the group level improved by about one percentage point, a result of a group-level RWA decrease of more than 10 percent—countered only by CET1 deductions of about 5 percent of overall CET1.

Until Basel IV rules are finalized

Each bank has its own capital-management plan. In our experience helping banks assess their needs in mitigating the capital impact of Basel IV, we have encountered varying levels of preparedness. Most banks have made good progress in improving RWA accuracy, eliminating data errors, and improving processes. Banks need now to focus on optimizing capital beyond RWA while sustaining their RWA improvements—by embedding the optimized balance sheet into strategic planning, selecting the right capital-steering metrics, and educating the front line on capital consumption and conservation.

As the impact of new regulations will vary by region, bank type, and business model, institutions should make specific impact assessments, identifying the portfolios and business segments that will be most affected. This requires banks to make a self-examination, testing for sensitivity to the new regulatory rules and accounting for new business economics in strategic considerations. Actions should be tailored accordingly, well analyzed in advance, and rigorously implemented. Some regulations will permit banks flexibility and adjustments in compliance (such as the gradual phase-in of the IRB capital floor, now forecast to extend from 2021 to 2027); nevertheless, banks should develop their mitigation plan without delay, at least to reassure forward-looking rating agencies and investors.

No-regrets moves

As final Basel IV rules are still pending, banks cannot fully develop their strategic response. Regardless of the final regulatory outcome, however, certain

actions will improve banks' capital position and risk-return profile in any scenario. These actions will also give banks a timely edge over the competition.

One such move would be to develop originate-to-distribute capabilities. This creates more balance-sheet flexibility through the distribution of assets to yield-searching buy-side firms. Banks could also make certain business changes, such as requiring more collateral, paring down under-utilized lines, adjusting contract clauses, and increasing profitability within underperforming customer segments. Banks should wait before implementing such commercial actions as adjusting product offerings or conducting repricing, however, as these will depend on the implied capital costs of finalized rules.

Correcting RWA accuracy and processes often reduces RWA under both the SMA and the IRB approach and will be beneficial whatever the regulatory outcome. At the same time, improving RWA calculations does not require high investments and is relatively easy to implement. Beyond RWA, other no-regrets moves involve the use of technical levers to improve capital deductions (minority interests, goodwill, intangibles, and nonconsolidated investments) and capital buffers (G-SIBs, Pillar 2, and countercyclical buffers). Finally, regardless of the results of their calculations, banks will need robust capital-steering mechanisms; they should begin considerations now on what adjustments are needed in these mechanisms to reflect stresses and regulatory capital (including diversification effects from IRB output floors). These moves can improve capitalization independent of the final Basel IV decisions.

The impact of Basel IV will, in our view, exceed most estimates, and banks will certainly need to raise more capital. Repercussions—and optimal responses—will vary by bank, according to location, positioning, and business model. The rules are not yet finalized, and the pace of implementation is yet being discussed. For now, banks should prepare for expected outcomes, define mitigating actions, and initiate no-regrets measures. This will give banks a running start on implementing a more efficient capital-management approach, so they can meet capital requirements without suffering a loss in profitability. ■

[1] Basel Committee on Banking Supervision, *Standards: Minimum capital requirements for market risk*, Bank for International Settlements, January 2016, bis.org.

[2] "Governors and heads of supervision announce progress in finalising post-crisis regulatory reforms," Bank for International Settlements, September 11, 2016, bis.org; Stefan Ingves, "Reflections of a Basel Committee chairman," keynote address, 19th International Conference of Banking Supervisors, November 30, 2016, bis.org; for more on McKinsey research on Basel IV, see Stefan Koch, Roland Schneider, Sebastian Schneider, and Gerhard Schröck, "Basel 'IV': What's next for banks?," April 2017, McKinsey.com.

[3] Koch, et al., "Basel 'IV'," pp. 9–14.

[4] As indicated, the standards for the output floor, which would replace the transitional capital floor adopted in Basel I, are still being discussed. The new rule is meant to mitigate model risk and measurement error resulting from internally modeled approaches and would limit the benefit in using internal models for estimating regulatory capital.

[5] Methods such as McKinsey's Balance Sheet Optimizer can model the effects of particular responses to regulatory and business constraints to derive optimal strategies on balance-sheet size and composition as well as the potential to improve profitability.

[6] Erik Lüders, Max Neukirchen, and Sebastian Schneider, "Hidden in plain sight: The hunt for banking capital," January 2010, McKinsey.com; Bernhard Babel, Daniela Glus, Alexander Gräwert, Erik Lüders, Alfonso Natale, Björn Nilsson, and Sebastian Schneider, *Capital management: Banking's new imperative*, McKinsey Working Papers on Risk, Number 38, November 2012, McKinsey.com.

[7] Assuming a capital ratio of 13 to 14 percent and a cost of equity of approximately 10 percent.

Stefan Koch is an alumnus of McKinsey's Cologne office. **Roland Schneider** is an expert in the Frankfurt office, where **Gerhard Schröck** is a partner. **Sebastian Schneider** is a partner in the Munich office.

Risk Practice leadership

Cindy Levy
Global
Cindy_Levy@McKinsey.com

Hamid Samandari
Americas
Hamid_Samandari@McKinsey.com

Philipp Härle
Western Europe
Philipp_Haerle@McKinsey.com

Alok Kshirsagar
Asia
Alok_Kshirsagar@McKinsey.com

Gokhan Sari
Eastern Europe, Middle East, North Africa
Gokhan_Sari@McKinsey.com

Kevin Buehler
Risk Advanced Analytics
Kevin_Buehler@McKinsey.com

Thomas Poppensieker
Risk Knowledge
Thomas_Poppensieker@McKinsey.com

January 2018
Designed by Global Editorial Services

This McKinsey Practice Publication meets the Forest Stewardship Council® (FSC®) chain-of-custody standards. The paper used in this publication is certified as being produced in an environmentally responsible, socially beneficial, and economically viable way.

Printed in the United States of America.