McKinsey&Company

The advanced-analytics solution for monitoring conduct risk

Advanced analytics and machine learning can help institutions “connect the dots” across customer and other data to detect conduct risk comprehensively and cost-effectively.

Juan Aristi Baquero, Joseba Eceiza, Dmitry Krivin, and Chetan Venkatesh

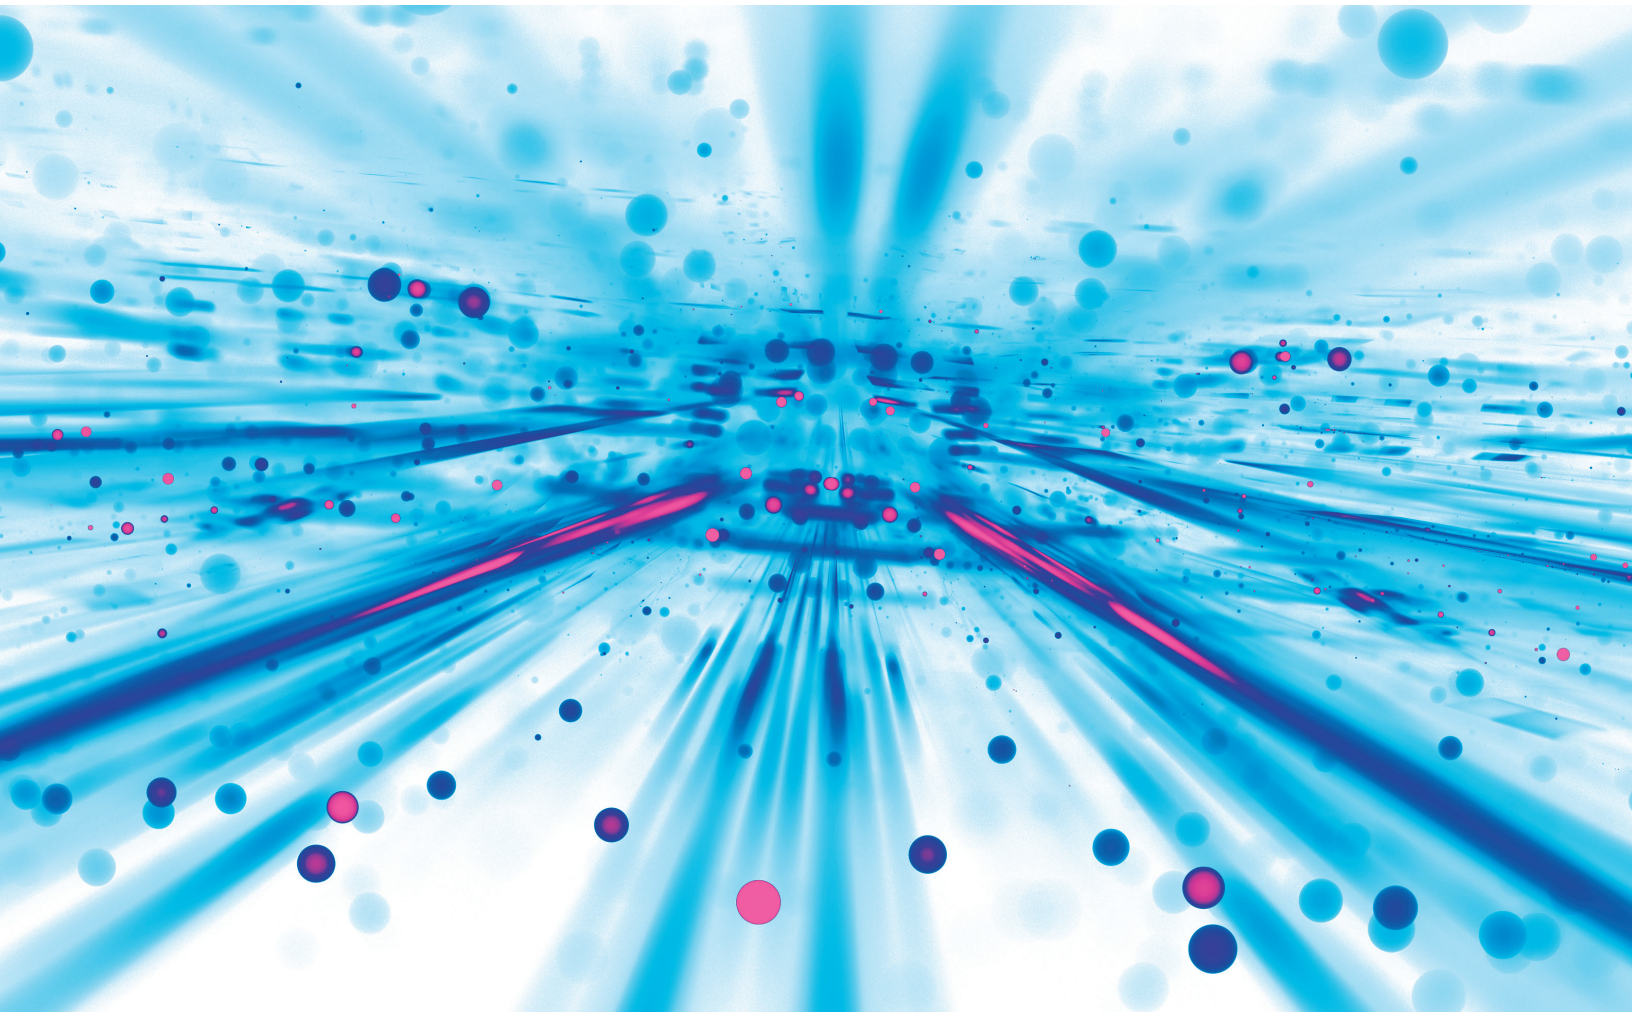

NOVEMBER 2018 • RISK PRACTICE

The fallout from highly visible instances of misconduct—including reputational damage, material losses, and increased regulatory focus—have led financial institutions to treat conduct risk as an important priority. As a risk category, however, conduct has proved difficult to monitor effectively with traditional controls and testing. The varieties of potential misconduct are numerous, and transgressing individuals or whole departments find ever-changing ways to circumvent rules. In addition, sample-based tests such as transactional reviews are not effective in finding isolated instances of misconduct.

Effective misconduct detection requires a new approach, one that can "connect the dots" across individual and team activities. These connections are often hidden in data that derive from multiple sources. They can be revealed by deploying advanced analytics and machine learning to mine the rich data and thereby identify incongruous sales or transaction patterns, misaligned incentives, and inappropriate customer interactions. Frequently underutilized records (such as the transcripts of customer interactions), can be automatically analyzed for potentially inappropriate treatment that customers may have experienced. But advanced-analytics solutions go beyond the detection of past instances of misconduct—by which the damage to an institution, if any, has already been done—to intercept the outlying patterns of activity that could lead to future losses.

What is conduct risk?

The definition of conduct risk varies somewhat by industry and region but can be commonly understood as individual or group actions that could cause unfair outcomes for customers, undermine market integrity, and damage the firm's reputation and competitive position.

Conduct risk has only recently become recognized as a stand-alone risk category, in the aftermath of a number of high-profile incidents of misconduct (and regulatory responses) in retail and commercial banking, capital markets, and wealth management:

- In the United Kingdom, the discovery of a number of episodes of questionable selling practices in retail and small business led the Financial Conduct Authority to publish new regulatory guidance discouraging staff incentives based on sales targets. These changes coincided with a decline of around 40 percent in sales productivity by branch advisers.
- On trading floors, certain individual rogue traders have caused hundreds of millions of dollars in trading losses at several firms around the globe, while others have colluded to undermine market integrity and gain unfair advantages—such as in the London Interbank Offered Rate (LIBOR) fixing case in 2012. The threat of such behavior continues to pose significant reputational and financial risks to major investment banks.
- In the United States, in 2016, the Consumer Financial Protection Bureau and the Office of the Comptroller of the Currency reinforced customer protections against unauthorized account opening or unrequested enrollment in bank services.
- In Canada, several articles were published in national media outlets in March 2017 claiming aggressive sales tactics by the major banks; the allegations led to internal investigations, hearings before the finance committee of the House of Commons, and a review of sales practices by the Financial Consumer Agency of Canada.

- In Australia, the Royal Commission into Misconduct in the Banking Industry was established in 2017 and is scheduled to make its final report in early 2019.

As a result of these incidents and regulatory responses, many banks are reviewing their conduct-risk approaches and some are revising key components, including the definition of conduct risk, risk taxonomy, risk detection and monitoring, policies and procedures, roles and responsibilities, and issue remediation.

Common risk-monitoring approaches are inadequate for conduct risk

Conduct risk is different from most other types of risk because it entails the great variety and complexity of aberrant human (and organizational) behavior. Since its causes are idiosyncratic, it is impossible to capture the essence of conduct risk in a few quantitative measures—such as those employed for other major risk types (whether value at risk for market risk or expected loss for credit risk). Nevertheless, a single instance of misconduct can have severe negative effects on an institution.

Because it is impossible to quantify simply, conduct risk has not been adequately addressed by traditional methods of risk detection. These methods are generally incapable of actively isolating rare instances of misconduct—instances that can nonetheless cause significant harm to financial institutions.

Apart from communications surveillance, an emerging approach primarily used in trading, conduct risk has mainly been monitored with three approaches: monitoring of customer complaints or internal whistle-blower reports; activity testing, such as verification of customer signatures for new-account opening in branches or sample-based call monitoring in contact centers; and rules-based analytics, such as trade alerts used to monitor activities of wealth advisers or trade surveillance in trading.

No matter how institutions apply these three approaches, they find themselves failing to detect conduct-risk issues comprehensively. For instance, while customer complaints and whistle-blower reporting are necessary elements of conduct monitoring, they cannot substitute for a more complete program. They are lagging indicators first of all, providing signals only after damage to an institution has been done. Additionally, a majority of conduct issues, such as inappropriate selling, go unreported by customers, who may not even be aware of the issue. Incidents that are reported are frequently resolved as soon as the customer contacts the institution, without a complaint ever having been filed. Finally, from the slight percentage of incidents for which a complaint is filed, an institution will find it difficult to identify the sources of the majority of incidents. Typically, a small number of employees or departments will have been responsible, but the common approaches won't find them (Exhibit 1).

Similarly, activity-testing controls, such as branch audits or verification of customer consent, have a number of problematic limitations. First, unless conducted on all or nearly all transactions, they will be ineffective in identifying rare instances of sales misconduct. Second, employees are generally aware of controls in place and can avoid them by focusing on untargeted transactional patterns. Finally, controls are usually manual, making coverage of all possible patterns of sales misconduct prohibitively expensive.

Exhibit 1 **Even with several levels of manual checkers and investigators, most methods for monitoring customer complaints fail to identify the few common sources of misconduct.**

The third common approach, rules-based analytics, has also been proven to be an insufficient way to address conduct risk. This approach uses such techniques as trade alerts to monitor wealth-management advisers, or, in retail banking, thresholds to target those originating an excessive number of unfunded accounts. The approach tends to generate a large number of false positive alerts for further—and costly—investigation. It is also easy to dodge, for example, by selling marginally below known thresholds.

These approaches to conduct-risk identification are beset by additional shortcomings. One is that they rely one-sidedly on numerical data. Unstructured data such as customer-call recordings or surveys are rarely used to good effect in the monitoring framework. Another is a tendency to treat individual factors in isolation, without connecting them in sequence. Given the diversity of conduct-risk activities, the most powerful insights lie in the discovery of patterns across multiple sources; for example, employee sales, customer calls, and incentive plans. Together, structured and unstructured sources of data can help institutions address misconduct more accurately, with far fewer false positives.

A better way

To effectively monitor and detect conduct risk, institutions need a new method, one that leverages the power of data from diverse sources, including customer feedback, sales and product data, and performance-management data. An inclusive data model—one that respects all local laws and regulations—will permit institutions to "connect the dots" across the activities of individuals and departments. Machine-learning algorithms can mine a complex data terrain to establish outlying activities and identify potential instances of misconduct. Designated outliers can then be captured automatically from all recorded customer interactions.

Making use of advances in data and analytics, institutions can transform conduct detection and replace extensive manual controls and verification activities. A number of leading institutions have started on this journey, putting in place monitoring analytics that detect infrequent instances of misconduct, such as inappropriate sales, before significant financial and reputational damage is sustained. An effective conduct-risk analytics monitoring program will be defined by the following capabilities:

1. Connecting the dots across individual activities

A data model can link sales, transaction, and performance data with other sources of insight at the level of individuals and departments (or teams and branches). To build it, organizations need to create data lakes or repositories of structured data (such as sales and account usage) and unstructured data (such as customer-call transcripts, surveys, and complaints). Captured are transactional and sales-performance data, customer patterns (such as portfolio activity for wealth management), and customer intelligence (such as call records to service centers, surveys, complaints). The sources of insight are assembled in granular detail, providing views of the activities of individuals and sales teams or branches. The data are contextualized with additional details, such as management hierarchy and office location. Taken together, the data should fully represent the sales and account activity associated with each individual and can be analyzed systemically to identify suspicious patterns.

Depending on the line of business or type of conduct risk being monitored, other data sources can be included. At the center of the fraud and collusion scandal surrounding the London Interbank Offered Rate in 2012, for example, were traders using chat rooms, instant messaging, and email. Key information from such communications can often be retrieved from corporate-communication platforms and added to the data model. Techniques using advanced analytics can work through this type of data and identify potentially suspicious behavior (Exhibit 2).

Exhibit 2 **The advanced-analytics approach to conduct monitoring 'connects the dots' drawn from many sources of data.**

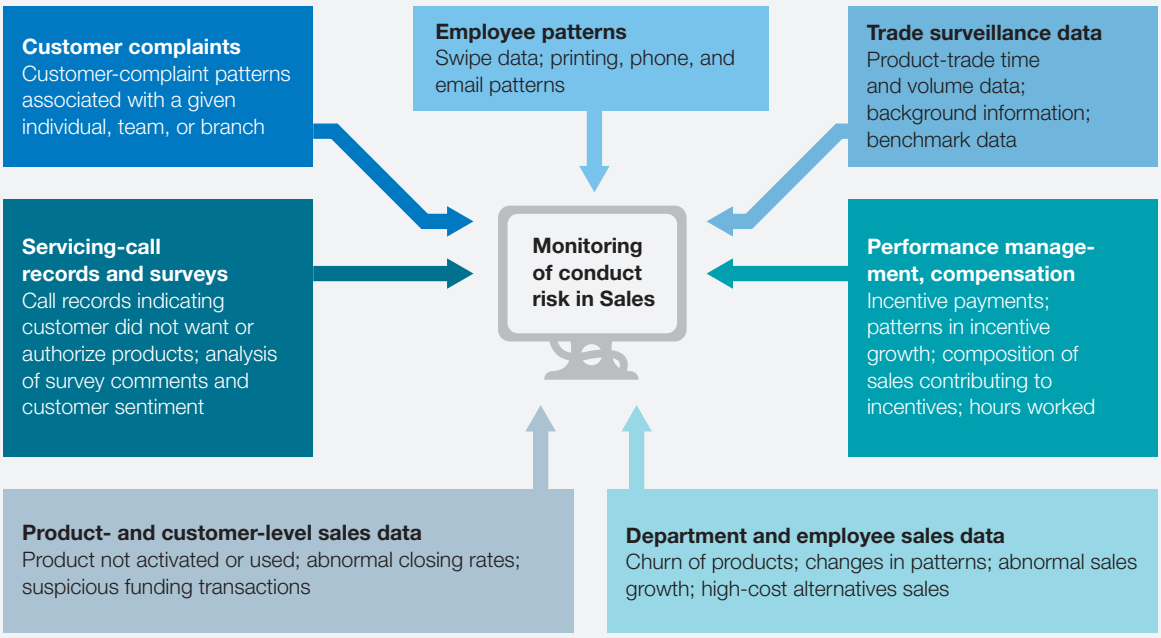

2. Finding the needle in the haystack

Advanced analytics applied across structured and unstructured data can help classify behavior and detect suspicious or outlying patterns. To develop robust analytics, including machine learning, data scientists work closely with the businesses and control functions to test hypotheses for behaviors defined as indicators for misconduct. This set of analytics will form a control group for identifying outliers. The algorithms are designed to detect these broad patterns, rather than individual instances of misconduct. The approach isolates patterns of activity that the business knows are associated with conduct risk (see sidebar "Mining for patterns associated with conduct risk in wealth management and retail banking").

In addition to mining for known misconduct patterns, machine learning can also be applied to detect previously unknown pattern anomalies. This application, unsupervised learning, is used to mine employee data for new suspicious patterns not identified in the past. For example, a company using unsupervised machine learning identified an employee who was sharing his ID credentials, after the outlier-detection algorithm detected two logins occurring close in time but at locations that were far apart in distance.

Known conduct-risk markers, coupled with unsupervised techniques and additional purely unsupervised techniques for anomaly detection, can be a powerful combination for managing known potential risks and uncovering new and emerging risks. Exhibit 3 illustrates a machine-learning algorithm called the isolation forest, which can identify outlying patterns while distinguishing between positive and negative outliers.

3. Mining customer interactions with natural-language processing

A great amount of data collected from customer interactions with financial institutions is text based, including transcribed phone conversations (see sidebar "Voice-to-text technology"). This kind of data, which is often underutilized, can provide rich insights for conduct-risk detection while also improving the customer experience.

Natural-language processing (NLP) is the branch of artificial intelligence devoted to enabling computers to respond to written or spoken comments and commands given in "natural languages," such as English or Chinese. NLP converts linguistic syntax into computer-readable numeric codes and responds using machine-learning algorithms. Increasingly sophisticated language models have enabled pattern identification within highly specific, tailored contexts. The capabilities of NLP have grown dramatically in the past decade, as has public awareness, with the proliferation of customer-support chatbots and virtual assistants. The application of NLP to text data is a proven approach for analyzing and interpreting customer interactions. The technology can be used to

To develop robust analytics, including machine learning, data scientists work closely with the businesses and control functions to test hypotheses for behaviors defined as indicators for misconduct.

Exhibit 3

Unsupervised machine-learning algorithms such as the isolation forest can detect irregular patterns and filter for potential misconduct.

classify these interactions and identify misconduct, capturing the context of customer dissatisfaction, including its immediate or original source.

In one application, an NLP model worked on a large number of customer surveys (more than 100,000). With and without the use of key words, the model was able to capture instances of potential sales misconduct in customer complaints from several dozen unique surveys. The model highlighted complaints of representatives changing customers' plans without their knowledge, failing sufficiently to explain how products work, and pressuring them to purchase ill-suited products. One complaint specified a hard sell for overdraft protection—including a voluble claim by the teller that it would help the customer's credit rating—despite the customer's insistence that he would

Mining for patterns associated with conduct risk in wealth management and retail banking

Wealth management

Trading and pricing activity

- Excessive churn
- Generation of new sales from rollover products
- Insufficient attention or reverse churn in managed accounts
- Abnormal investment performance compared with clients with similar portfolio composition

Product activation and reversals

- Abnormally high claw-back levels
- Cancellations or reversals of products sold (for example, overdraft, credit protections)
- Excessive sales funded with short-term transfers

Compensation trends

- Abnormal compensation composition or patterns (such as the same or increasing compensation levels with decreasing assets in the portfolio)
- Excess trading near next compensation grid
- Abnormal claw backs due to reversed trades or client gestures

Concentration risk

- Abnormal single-position concentration risk
- Excessive number of unique positions (outside separately managed accounts)
- Suitability: misalignment of portfolio holdings with client risk-appetite statement

Sales-pattern activities

- Outlying growth in sales and incentives, not explained by tenure or hours
- Product and account churn
- Profile of sales skewed toward products with verbal consent

Own account and outside activities

- Excess profitability of own account
- Outside business activities
- Inappropriate investment in private securities

Retail banking

Sales-pattern activities

- Outlying growth in sales or incentives not explained by tenure or hours
- Profile of sales skewed toward products with verbal consent
- Excessive reliance on sales of secondary and tertiary accounts
- Product or service churn on the same account

Product activation and reversals

- Abnormally high claw-back levels
- Downgrades from higher-incentive products
- Cancellations or reversals of products sold (for example, overdraft, credit protection)
- High share of sales of unused products (for example, cards, deposits)

Team-level risk markers

- Team results from operational and compliance reviews
- Team-level complaint trends
- Network shift between branches

Internal risk markers

- Abnormal number of address changes
- High level of products sent to branches at time of sale
- HR referrals related to ethics or sales integrity

Voice-to-text technology

Voice- (or speech-) to-text technology converts audio files of speech to text. The technology has been greatly enhanced recently, through improvements in computing power and the refinement of "deep neural networks." These are sets of algorithms, named after the physiology of thinking, that can cluster and classify large quantities of data in highly sophisticated, customizable ways. The accuracy of transcriptions created using voice-to-text technology has consequently improved dramatically: greater than 90 percent accuracy has been achieved in some test data sets. Better transcriptions of telephone conversations—approaching human accuracy—opens the way to an application of voice-to-text in many use cases, including conduct risk and sales performance.

never use this protection. Another suggested that customers were being signed up for credit insurance without their knowledge. Even taking into account the time needed to develop the model, the automated process saved many hours of human labor in precisely identifying outlying patterns for further investigation.

4. Employee-conduct transparency

By capturing all data on employee conduct and coupling the data with contextual details such as branch, supervisor, and tenure, organizations can build a comprehensive picture of employee performance. As the analyses are targeted at specific behaviors, run on all employees, and normalized through analytic methods, each employee can be compared against the rest of the organization, or a cohort (such as tenure bands). Three useful outcomes can be gleaned from this type of reporting:

- Trends can be identified and specific interventions can be developed before a misconduct case occurs.
- Chronic behavior patterns can be mined and treated with appropriate behavior-improvement training or product controls.
- Systemic and prevalent behaviors can be identified by aggregating the standardized data to supervisor, branch, or district levels.

Overcoming practical challenges

Technological and psychological challenges to developing an analytics-based conduct-risk program may arise. Described here are some common challenges and how they can be resolved.

Insufficient or siloed data. Early in their application of a data strategy, many organizations are unable to monitor employee conduct effectively because they have integrated too little data to build a full picture of employee activity. The initial limitation can be addressed by building out the program gradually, beginning with the most critical data (such as sales and product data, account-activity patterns, and data relating to incentives), and incorporating additional data and analytics over time. Data marts— subject-oriented

databases created for specific purposes—can also be developed, with a view to incorporating them comprehensively or in part, into a data lake over the long term.

Insufficient expertise or resources in data science and advanced analytics. Commonly, banks find that ingoing levels of expertise and resources are insufficient for their analytics ambitions. Overcoming such capability deficits is not difficult, however. Most banks have already created groups of data scientists to develop many forms of machine- learning code, even if they are less familiar with more advanced deep-learning or natural-language-processing algorithms needed for advanced conduct monitoring. These internal resources, complemented with initial external support and/or specific recruiting, can quickly add the required skills as individual use cases are built. Before long, the technical side of conduct-risk monitoring can be managed internally.

Organizational reluctance. Some organizations are reluctant to make a large risk-management investment without evident business benefits. A number of banks have expressed concerns about building conduct-analytics infrastructure for what they see as a purely defensive play, particularly if regulatory examinations did not discover a widespread cause for concern. While the decisions to invest ultimately depend on the risk appetite of each institution, the data and analytics investment described above can generate positive profit-and-loss impact in addition to mitigating risks. First, deploying analytics-based conduct-risk monitoring allows institutions to retire expensive manual controls, testing activities, and investigations of false positives associated with traditional risk-management methods. Second, the same analytics used to mine for conduct risk can also unearth business insights. For instance, the same natural-language-processing engine used to find sales misconduct can also find insights into customer perception of specific products and services. Likewise, data and analytics used to detect suspicious patterns can also identify the sales and behavioral patterns associated with top-performing sales associates. These insights can then be incorporated into training and performance-improvement programs.

■ ■ ■

Heightened awareness of business misconduct has affected financial services in challenging ways, exposing even the largest institutions to reputation risk and regulatory scrutiny. Institutions have responded to the challenges, but monitoring the conduct of thousands of employees across many activities, locations, and business units is a complex problem. Fortunately, controls based on advanced analytics and machine learning offer institutions an alternative to a costly infrastructure of manual checkers and investigators. The new approaches enable the effective and efficient monitoring and detection of employee conduct-risk issues before they become serious incidents. A number of advantages accrue to institutions implementing these advanced control programs: they avoid losses of various kinds, they instill confidence in the front line and in regulators by addressing conduct risk in a timely manner, and they create value too, through improved customer relationships. Can you say, "Win-win-win"?■

Juan Aristi Baquero is a partner in McKinsey's New York office, where **Dmitry Krivin** is a partner and **Chetan Venkatesh** is a consultant. **Joseba Eceiza** is a partner in the Madrid office.