McKinsey&Company

McKinsey on Payments

Special Edition on Advanced Analytics in Banking

Highlights



4 The analytics-enabled collections model

13 How machine learning can improve pricing performance

36 Hidden figures: The quiet discipline of managing people using data


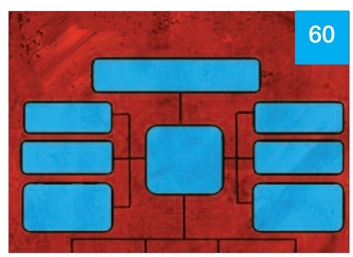
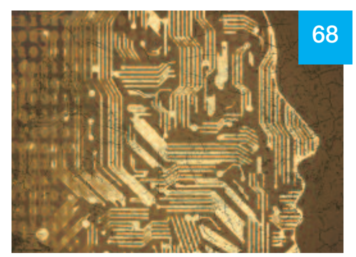

52 Designing a data transformation that delivers value right from the start

60 Building an effective analytics organization

68 "All in the mind": Harnessing psychology and analytics to counter bias and reduce risk

Volume 11 Number 28 August 2018

McKinsey on Payments is written by experts and practitioners in the McKinsey & Company Global Payments Practice.

To send comments or request copies, email us at:
paymentspractice@mckinsey.com

To download selected articles from previous issues, visit:
https://www.mckinsey.com/industries/financial-services/our-insights/payments

Contents

4 The analytics-enabled collections model

13 How machine learning can improve pricing performance

20 Combating payments fraud and enhancing customer experience

28 Using data to unlock the potential of an SME and mid-corporate franchise

36 Hidden figures: The quiet discipline of managing people using data

44 Using analytics to increase satisfaction, efficiency, and revenue in customer service

52 Designing a data transformation that delivers value right from the start

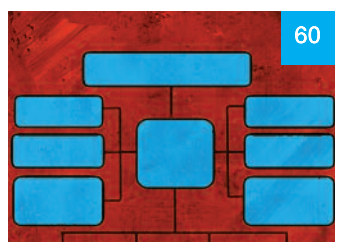

60 Building an effective analytics organization

68 "All in the mind": Harnessing psychology and analytics to counter bias and reduce risk

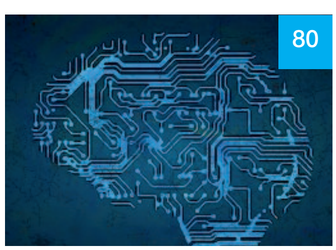

80 Mapping AI techniques to problem types

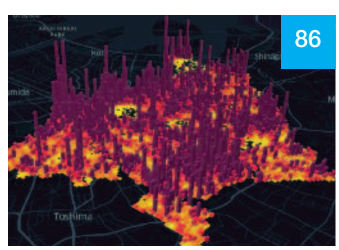

86 Data sheet: Advanced analytics

Foreword

In this special edition of *McKinsey on Payments*, we examine the emergence of advanced analytics as one of the biggest disruptions in financial services, influencing almost every part of the industry. A recent McKinsey Global Institute study estimated the annual potential value of artificial intelligence in banking at as much as 2.5 to 5.2 percent of revenues, or $200 billion to $300 billion annually, based on a detailed look at over four hundred use cases.

The importance of data and analytics in banking is not new. The 1950s and 1960s saw innovations such as credit scoring in consumer credit, and the use of market data for securities trading, driven by the desire for more data-driven decisioning. The 1970s and 1980s unleashed direct marketing of credit cards (Citibank's "drop" of pre-approved credit cards to 20 million Americans in 1977 instantly made it one of the largest US issuers) and the emergence of new financial derivatives and risk-transfer vehicles, enabled by the growth of mainframe computers. The 1990s and 2000s saw the institutionalization of data and analytics, with the launch of the FICO score in 1989 by Fair, Isaac and Co, and the introduction in 2004 of Basel II, which embedded a data-driven analytical framework into bank regulation. This period, during which computing and storage power skyrocketed, also featured the emergence of nimble credit card players in the US such as Capital One, Providian, MBNA, and First USA, and of hundreds of quant-driven hedge funds globally in wholesale banking.

But it is this decade's explosion of available data and ubiquitous channels that is starting to fundamentally change the industry. It was once estimated that five exabytes of data were created from the dawn of civilization through 2003. By 2016, the world was producing two exabytes an hour, a pace that is still increasing. It is therefore no surprise that banks and payments firms globally are struggling to leverage the new data available, while hampered by legacy infrastructure, outmoded organizational structures and skills, and the limitations of regulatory guardrails.

To truly understand how some banks are starting to change, it helps to look at four distinct pillars:

1. **Organizing for use cases.** Some banks believe in starting by building the right infrastructure. Unfortunately, "build-it-and-they-will-come" strategies rarely succeed. While there is no available estimate of wasted spend, we do know that 70 percent of large projects fail overall. Far more effective is a "customer-back" use case-driven approach, often clustered into domains. For banking, we have identified hundreds of cases of analytics-driven impact, which range from using simple "random forest" algorithms to predict call traf-

McKinsey Analytics

McKinsey Analytics helps clients achieve better performance through data, working with more than 2,000 clients to build analytics-driven organizations, and help them develop the data strategies, operations, and analytical capabilities to deliver sustained impact.

In the last few years, McKinsey Analytics has grown rapidly, acquiring seven specialty analytics firms and partnering with more than 150 other providers. Today, McKinsey Analytics brings together over 2,000 advanced analytics and AI experts and spans more than 125 domains (industry- and function-specific teams with people, data, and tools focused on unique applications of analytics), working with client, external, and McKinsey proprietary data in a secure environment. QuantumBlack, an advanced analytics firm that McKinsey acquired in 2015, is now a core part of McKinsey's research, development, and delivery of AI techniques and real-world applications for clients.

Learn more at www.mckinsey.com/business-functions/mckinsey-analytics/our-insights and https://www.quantumblack.com/

fic, to leveraging natural language processing to scan résumés (as McKinsey does) or contracts, to taking advantage of deep learning and image processing to detect fraud. Mapping out opportunities and their business cases is the right first step, along with the front-line adoption necessary to realize value.

2. **Building a useable data and analytics infrastructure.** Ubiquitous data does not always translate to useable information, though most banks have a treasure trove of data that they could be using today. Many institutions are still early in the process of defining their data infrastructure, and the risk of bad data and models cannot be overstated. Based on which use cases have the greatest priority, issues that banks are sorting out are (1) the economic, service, and regulatory trade-offs between public cloud, private cloud, and other on-premise solutions, (2) challenges in defining and enforcing an effective data model which requires front-line and business commitment and shared investments, and (3) how to build an effective chief data officer (CDO) function with real teeth to manage the process and maintain effective quality control.

3. **Developing a standardized analytical environment.** There is a proliferation of analytical toolkits today. Some firms are literally "boiling the ocean" of available data to gain insight. For instance, McKinsey's automated analytical engines can now test millions of permutations in minutes, incorporating external data sources like market information, census data, the weather, and mobile location data. Mature banks are now focusing on how best to balance

toolkit innovation with standardization to move from "one offs" to enabling analytics at scale. Agile approaches are particularly effective, especially when small differences can mean big money—for instance, when response rates to digital marketing campaigns are less than five basis points, the slightest edge is valuable.

4. **Building an effective organization.** Banks are becoming more thoughtful about the organizational infrastructure and people to support a data and analytics transformation. These are some of the most in-demand roles today and firms are using a combination of build versus rent models to accelerate execution. Universities are trying to fill the gap but not fast enough—last year, MIT's Masters in Business Analytics program received the highest application rate per open seat of any master's program at the institute. Within banks, the debate is about how best to organize and deploy these resources—ranging from centers of competence to distributed groups closer to the business. When banks do set up central analytics groups, these average 2.8 percent of the total employee base—nearly triple the percentage for other industries, according to McKinsey's Analytics Quotient tool.

For this issue, we have selected a set of complementary articles to exemplify both the opportunities and challenges faced by financial firms. We start with six innovative use cases demonstrating the potential of next-generation analytics to serve as a differentiator in collections, pricing, fraud detection, segment revenue growth, talent management, and customer service. We then pivot to three foundational articles on designing a data transformation for value creation, building an effective analytics organization, and detecting and reducing bias in implementation. We end the issue with a primer on mapping AI techniques to problem-types, and our customary data pages.

We hope you find this issue thought-provoking and that these articles generate constructive conversations. As always, we welcome your feedback at *paymentspractice@mckinsey.com*.

Vijay D'Silva *is a senior partner in McKinsey's New York office.*

The analytics-enabled collections model

The global credit environment absorbed the effects of the financial crisis at varying speeds from market to market. In some places, loss rates have remained relatively high since 2008–09; in others, the past decade has been one of steady improvement, with tapering losses that have only recently begun to climb again. In the expanding markets, lenders increased their risk exposure, issuing new products designed around easier underwriting guidelines. Little attention was paid to maintaining or improving collections capabilities. As debt loads rise, however, institutions in these markets are beginning to rebuild collections staff and skills that eroded in the previous period. Meanwhile, in the more stressed markets, the need for more efficient and effective collections operations is likewise becoming a priority.

Ignacio Crespo
Arvind Govindarajan

The need to renew collections operations provides lenders with an ideal occasion to build in new technologies and approaches that were unavailable when the financial crisis hit. The most important advances in collections are being enabled by advanced analytics and machine learning. These powerful digital innovations are transforming collections operations, helping to improve performance at a lower cost. Better criteria for customer segmentation and more effective contact strategies are being developed. Individual collector performance is being improved with better credit-management information and other tools. Contact can be managed through an array of channels, some allowing customers a greater sense of control over their finances. Loss-forecasting strategies can also be made more accurate and predelinquency outreach made more effective with enhanced financial tools and mobile apps.

Some of the most significant advances brought about by advanced analytics and machine learning are in customer segmentation, which is becoming much more sophisticated and productive. Better segmentation—including innovative behavioral segmentation, discussed in detail in "'All in the mind': Harnessing psychology and analytics to counter bias and reduce risk," on page 68—is providing the basis for more effective collections processes and strategy. The improvements affect the complete collections agenda, beginning with the prevention and management of bad debt and extending through to internal and external account resolution.

A next-generation collections model

In traditional collections processes, banks segregate customers into a few simple risk categories, based either on delinquency buckets or

Editor's note: This article first appeared in *McKinsey on Risk* in June 2018. Since then, delinquencies across many asset classes have continued to rise, and the importance of analytics in collections has only increased.

on simple analytics, and assign customer-service teams accordingly. Low-risk customers are usually given to newer collections agents based on availability; the agents follow standardized scripts without being asked to evaluate customer behavior. Agents with moderate experience, training, and skills are assigned, again based on availability, to medium-risk customers. These agents also follow a standardized script but are trained to assess customer behavior based on ability and willingness to pay. High-risk customers are assigned to the most skilled agents, who own their accounts and use less standardized approaches to develop assessments of customer behavior. Contact strategies and treatment offerings are fairly varied across the risk categories.

By using advanced analytics and applying machine-learning algorithms, banks can move to a deeper, more nuanced understanding of their at-risk customers. With this more complex picture, customers can be classified into microsegments and more targeted—and effective—interventions can be designed for them (Exhibit 1).

Using analytics in the new model

Analytics-based customer segmentation is at the center of the next-generation collections model. The transformed collections model will allow lenders to move away from decision making based on static classifications, whether these are standard delinquency stages or simple risk scores. Early identification of self-cure customers will be one bene-

Exhibit 1

Advanced analytics and machine learning can classify customers into microsegments for more targeted interventions.

Customer type	**True low-risk**	**Absent-minded**	**Dialer-based**	**True high-touch**	**Unable to cure**
Targeted intervention	Use least experienced agents provided with set scripts.	Ignore or use interactive voice message (segment will probably self-cure).	Matching agents to customers and live prompts to agents to modify scripts.	Focus on customers able to pay and at high risk of not paying.	Offer debt restructuring settlements early for those truly underwater.
Impact	Agent–client conversation guided by onscreen prompts based on probability of breaking promises.	10% time savings allows agents to be reassigned to more difficult customers and specific campaigns.	Can lead to increased “connection” and higher likelihood of paying.	Added focus addresses higher probability of default rates in this segment.	Significant increase in restructuring and settlements enhances chance of collecting at least part of debt.

Source: McKinsey analysis

fit. Another will be an approach based on value at risk, rather than blanket decisions based on standardized criteria. The aspiration is to have every customer as a "segment of one" with customized treatments.

Most banks can achieve significant results in all key collections areas by introducing an analytics-based solution quickly and then making needed improvements as they go.

Leaders taking the analytics-based actions that define the new model have already begun to realize gains in efficiency and effectiveness. One European bank automated 90 percent of communications with clients by developing two advanced-analytics models using machine-learning algorithms. A binary model identifies self-curers and non-self-curers, and a multiclass model recommends collections strategies for the non-self-curers, including soft measures, restructuring, or work outs. The models use around 800 variables, including client demographics and information on overdrafts, client transactions, contracts, and collaterals. The bank has realized more than 30 percent in savings with no loss in operational performance.

Another European bank set out to develop a top-notch recovery process using advanced analytics. The goals were to minimize the number of clients falling 90 or more days' past due while maximizing the economic impact of exits, focusing on retail and small-and-medium enterprise portfolios. As the bank gained a deeper understanding of its nonperforming loans, it was able immediately to address certain borrowers (such as recurring defaulting clients) with effective initiatives. Other groups of clients were identified, and exit strategies based on economic value were developed for each group. The results are compelling. The bank reduced its 90-day-or-more portfolio by more than €100 million, with €50 million in fewer past-due entries and the remainder in exit acceleration. Moreover, a reduction of 10 percent in past-due volumes was achieved across the board, worth around €300 million less in past-due exposure.

A leading North American bank has rolled out a number of machine-learning models that improve the estimation of customer risk, identifying customers with a high propensity to self-cure as well as those suitable for early offers. These models have so far enabled the bank to save $25 million on a $1 billion portfolio.

Most banks can achieve results of this magnitude by introducing an analytics-based solution quickly and then making needed improvements as they go. Value can be gained in almost all of the key areas in the collections environment.

- **Early self-cure identification.** Some banks use rudimentary heuristics (rules of thumb) or simple models to identify self-cure customers, while others have adopted simple self-cure models with limited variables. The new self-cure model based on machine learning and big data can save collectors a lot of time. By using many variables to better identify self-cure accounts, banks can increase collector capacity by 5 to 10 percent, allowing agents to be reassigned to more complex collections cases.
- **Value-at-risk assessment.** While many banks use time in delinquency as the primary measure of default risk, some lenders

are taking a more sophisticated approach, building a risk model to determine value at risk. Many of these are simple trees and logistic regressions, however, with limited data. Leaders are moving to a future state in which models project conditional probability rather than assign customers single risk scores. The conditional score is dependent on a range of tailored approaches to customer contact and engagement: every borrower has several scores depending on the contact strategy and offer. Lenders would then use the strategy and offer that optimizes recoveries. The approach better calibrates the intensity of contact with each account, thus optimizing resources. A next-generation value-at-risk assessment can further reduce charge-offs by 5 to 15 percent depending on maturity of current operations, analytics, and availability of data.

Lenders at the analytics forefront are assembling masses of data from many kinds of sources and developing different models to serve collections goals.

- **Cure assessment versus pre-charge-off offers.** At most banks, agents determine whether a customer will cure or will need an offer of some sort; some banks have heuristic rules for agents to follow. The new approach is to use models that ascertain a customer's ability and willingness to pay and gauge whether the better path is a cure or an offer. Banks can resegment delinquent accounts to improve their decisions to offer early settlement, an approach that increases the uptake of offers while reducing charge-offs by 10 to 20 percent.
- **Optimizing pre-charge-off offers.** Banks are currently using rules or simple analytics to create offers for customers, often without determining the likelihood that they will accept. Models will predict the best offer, optimized for the needs of the bank and the customer. Banks can change the prompt, adjusting loan characteristics and offerings to those most likely to reduce charge- offs, including reamortizing the term or interest rate, consolidating loans, or settling. Making the right offer early, before accounts enter late-stage delinquency, can improve acceptance rates.
- **Post-charge-off decision.** Most banks use simple models or heuristics to determine which agencies to send accounts to and at what price. To refine these decisions, models will determine the best agency for each account and tailor prices accordingly. The model will also determine the optimal pricing segmentation for third-party agencies and identify the accounts to retain in-house longer (based on products retained with the bank, for example). The strategic use of third parties can help with accounts that cannot be cured internally.

Integrated analytics models

Lenders at the forefront of the analytics transformation are assembling masses of data from many kinds of sources and developing different models to serve collections goals. The data sources can include customer demographics, collections and account activity, and risk ratings. The most sophisticated lenders are creating "synthetic" variables from the raw data to further enrich their data. Machine learning helps identify markers for high-risk accounts from such variables as cash-flow status, ownership of banking products, collections history, and banking and investment balances. By using so many inputs from many different systems, lenders can dramatically improve model accuracy, lower charge-off losses, and increase recovered amounts. Two separate institutions recently adopted similar approaches using more than 100 variables to support numerous machine-learning models. These issuers used machine learning to identify the optimal treatment and contact strategy for each delinquent account, deployed the solution inside the existing collections work-flow environment, and trained collectors to use the system and collect additional data to improve model performance. The initiatives were up and running in about four months (Exhibit 2).

Contact strategies and treatment approaches

Institutions adopting the most analytics-forward approaches have been intensifying the development of new treatment and contact strategies, expanding the limits of digital capabilities. By applying advanced analytics and machine learning, banks can identify the

Exhibit 2

Two major issuers used machine learning and more than 100 variables to accelerate development of treatment and contact strategies.

Integrated data sources on client behavior

- Customer demographics
- Collections activity
- Account activity
- Payments
- Risk ratings

• 140 inputs across 15 systems describe profiles of each client on each day in collections

Multiple machine-learning models used to identify features of high-risk accounts

- Low cash flow
- More banking products
- Previously in collections
- High total balance across products
- Low investment balance

• 30–40% improvement in model accuracy versus previously existing models

Implemented in the existing collections environment

- Contact-center interface
- Banking machine (automated touchpoint)
- Contact customer via call or text (automated touchpoint)
- Interactive voice response (automated touchpoint)

• Implemented in 12–16 weeks

Source: McKinsey analysis

most promising contact channels while also developing digital channels to define innovative and regulatory-compliant contact strategies. The same digital channels can be used to build awareness of payment options.

- **Websites.** Display messages and repayment options as soon as customers log in, increasing awareness and providing opportunities for early delinquency reduction.
- **Messenger and chat.** Where legally permissible, collectors can contact customers and negotiate payment options with chat functionality and free messenger applications (such as WhatsApp).
- **Mobile apps.** Build collections functionality into the mobile app, reminding customers in early delinquency stages to pay and offering payment options.
- **Virtual agent.** Create capacity by developing virtual agent functionality to call customers in early delinquency stages.
- **Voice-response unit.** Enhance current voice-response capability, offering basic repayment options when customers call, which frees collector capacity.

Most banks use heuristics to establish the best times to call. Usually, however, agents are inadequately supported on questions of which channel to use, when to use it, and what the message should be. Advanced models can project a full channel strategy, including channel usage, timing, and messaging. Banks will be able to control contact down to the hour and minute, as well as the sequence of communications—including voice, text, email, letter, and interactive voice message. The approach is developed to maximize the right-party-contact rate and influence customer behavior to prioritize payment. Such optimal contact sequencing can increase success in early stages of delinquency.

The analytics focus on the front line

Leading companies in many sectors—digital giants, healthcare providers, retailers, and manufacturers— are using data and analytics to develop a workforce optimized to business goals. Analytics is now the source of improved performance in realizing talent strategies as well as a means for linking talent strategy to business needs. (See “Hidden figures: The quiet discipline of managing people using data”, page 36.) Presently, recruiting and retention are often based on legacy processes, including résumé screening and interviews; retention is based solely on performance. Analytics can improve hiring, finding agents with affinities to the most valuable at-risk segments, as well as help identify collectors at risk of leaving. Companies are using machine-learning algorithms to screen résumés and to determine the value of external hiring compared with internal promotion. One global digital company used analytics to create a checklist that boosted onboarding speed by 15 percent. The algorithms, it should be stressed, are not replacing human judgment but are rather providing a deeper fact base for the exercise of informed judgment.

Companies are also using algorithms to uncover the bottom-line impact of employee engagement and to drive deeper engagement across the organization. In collections, where retention of talent is a recurring issue, people analytics can be used to find the drivers of performance, including personality profiles and risk factors for low performance and engagement. By identifying individuals most at risk of leaving, for example, banks can take responsive measures to optimize their talent pool for sustained performance improvement.

Machine learning and nontraditional data have become the new frontier in collections-decision support. Audio analytics, for example, is now an important tool for understanding front-line effectiveness. By allowing algorithms to work through thousands of conversations, banks can discover the most productive and engaging approaches. With hypotheses informed by insights from the field of behavioral science, banks are also using machine learning to diagnose and neutralize the biases that affect collector and customer decision making. At the same time, the machine-learning approach is enabling automation of larger classes of decisions. By giving agents more prescriptive decision support in certain situations, including a wider range of set script elements and narrower parameters for negotiations, banks can free capacity and redirect resources toward the most valuable accounts. In this vein, one card issuer achieved dramatic improvements in the rate of promises kept in its high-risk segment by using an approach enabled by data and analytics to script elements, including behavioral insights (Exhibit 3).

Behavioral pairing and agent coaching

Many banks do not apply agent–customer pairing uniformly or deliberately. When it is

Exhibit 3

One card issuer systematically identified and implemented "assertive" script elements, doubling the promise-kept rate to 95 percent.

Diagnostic and hypothesis generation	Validation	Implementation
Hypotheses on psychological levers that could increase promise-kept rate derived from call shadowing, collector debriefing, and customer interviews Eight hypotheses identified for testing	Impact of new psychological levers tested with 200 credit card collection calls Top levers selected for roll-out	Levers embedded in a new script for a high-risk segment with a low promise-kept rate Ongoing coaching of collectors

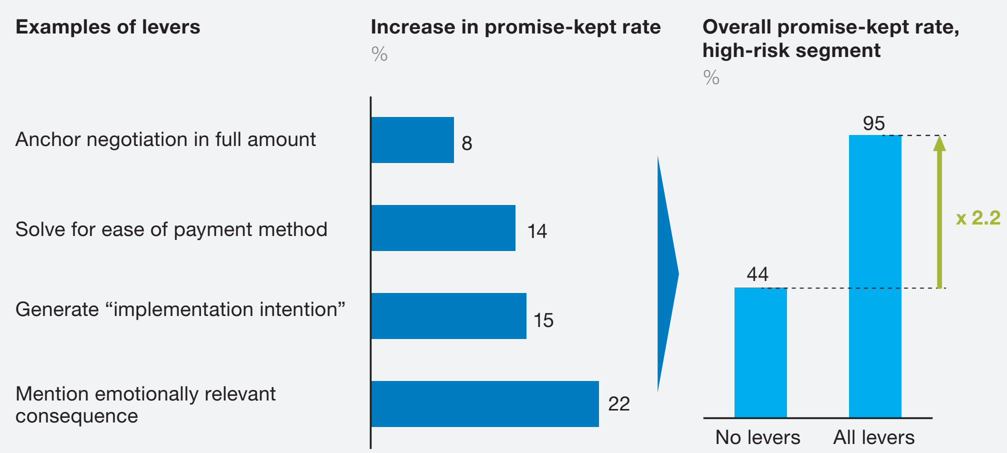

Source: McKinsey analysis

applied, high-risk customers are usually given to experienced, high-performing collectors, while low-risk customers are assigned to new collectors. Analytics-aided pairing helps match collectors and customers who have similar personal profiles. By smarter pairing—matching delinquent clients with the agent expected to be most effective—outcomes can be improved and call times reduced. As for coaching, this has often occurred in training sessions, huddles, and call monitoring by managers. Analytics-aided coaching permits real-time feedback and analysis in live phone calls.

Breaking through artificial barriers to transformation

Most banks understand that analytics and digital automation will transform their collections operations. Some have been reluctant to get started, however, due to the following persisting myths about the new technologies.

- *"Sophisticated data infrastructure is a prerequisite."* While this is an advantage, it can be developed over time. The truth is that banks can build value-enhancing collections models with available data. As the data are improved, the models can be updated accordingly.
- *"Both the collections front line and the digital infrastructure need to be in place before analytics models can be implemented."* Actually, models can be implemented using legacy infrastructure, and the value they generate can be used to invest in the needed infrastructure improvements.
- *"The development and implementation of models take a long time."* Banks can get started using agile model development with minimum viable products subject to continuous improvement. Without rapid iteration and deployment of models, value is left unrealized.
- *"Given compliance and regulatory issues, models are too opaque to use."* Banks can select among a range of modeling techniques with different levels of transparency. They can balance demands for transparency and performance by choosing the most appropriate algorithms.
- *"Success depends on nontraditional data."* For most collections applications, banks' internal data can provide the majority of the gains from advanced analytics. Banks can begin by utilizing all internal data and supplement with external data subsequently as needed.
- *"Regulations and compliance negate many of the benefits of advanced analytics and machine learning."* A number of banks in highly regulated jurisdictions have already successfully deployed machine learning. Indeed, machine learning can improve compliance by better matching the right treatment with the right customer and avoiding biases.

None of these myths should prevent banks from beginning the analytics-enabled transformation of their collections operations. There is no perfect way to start a transformation—some of the implementation might even be messy at first. The essentials of the analytics transformation in collections are clear, however. First, set a long-term vision but also a path toward it that generates value continuously. Second, work in an agile manner, with teams from all dimensions of the transformation. Focus on implementing working models from day one, avoiding an overly complex academic approach. Use synthetic variables to enhance model performance, and continuously experiment with strategies to generate additional data for the next generation of models.

* * *

The next-generation collections environment will be built around advanced analytics and machine learning. These approaches will help lending institutions meet the new delinquency challenges that market analysts predict are on the horizon. The transformation of collections has in fact already begun, as leading institutions assemble the data and develop algorithms to attain improvements in their existing collections context within a few months' time. These leaders are showing the way by applying the new approaches and making improvements as they go. And they are already generating bottom-line results.

Ignacio Crespo is a partner in McKinsey's Madrid office, and **Arvind Govindarajan** is a partner in the Boston office.

In the next issue of McKinsey on Payments*, we will explore how increasing digitization and advanced analytics are reshaping the collections discipline in ways that significantly improve overall effectiveness, efficiency, and the customer experience.*

How machine learning can improve pricing performance

Obtaining fair compensation for complex payments products, such as corporate cards, merchant acquisition, and treasury management services, has long been a major challenge. This is primarily because these products tend to be complex, offered in myriad forms, and implemented across diverse markets. Treasury services, for instance, might have 1,000 or more different fees, and prices are often embedded in private contracts not shared within the organization. Throughout the payments industry, these problems are further complicated by ever-changing payments methods and platforms created by the rapid evolution of payments technologies. And now, expectations of rising interest rates are compounding the situation, increasing uncertainty in product pricing performance for both the short and long term.

Walter Rizzi
Z. Maria Wang
Kuba Zielinski

However, there is also good news on the technology front. Just as technological advances are reshaping the payments landscape, they are also delivering powerful new analytical capabilities that have the potential to transform the way banks and other payments providers price products and services. In fact, early users are already reporting reduced volume loss and customer attrition rates attributable to their use of advanced analytics.

Mining diverse data sets for deeper insight

The complex nature of financial services presents substantial hurdles to those charged with pricing strategically. Merchant acquiring, corporate cards, and treasury services, for example, often include hundreds of products, each with their own distinct fees. Service contracts also differ, and might begin and end at different times. Moreover, prices tend

How effective is your current pricing strategy?

- When was the last time the company reviewed its current pricing strategy?
- Is the full market value of its products and services being captured?
- Does present pricing cover the marginal cost of each transaction?
- What internal and external data sources are being called upon to understand customer and prospect needs and behaviors?
- Does the company make full use of new technologies to mine deep customer and prospect insights from diverse data sources?
- Is the company applying appropriate disciplines and interactive tools to translate newfound insights into actionable strategies?

Exhibit 1

Treasury repricing programs are nearly as likely to be value destructive as value additive.

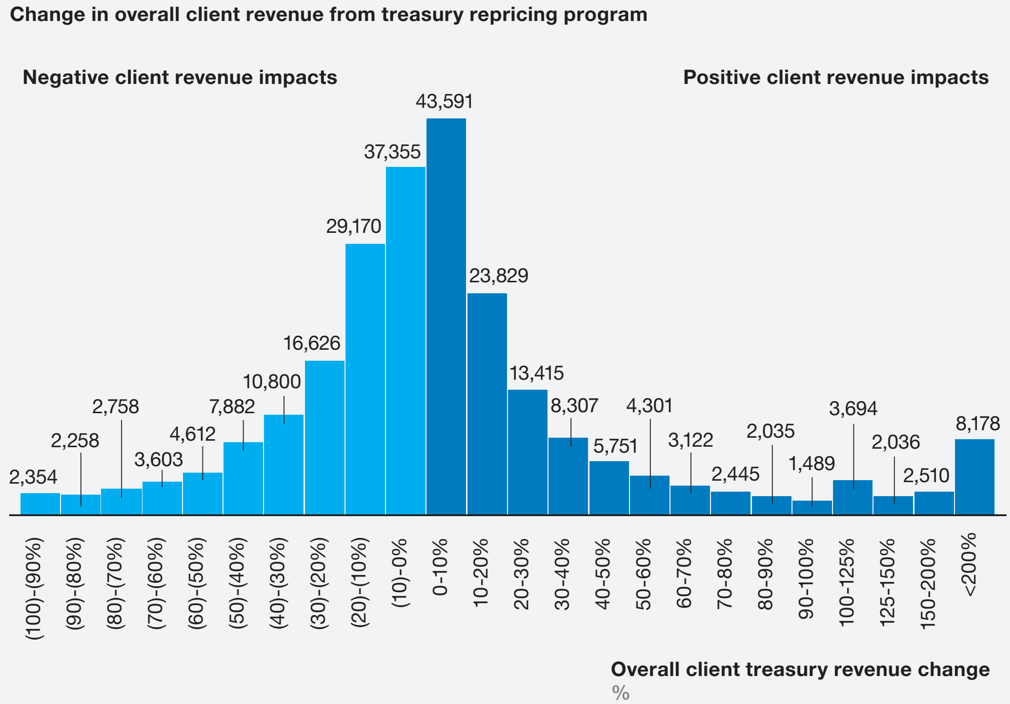

Source: GCInsights; McKinsey Analytics

to be set within the context of the respective client relationship and transparency within the industry and individual institutions is frequently rare or nonexistent. In light of such complexities, most payments providers struggle to systematically devise fair and effective price strategies. For instance, a recent McKinsey study of treasury management services in North America suggests that, over the long term, re-pricing services leads to value destruction about as frequently as it does to value creation (Exhibit 1). In the study, price increases resulted in revenue declines a year later at more than half of the subject institutions, suggesting the outcome of pricing adjustments is highly unpredictable.

New solutions are emerging to cope with pricing complexity. Developments in computational technology, data engineering, and digitization of general processes can now transform how banks and other payments providers create and implement pricing structures. Rapidly declining costs in high-performance computing and data storage, for example, are enabling them to use larger and more diverse data sets to build more sophisticated analytical pricing models. Unsurprisingly, several industry leaders are already capitalizing on the benefits of these developments.

An especially useful new tool has been SparkBeyond. This application can automate feature engineering by creating a wide range of variable transformations, and is highly efficient in identifying the most effective machine learning algorithms, such as random forest or XGBoost. The application also enables users to export selected algorithms and features for out-of-sample testing and other modeling needs in an external environment.

Banks that adopted advanced analytics early on have been building massive data sets that integrate customer and prospect details drawn from diverse internal and external data sources. The resulting content-rich data sets are yielding deeper customer and market insights that are unobtainable using traditional data. For instance, government-published econometric data can yield economic wellbeing information and thereby better guide a bank's budgeting process. And adding commercial and benchmark data can help banks more accurately determine their current business share in large corporate relationships. To obtain richer and more actionable insight at a granular level, some institutions are adopting a variety of advanced technological capabilities, including machine learning, deep learning, and more generally, artificial intelligence (AI). AI uses algorithms that range from unsupervised (such as clustering and principal component analysis) to supervised (such as random forest and neural networks) to reinforcement learning.

Some payments leaders are also venturing into interactive digital pricing, either by subscribing to third-party services or building their own digital pricing tools. Using new data sources, technologies, and modeling techniques, these early adopters are providing front-end staff with in-depth views of customers and prospects, including such information as their product acceptance probabilities, price sensitivities, propensity to churn, and life-time value. These new insights allow management to identify micro-market segments, and thus target pricing more narrowly—down to the individual customer level when data permits. Closely attuning pricing to customer and prospect needs maximizes price performance while minimizing customer attrition and volume loss (Exhibit 2, page 16).

Developing a holistic multi-step approach

Although new tools and capabilities bring opportunities for payments providers to significantly enhance their pricing performance, real success will only come through systematic and comprehensive execution. Achieving maximum effectiveness requires an enterprise-wide pricing transformation. One way to initiate a pricing transformation is to develop incremental price changes with selected markets or segments using pilot programs that can be quickly learned and iterated before rolling out a broad pricing program. Once the proof-of-concept is established, the full program can then be deployed through a three-step approach that optionally can use early revenue gains to fund subsequent steps.

Step 1: Early transformation success tends to include, but not necessarily be limited to, the following initiatives:

- Using advanced analytics technologies, such as machine learning, to establish pricing benchmarks at a granular level. For example, banks can drill down from traditional segmentation levels (such as geography, industry, and deal size) to postal code or business unit, and thus more quickly identify fee leakage at the level of individual customer pricing.

Exhibit 2

An interactive tool can give relationship managers precise price recommendations, helping to minimize client attrition.

	Current price	GCI benchmark price	Target price	Volume
	Average of Current Price	Average of GCI bench price	Average of Target price	Sum of Volume
Credit			9.57%	31,959,115
Deposits	0.48%		0.32%	242,392
TM	8.01	6.26	8.21	259788
	15.00	30.00	15.00	6
	0.07	0.19	0.08	46817
	0.15	0.16	0.19	6681
	0.00		0.00	165673
	0.03	0.10	0.03	6550
	0.13	0.10	0.13	1223
	2.50	3.00	3.13	13
	0.50	0.60	0.63	15
	0.50		0.63	332
	2.00	1.10	2.00	8
	3.25	3.25	3.25	30
	7.00	6.55	7.00	4
	0.15	0.11	0.15	1
	0.08	0.11	0.09	164
	0.05		0.05	165
	0.07		0.09	45
	2.00		2.00	18
	8.50		8.50	28
	5.00		5.00	6
	25.00	24.37	25.00	12
	0.18	0.19	0.18	251
	0.59	1.39	0.74	2446
	5.00	5.00	5.00	12
	-		-	12
	-	0.30	-	20
	80.00	29.33	80.00	12
	70.00	25.28	70.00	12
	0.04	0.14	0.05	17154
	0.04	0.18	0.05	2399
	20.00		25.00	302
	0.50		0.63	9378
	7.76	6.26	7.72	32461295

Cash vault coin order per roll

Current price	GCI price	Internal benchmark	Target price	Potential
$0.13	$0.10	$0.13	$0.13	**$0**

Client priced at internal benchmark for this quartile

Transit check deposited

Current price	GCI price	Internal benchmark	Target price	Potential
$0.08	$0.11	$0.09	$0.09	**$3**

Client priced below both GCI and internal benchmarks, price increase up to internal benchmark

DDA deposit

Current price	GCI price	Internal benchmark	Target price	Potential
$0.59	$1.39	$1.30	$0.74	**$739**

Client priced below both GCI and internal benchmarks, price increase capped at 25% (still below internal benchmark)

Online information reporting previous day

Current price	GCI price	Internal benchmark	Target price	Potential
$0.04	$0.14	$0.05	$0.05	**$352**

Client priced below both GCI and internal benchmarks, price increase up to internal benchmark

Source: GCInsights; McKinsey Analytics

- In parallel with the above, developing interactive tools that enable field representatives to rapidly recognize pricing opportunities in client portfolios, and to simultaneously leverage other opportunities to expand share of business with the customer.

- Initiating price discussions throughout the organization and redesigning the pricing process so it can be implemented in carefully timed waves. A key component of this is building a disciplined exception-management process to strengthen pricing governance, and to identify and remedy flaws in current processes and policies.

Initial programs incorporating these three elements can yield revenue lifts of about 15 percent within six to nine months, yet incur only minimal client and volume attrition rates. And those who implement rigorous service re-pricing programs can apply early revenue gains to funding the overall journey.

Step 2: Begin developing new organization-wide pricing skills and capabilities. While this often becomes a longer-term journey, it is one that needs to be initiated and proactively

Pursuing insights across the payments landscape

To better track the global payments industry, McKinsey frequently turns to GCInsights, its wholly owned research subsidiary. Since 2014, GCInsights has been acquiring and analyzing monthly or quarterly data from approximately 30 North American banks that provide commercial card and treasury management services. Its extensive data pool provides transaction-level details, such as price and volume information for approximately 5 million pricing events across more than 250 service codes. These codes include treasury management (such as DDA deposit) and commercial card (such as travel and entertainment purchases). Some interchange details, such as a merchant's industry, region, and institutional size, are also included. Finalta, another McKinsey tool, provides global benchmarking services for the retail banking and insurance sectors. It covers approximately 300 banks and insurers in more than 55 countries.

managed from an early stage. Step 2 commonly includes:

- Improving and expanding skill sets throughout the pricing organization
- Significantly enhancing current pricing data sets
- Building strong pricing analytics capabilities
- Developing enterprise-grade tools to assist in such key areas as new-deal pricing, contract renewal pricing, and ongoing revenue portfolio management

Common related investments include acquiring new technology capabilities, such as voice recognition and automation, to reduce manual processing, human error rates, and technical leakage.

Step 3: Lastly, equipped with powerful analytical tools, immense data sets, and newfound skills, banks can continually enhance their pricing strategies through ongoing monitoring and scaling of new pricing constructs. Together, these actions are helping many institutions to improve the ways they address current and prospective client needs in diverse markets and segments. Pricing tactics, for instance, can be finely tuned to reflect evolving customer and prospect needs by drawing on a variety of pricing approaches, including bundled pricing, subscription pricing for small businesses, and unbundled granular pricing for corporate clients. Consequently, deeper understanding of ever-changing marketplace needs can provide a clear competitive advantage.

For example, when helping a global payments-network provider to develop a new pricing strategy, McKinsey heavily adopted machine learning while applying this three-step approach. Drawing on the large amount of transactional data from the last few years, McKinsey designed a new pricing construct, and subsequently simulated its implementation to determine the probable impact on both revenue and attrition. Simulation played a central role in forming the new pricing strategy.

In another case, a global merchant services company needed to more closely match its established pricing with the current value of its product and service offerings. Using gradient boosting machine (a machine-learning technique) proved to be highly effective in identifying and realigning numerous mismatched price-value occurrences while accounting for the current competitive environment.

Notably, this framework is expandable to fit a wide range of pricing scenarios. In small-to-medium enterprise lending, for instance, McKinsey worked with UniCredit to devise a value-based account management strategy; a core component of the effort was developing machine-learning models and a custom user interface that generates client benchmark profits. The customized solution enables relationship managers to simulate various deal scenarios at the full client relationship level in a test-and-learn environment to ensure continuously monitored and improved results. These changes resulted in value increase of up to 15 percent.

Leaping the hurdles of price transformation

Adopting machine learning and advanced analytics generally gives payments providers significant power to reshape their longstanding pricing strategies, but transformation can also present unique challenges.

Advanced analytics presents a variety of sophisticated tools, but their effectiveness depends largely on how the insights are actually derived and subsequently used. For example, traditional approaches to setting pricing targets, such as scoring or ranking customers on price sensitivity, are less actionable than employing a mathematical model that links offer acceptance probability to historically accepted offer rates. Pricing models based solely on statistical performance can deliver suboptimal guidance: maximum performance, by contrast, also requires the application of sound business principles and disciplined practices.

Aside from aspects of data and analytics modeling, another common obstacle to achieving full effectiveness in the use of advanced analytics is a siloed organizational structure. Organizational silos often lead to departmental misalignments—say between finance, marketing, and sales—when making strategic pricing decisions. In these situations, the best practice is usually to ensure from the outset that all stakeholders have integral roles in planning and implementing the pricing transformation, and participate regularly in transformation planning and progress review meetings.

To generate positive results even the best of strategies require seamless execution. Real or perceived flaws during the roll-out of a new pricing can quickly incite rejection among relationship mangers—a problem that successful institutions are overcoming by showcasing success stories and prominently recognizing champions of change within their organizations. Of course, it is also essential to promptly realign performance incentives with new pricing approaches and goals. Engaging relationship managers in codeveloping pricing strategies is a highly effective way to generate positive change attitudes. To instill manager confidence in a new pricing approach, one European bank devised algorithms that can instantly show managers the bank’s pricing structure on comparable deals.

Advanced analytics technologies are beginning to rapidly alter how businesses operate around the globe. Given the central role of banks and other payments industry participants, they are fast becoming subject to those same forces, and to remain competitive will therefore need to embrace them on a timely basis. Many in the payments industry might be hesitant to change their longstanding approaches to pricing, but those willing to adopt a comprehensive pricing transformation built on deep market insights will clearly be among tomorrow's industry leaders.

Walter Rizzi is a partner in McKinsey's Milan office. **Z. Maria Wang** is an analytics expert and **Kuba Zielinski** is a partner, both in the Boston office.

Robust analytical tools are available to help banks increase revenue at each phase of the customer life cycle. The next issue of McKinsey on Payments *will include an article examining how treasury management service providers are using analytics to optimize portfolio value, with special attention to product structure and pricing, cross-selling, and managing attrition.*

Combating payments fraud and enhancing customer experience

The fraud threat facing banks and payments firms has grown dramatically in recent years (Exhibit 1). Estimates of fraud's impact on consumers and financial institutions vary significantly but losses to banks alone are conservatively estimated to exceed $31 billion globally by 2018. Several converging trends have propelled the increasing scale, diversity, and complexity of fraud. Vulnerabilities in payments services have increased as the shift to digital and mobile customer platforms accelerates. New solutions have also led to payments transactions being executed more quickly, leaving banks and processors with less time to identify, counteract, and recover the underlying funds when necessary. Finally, the sophistication of fraud has increased, in part through greater collaboration among bad actors, including the exchange of stolen data, new techniques, and expertise on the dark web.

Salim Hasham
Rob Hayden
Rob Wavra

Increasingly agile fraud perpetrators have benefited from banks' and payments firms' limited ability to adapt. While most institutions have well-funded anti-fraud groups, key resources are often fragmented across the organization. Essential data, investigative and forensics expertise, and analytics talent are typically distributed across cyber, compliance, legal, IT, and fraud teams, with little to no coordination or data sharing.

Effectively combating fraud through analytics requires a mindset shift from a narrow focus on false positives and loss prevention to an appreciation that the same technological advancements making fraud more pervasive also enable the tools and environment to address it. With their shift to digital services, banks have access to exponentially more customer and transaction data than in the past. New technologies create the means to more accurately segment customers by risk, enabling lower-friction digital experiences—and higher satisfaction levels—for low-risk customers. And the explosion of industry verticals in cyber and data analytics has created a ready supply of talented, cross-disciplinary resources unencumbered by legacy organizational structures. Today's challenge is harnessing these components to reduce current losses, detect and prevent emerging fraud, and enhance customer experience.

The shifting fraud landscape

Fraud is not only growing but evolving (Exhibit 2, page 22), forcing countermeasures to shift from the transaction-centric assessment of fraudulent charges on a card or doctored checks deposited at an ATM, to preventing, detecting, and remediating increasingly sophisticated, long-term sleeper frauds and exotic concerns like manipulated synthetic identities. Some tactics have worked, with Visa estimating that chip technology reduced counterfeit card fraud in the US by 66 percent for EMV-enabled merchants in June 2017 compared to June 2015. Other typologies ("abuse cases") of fraud remain without effective countermeasures, straining traditional anti-fraud efforts, generating increasing losses, false positives, and negative customer experiences:

- **Account takeover** (ATO) is the theft or misuse of credentials to fraudulently gain access to an existing customer account. This can be a one-time funds transfer event or an ongoing access exploitation (e.g., adding a registered user, changing the contact email or mailing address) for criminal purposes. Successfully combatting ATO requires a mix of nontraditional data sources that expand customer identification beyond knowledge-based authentication (KBA), analytics to detect emerging trends and high-risk access events, and customer experience-sensitive authentication journeys, limiting customer challenges based on risk segmentation and other triggers.

- **Synthetic identity,** a scenario in which fraud perpetrators combine fragments of stolen or fake information to create a new identity and apply for financial products, is

Exhibit 1

Fraud is on the increase in the US.

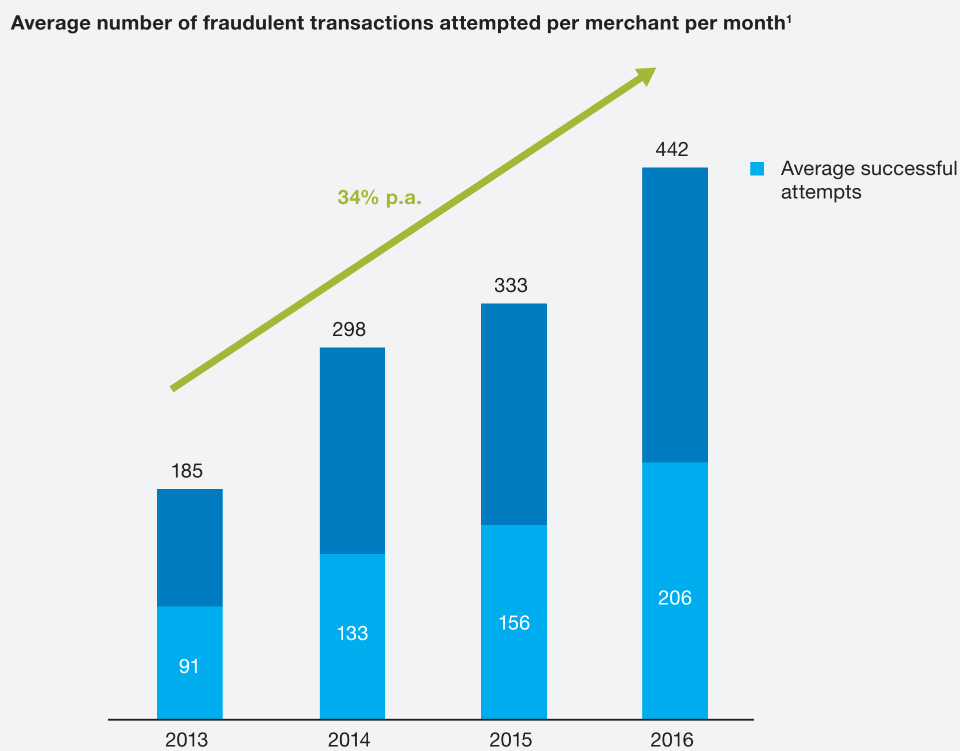

[1] Weighted merchant responses to LexisNexis survey question: In a typical month, approximately how many fraudulent transactions are prevented by your company / successfully completed by fraudsters? What is the average value of successful fraud transactions?

Source: US Department of Commerce; LexisNexis The True Cost of Fraud study, 2016

Exhibit 2

Payments transaction fraud takes many forms.

Fraud types/threats		Key examples	Trends impacting this fraud type
Account fraud (new or takeover)		Account creation using false or stolen identity Fraudulent **access to an existing account** (e.g., adding a registered user, changing email or mailing address)	Increasing **sophistication of tools** used to establish identity (e.g., IP/address geolocation, record matching algorithms)
Transaction/ payments fraud	Identity theft	Transactions using **stolen cards/accounts** **Suspicious transactions** (e.g., geography, counterparties)	Increases in **card-not-present fraud** as EMV is rolled out New **authentication measures** (e.g., biometrics) and **algorithms** being rolled out
	Chargeback fraud	Existing customer **fraudulently disputes a charge** (e.g., denies delivery while retaining the goods)	**Improved tracking** methods for delivery
	Friendly fraud	**Unintentional chargeback fraud** (e.g., forgetfulness, accidental order, not recognizing merchant name, misunderstanding return policy, family members ordering)	**More automation**, less manual customer service Increased **'on-demand' ecommerce**

■ Increasing threats

Source: McKinsey analysis

of growing concern in light of data exposed through the 2017 Equifax breach. All of synthetic identity fraud's forms—traditional (a fusion of valid information from multiple real people), manipulated (all real information about a single person with a fake national ID/SSN), and manufactured (wholly fake information, including national ID/SSN)—can exist only because of inadequate onboarding and customer due diligence. Filling these gaps will require cross-functional collaboration across lines of business and functional silos, expanded external data for validating multiple elements of customer applications and scoring their likelihood of authenticity, and the fusion of these external sources and existing internal customer data.

- **Business email compromise** invokes social engineering to lure an empowered employee to initiate a transfer to the fraudster's account, usually at the apparent request of an executive. A similar phenomenon, *invoice redirection*, leverages social engineering to alter payment information for legitimate payables accounts (often by claiming a new bank account has been opened), redirecting payment to a fraudster's account. These are growing fraud categories—business email compromise alone causes nearly $1.5 billion per year in losses according to the FBI—demanding institutions respond with tailored frontline training, re-architecture of existing controls (e.g., who can change payment information and based on what information, verification of new information with a known contact), sophisticated analytics to flag risky changes before payments are made, and new data and technologies like voice analytics.

In "Fraud management: Recovering value through next-generation solutions" (*McKinsey on Payments,* June 2018), our colleagues identified three concrete steps to effectively redefine fraud operating models to fight these emerging threats: re-engineering fraud case management; redesigning journeys to improve the customer experience; and employing advanced analytics. Given the vast potential of advanced analytics in the fraud arena and the significant barriers to its effective use, we will focus on this critical dimension. When done well, analytics can consistently reduce fraud losses by 3 to 5 percent in mature environments and by over 30 percent in evolving contexts. And yet we have seen even the most advanced firms struggle to attract and maintain analytics talent, transcend organizational and disciplinary boundaries to deploy the best solutions, and transition from analytics test cases to production capabilities.

False starts

In the face of the continuous evolution and increasing pace and volume of fraud threats, fraud teams find themselves hamstrung by ineffective triage of alerts, poor data quality, and non-existent or outdated intelligence. Compounding this fragmentation, many investments in fraud-related artificial intelligence can be characterized as "science projects," lacking the scale to deliver enterprise impact. In the meantime, institutions are dedicating additional resources to manually wade through low-value alerts or building increasingly aggressive rules and models that often hurt customer experience more than they mitigate fraud.

Fraud interventions driven by advanced analytics tend to follow a few archetypes:

- *Predictive detection, encompassing user authentication* (e.g., determining whether the transacting party is in fact a customer), customer due diligence (e.g., low/high-risk fraud profiling as a factor in exception decisioning), and transaction risk (e.g., whether hallmarks of fraud are present in the context of other transactions for the account, customer, and household). This can come in the form of in-house custom analytics models, commercial off-the-shelf software-enabled detection, or public partnerships with emerging technology companies, like HSBC's relationship with Ayasdi.
- *Enhanced internal process efficiency,* such as capacity forecasting and providing analysts with context detailing the reasons a transaction failed an initial screen.
- *Automated fraud triage and other robotic process automation (RPA).* The London School of Economics examined 16 case studies of RPA, finding first-year returns on investment of 30 to 200 percent. The longer-term value—including enhanced compliance and the reallocation of employees to higher-value tasks—is likely even greater.

Many banks, however, have faced serious challenges when attempting to effectively integrate advanced analytics into their fraud defense. Common pitfalls include:

- Building models that do not take advantage of all available data, overlooking siloed risk scoring inputs residing in cyber, customer relationship and product sales groups.

Such inputs can be as simple as determining whether cross-ownership of mortgage or card products correlates to lower fraud risk or exploiting device geolocation data to inform mobile deposit fraud screening—which enabled a US bank to identify deposits typologies with higher fraud incidence of 25 to 1,000 times. More ambitious enhancements include holistic realignment of a bank's financial crime structures, people, and technology, as undertaken by HSBC in 2015 with its creation of a unified Financial Crime Threat Mitigation organization.

- Deploying "crime- and institution-ignorant" models, which are statistically compelling but hobbled by a lack of understanding of underlying fraud mechanisms, institutional controls, and intervention options. While staffing fraud analytics efforts with cross-disciplinary teams of data scientists, data engineers, translators, and financial crime and fraud subject matter experts is a powerful solution, Citigroup went one step further, empowering a permanent Global Investigations Unit to proactively analyze and combat emerging financial crimes with a full range of experts and technical staff.
- Not addressing the growing model risk management (MRM) demands in fraud mitigation. The increasingly opaque and sophisticated models used to detect fraud and the rapid pace at which fraud is evolving combine to create model risk. Some causes are easily addressable—assumptions about the markers of fraud and the scale of potential losses can become strained—but others stem from well-meaning attempts to use cutting-edge deep learning and neural network algorithms which are difficult, if not impossible, to interpret. Techniques like Locally Interpretable Model-Agnostic Explanations (LIME) provide some insights into sophisticated models, but do not mitigate the increased model risk that the push for performance and innovation has created.
- Not accounting for the increasing interest of regulators in fraud models. This scrutiny is likely to accelerate, given the opaque nature of fraud rules and concern over whether they impose disparate impact on members of a protected class. Loss ratios and raw statistical performance cannot be the only metrics by which modern fraud models are measured.
- Grafting advanced analytics tools onto existing processes and policy frameworks rather than leveraging analytics to transform the business. Analytics should not be deployed merely to dig out of a false positive hole created by bad policies and inefficient processes. While many frauds are driven by control weaknesses, fast-growing threats like synthetic identity fraud exist only because of insufficient onboarding processes and customer due diligence at the application stage. Using advanced analytics to detect these frauds or reduce false positives being generated misses the real opportunity to fix outmoded policies and underperforming processes.

The best analytics interventions leverage cross-disciplinary expertise, fusing analytics with deep industry and client organization context. At a regional bank in the United States, the breakthrough came from shifting its focus from identifying fraudulent transactions to minimizing dollar losses from a specific fraud typology. Pairing this approach with risk-

Accelerating analytics-driven fraud defenses

"Money mule" accounts are often recruited via unwitting accomplices (e.g., through work-from-home schemes) and exploited to launder illicit funds, rapidly moving sums through multiple accounts to obfuscate sources and frustrate identification and repatriation efforts. Advanced network analytics and machine-learning techniques can discern patterns in the noise, exposing suspicious accounts with impressive efficacy. For instance, QuantumBlack, a McKinsey company focused on advanced analytics, analyzed over 18 billion transactions across multiple banks, creating a "mule-inesque" score integrating indicators of mule activity (e.g., account age, economic relationships, direct debit frequency). QuantumBlack analyzed over 10,000 suspected criminal account networks through an investigator analytics support tool, visually tracing dispersion networks to allow for real-time detection and timely repatriation. The exercise ultimately identified 15,000 mule accounts across multiple banks.

Although signature fraud has been a common tactic for generations, it has taken on new dimensions in certain markets. A bank in Latin America was overwhelmed by both traditional loan application fraud (e.g., for recently deceased relatives) and "auto-fraud," where an applicant intentionally modifies their own signature with the intention of later claiming not to have initiated the loan. Using deep learning-based image analytics techniques, McKinsey identified the subtle indicators of both types of fraudulent signatures. The new model improved fraud detection by over 31 percent when compared to the bank's existing model.

driven policy changes and data science-driven enhancements to tune their detection model, the bank was able to create a combination of model enhancements and policy change efforts projected to reduce annual losses in the target category by over 32 percent.

Succeeding in fraud analytics

Effectively deploying analytics to combat fraud requires a shift in thinking from a narrow focus on false positives and losses to an appreciation that the same trends making fraud more pervasive also enable the tools and environment necessary to combat it. With their shift to digital services, banks have exponentially more customer and transaction data than in the past. New technologies also create the means to more accurately segment customers by risk, enabling lower-friction digital experiences—and higher satisfaction levels—for low-risk customers. Many of the technological advances that have sped the pace of payments can also be leveraged to increase the speed and efficiency of anti-fraud processes. And the explosion of the cyber and data analytics verticals has created a ready supply of talented, cross-disciplinary resources unencumbered by legacy organizational structures.

Analytics provide a unique and powerful means to transform fraud operations. The most successful fraud analytics programs are designed to be:

- **Business-back:** Anti-fraud analytics efforts must be built on a unified, cross enterprise foundation, breaking down silos between channels, products, and fraud types. This is usually best accomplished with an overarching fraud operations transformation mandate from senior management, transcending analytics. Given the increasing impact of fraud on bottom lines and reputations, the business case to secure such a broad mandate should be fairly straightforward. The goal should be a process seamlessly integrated across the fraud lifecycle, incorporating data spanning business units and functional silos to create a holistic view.

- **Criminal-forward:** Applying a criminal mindset to fraud analytics—a common tactic used by law enforcement agencies—can provide inputs to better understand the motivations and methods of perpetrators of fraud. From this starting point, models can be designed to predict, prevent, and detect crime based on powerful data-driven insights and expert-created indicators created from more nuanced and comprehensive understanding of the criminal. By mapping typologies to indicators of fraud, analytics can be better targeted and prioritized. Such a focus requires more than just fraud experts and data scientists; it demands a rigorous, evidence-based method to testing expert hypotheses with large data sets on past fraud and a culture that embraces the power of such a hybrid approach.

- **Intelligence-driven:** Rather than building models that chase historical fraud threats after the fact, banks must continuously evolve their analytics-centered defenses based on detailed up-to-the-minute understanding of the criminal environment. Such knowledge is best developed through intelligence operations and sharing, including monitoring of the dark web. Rather than interrogating fraud incidents in isolation, institutions must take a broader look at the patterns of crime. Industry-wide objectives such as FS-ISAC in the United States provide a more robust data set from which to identify such patterns. The goal should be to shift risk identification from regulatory rules-based detection and predictive models built on past frauds to forward-looking analytics built on well-founded indicators of crime. This creates the means to spot broader patterns of suspicious behavior—such as campaigns by criminal networks as opposed to lone fraudsters—and to look for emerging fraud typologies before significant losses result.

- **Customer-focused:** While constantly evolving to counter the fraud threat, countermeasures should be designed in ways that create a distinctive customer experience balancing trust and convenience to accelerate insight into fraud. Analytics should play as critical a role in facilitating low-risk customers and transactions as they do in thwarting potential fraud, enabling institutions to create customized, analytics-informed journeys balancing security and convenience. Models must be built on the proper foundation, integrating customer behaviors across accounts and transactions into a single view that enhances the power of prediction and detection.

Cutting-edge efforts integrate these themes, pairing a mandate to improve customer experience with improvements in fraud identification. One global bank undertook such a hybrid effort, redesigning customer authentication journeys in its digital channel to simultaneously improve its confidence in customer identification while dramatically improving experience. Beyond achieving its security-related goals, this effort reduced costs related to customer lock-out by $5 million and improved Net Promoter Scores in the online channel by 29 points.

Getting started

To get the most from advanced analytics, organizations should begin by clearly articulating their operational objectives. This critical foundation provides the proper screens against which to evaluate analytics efforts and investments. It also aligns analytics interventions with business unit goals, identifying the core decisions requiring analytics support, prioritizing those decisions best informed by advanced analytics, mapping data to inform those decisions, designing models leveraging that data, and establishing the metrics against which to evaluate analytics success.

Building from this base, firms should approach advanced analytics as a transformation rather than a one-off event. In the near term, this involves a focus on:

- Identifying the universe of possible interventions, connecting the business with analytics and compliance to prioritize based on potential impact and technical feasibility.
- Articulating clear operational goals, understanding where internal analytics capabilities stand today, where they should be, the investments required, and developing plan to transition from outside support to a reliance on internal resources.
- Cataloguing current capabilities and ensuring they are being leveraged to their maximum potential. Banks often have many of the tools required for an effective initial defense but have not yet aligned them properly.

Individual use cases, pilots, and other traditional means of intervening through analytics should be used to enhance these base capabilities and push the institution's capacity, rather than simply as a means to deliver point solutions. In the medium to long term, organizations must build organic capabilities to constantly reassess evolving fraud threats, revisit and improve the operating model, and design fit-for-purpose advanced analytics and fused data sets.

* * *

The perpetrators of fraud are highly adept at exploiting advances in technology, collaboration, and specialization. Legacy approaches to fraud prevention have not kept pace, with financial institutions stubbornly dependent on siloed data and manual processes. Banks and payments firms looking to establish a competitive edge—and avoid increasing loss exposure and mitigation expense—must harness these same trends. Advanced analytics provide a tangible reason to integrate data across siloes, a means to automate and enhance expert knowledge, and the right tools to prevent, predict, detect, and remediate fraud. Analytics is not an overnight fix, but it can pay immediate benefits while creating the foundation for anti-fraud operating models of the future.

Salim Hasham is a partner in McKinsey's New York office, and **Rob Wavra** is an expert associate partner with McKinsey's QuantumBlack in Boston. **Rob Hayden** was a senior expert in McKinsey's Cleveland office. Rob passed away suddenly and unexpectedly earlier this year, and is deeply missed by all whose lives he touched. Please see McKinsey on Payments, Issue 27 for a remembrance of Rob.

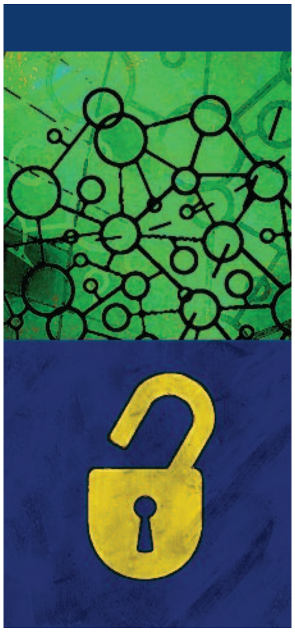

Using data to unlock the potential of an SME and mid-corporate franchise

Banks have long pondered the untapped value of the commercial segment but often lack the means to identify the precise needs of individual companies in this large and diverse population. This is changing, however. By mining huge reserves of customer data, banking analytics leaders are meeting the needs of hundreds of thousands of commercial customers—from small businesses to medium-size corporations—with new levels of convenience and cost efficiency. Several banks have achieved a ten-fold increase in the success rate of product recommendations, thus delivering highly relevant offers with clear economic benefit. This article highlights recent examples of how "next-product-to-buy" (NPtB) recommendation engines are identifying time-critical needs for their small- and medium-size enterprise (SME) and mid-corporate clients.

Ignacio Crespo
Carlos Fernandez
Huw Kwon

The problem: Most businesses are "invisible"

Many banks currently use rules-based models to generate recommendations for SMEs and mid-corporate companies (with annual sales up to $100 million), but with limited success. Relationship managers often view these recommendations with skepticism, as conversion rates typically range between three and five percent. They resort to general propositions designed for the consumer segment and devote most of their energy to those clients whose businesses they already know well and whose needs they can anticipate reliably. The result is that 25 percent of a bank's commercial customers usually account for 85 percent of the revenues, and the remaining 75 percent represents the "long tail" of untapped potential. These companies are effectively invisible to the bank's sales force (Exhibit 1).

The solution: Anticipate the client's next step

Banks are investing in building up predictive models globally: US Bank and TD Bank in North America; Itau and Banco do Brasil in Latin America; Barclays Bank and Lloyds Bank in the UK; ING, Banco Santander, and BBVA in Europe are just some examples of banks improving their commercial performance by leveraging machine learning. These advanced techniques have proven effective in diverse customer segments, from self-employed individuals to large corporate customers. SMEs and mid-corps are the sweet-spot for NPtB, as they generate massive amounts of data, which are typically underused. With the help of advanced analytics decisioning engines, banks have demonstrated that it is now practical to mine vast (and often messy) amounts of data, separating signal from noise, to arrive at precise recommendations for a client's next action. In addition, by broadening the types of data collected for the commercial segment, banks are also analyzing customer behaviors, transactions, and customer preferences across more extensive databases.

Successful implementation of NPtB engines has boosted new sales upwards of 30 percent and increased commercial segment revenues by between two to three percent. The impact on sales efficiency has been radical in some cases, with an increase of more than 50 percent in the number of leads offered per client and as many as six out of ten customers purchasing a new product in response to a sales call.

Leveraging data for NPTB recommendations

More than a decade ago, Amazon and Netflix began leveraging data and analytics to improve their cross-selling efforts. They started with simple analytics, dividing huge customer populations into several dozens of microsegments according to key behaviors (inputs). In order to achieve this new level of precision, they used singular value decomposition (SVD) to classify customers according to patterns in their purchase histories, each pattern culminating in a target output, that is, the "next product to buy." The number of inputs and the complexity of the algorithms used to analyze these inputs have been increasing in recent years, achieving outputs that have much greater precision than was possible with next-best-action (NBA) models. (See sidebar on page 33 for a summary of the evolution of NPtB from NBA.)

The NBA engines employed by Amazon (with more than 300 million customers reported in 2016) and Netflix (125 million subscribers re-

Exhibit 1

Cross-selling efforts often focus on the companies that relationship managers know well.

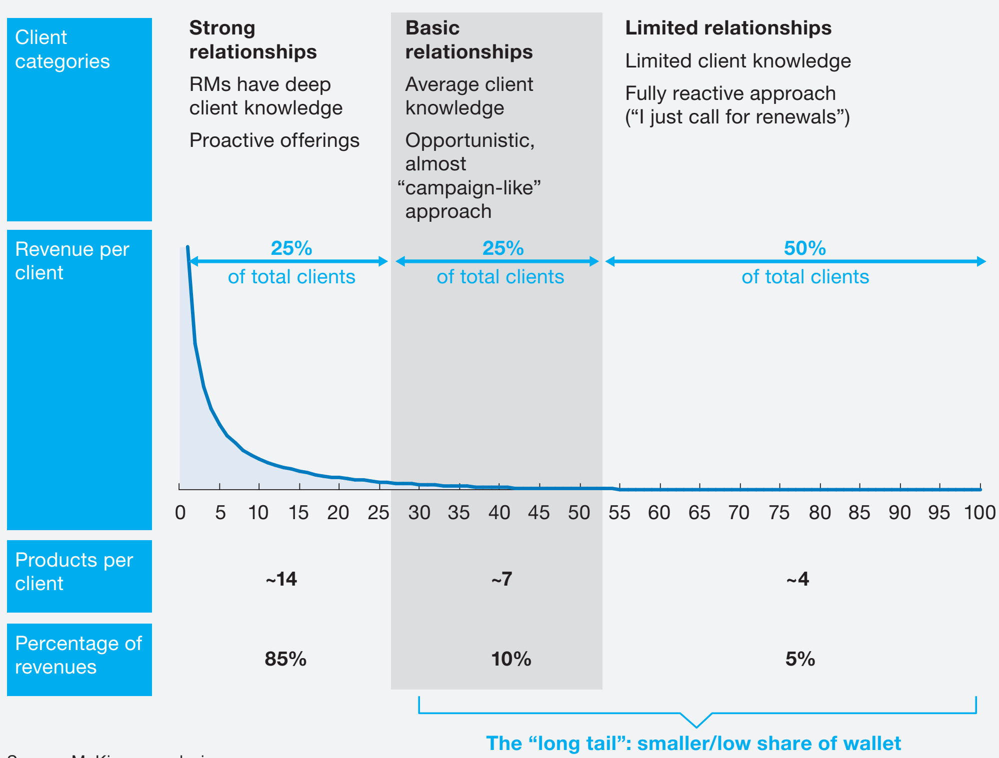

Source: McKinsey analysis

ported in the first quarter of 2018) are not entirely suited to banks serving SMEs and large corporations with products/services that address a relatively narrow range of business activity—domestic and cross-border payments, financing, documentary credit, investments, and insurance. By building NBA recommendation engines designed specifically for transaction banking, banks have increased service levels and profitability, improving their responsiveness to SMEs and helping large corporate clients cut through complex banking relationships and account structures to optimize liquidity.

To maximize the impact of each recommendation, decision engines should identify both customer needs and the preferred channel(s) for delivering the proposal and related communication. In some markets, companies tend to rely more heavily on direct communication with relationship managers, who play a key role in following up on recommendations. In other markets, such as the Nordics and the UK, the digital channel is the primary means both for alerting a customer to a recommended action and for delivering more detailed information about the opportunity.

Consolidate data for analysis in three waves

The data reserves required to power an NPtB engine are consolidated in three waves. As the volume and complexity of data increase across the three waves, analytical algorithms become progressively more sophisticated and accurate in predicting precise, time-critical needs of individual customers.

The first wave starts with the aggregation and analysis of internal structured data of various formats, including customer demographics, product usage, profitability, and transaction history. For example, one bank in Europe started by consolidating the information it had for 1.3 million SME customers, ranging from beauty salons, doctor's offices, and family-owned stores to small manufacturing companies and technology start-ups. This data set yielded 1,200 variables for analysis.

Continuing the focus on internal data, the second wave introduces algorithms capable of digesting unstructured data (e.g., call records, email communication), as well as a broader range of structured data from CRM systems (e.g., share of wallet, historical risk scoring, maturities, customer relationship lifecycle, company value chain, and suppliers). Fast-evolving algorithms augment the value of data already at hand by learning to recognize unanticipated clusters and associations in increasingly complex data sets. The algorithms generate actionable insights into a company's current needs, from payables management to financing for new equipment, based on information coming from transactions and payments along the customer value chain.

The third wave analyzes a broad range of data from point-of-sale transactions to industry news and comments on social media to generate ever more precise recommendations. As machine learning algorithms become more sophisticated, it is possible to produce recommendations from increasingly diverse types of unstructured data (including, voice, image, and video files) extracted from industry and company web sites, as well as news and social media (Exhibit 2).

How to build an NPtB engine: Design, develop, deploy

The preparation of an NPtB model moves through three phases: design, development, and deployment. Clear milestones mark the advance from analytics to proof-of-concept to implementation in the front line.

Design the NPtB engine

This first phase is the preparation and design of the NPtB engine—adjusting the scope, mapping pitfalls, and elaborating the business case to convince stakeholders throughout the organization of the value of the effort. This phase has five key activities:

1. **Prepare the data.** Looking across the three waves of data consolidation, the first step is to identify the types and sources of data and map the variables that can be analyzed. The goal is to locate data that can be combined with unique customer identifiers and provide sufficient record history (at least two years of data). It is also necessary during this phase to check data quality (that is, data consistency, format validation) to identify points for improvement and to restructure data ingestion (for example, defining treatment of unavailable values and primary keys, synchronizing time periods).

2. **Ensure the IT infrastructure meets the processing requirements to run the model.** It is important to design the recommendation engine so that it runs efficiently on the current IT infrastructure. Most banks have adequate processing and storage capacity to run basic NBA models that will generate actionable recommendations. For example, a CPU system with expanded memory capacity is adequate to run induction algorithms based on decision trees. However, deep learning algo-

Exhibit 2

Data from diverse sources are consolidated in successive waves of increasing complexity.

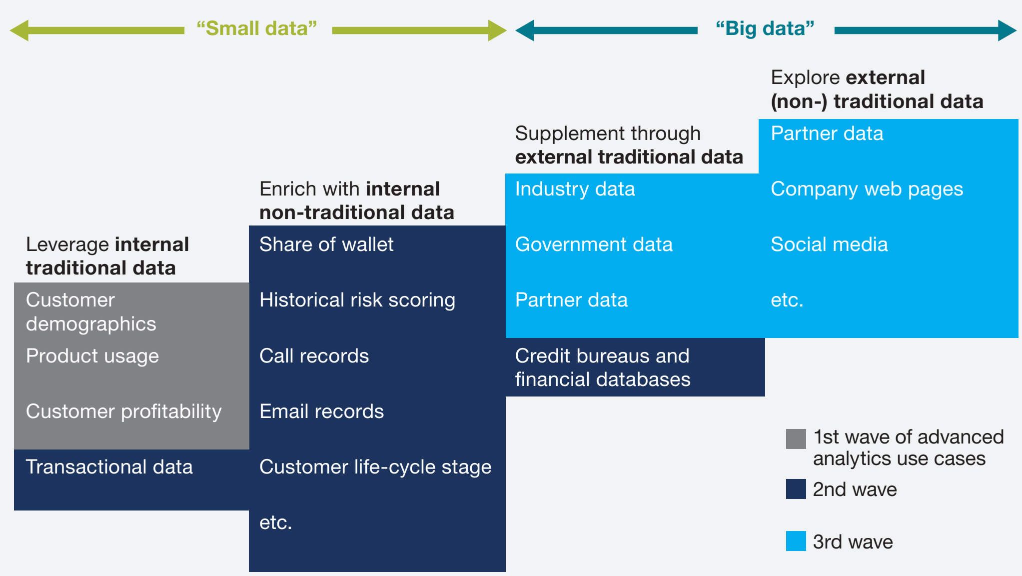

Source: McKinsey analysis

rithms require a GPU system to support complex artificial neural networks.

3. **Identify business sponsor and form multidisciplinary team.** The sponsor should be a business-line owner authorized to make binding decisions. The team should include data scientists, data engineers, business translators, UX designers, and data architects. The team identifies the variables to be analyzed and builds the analytical model based on precise understanding of business goals, including familiarity with the market segment to be addressed and the products to be targeted.

4. **Scope the effort and build the business case.** The executive team must reach a shared understanding of the problem to be solved and agree on a "back of the envelope" business case for the NPtB effort. It should also identify the main elements for evaluation (model performance, adoption rate, conversion rate improvement, etc.).

5. **Identify potential roadblocks.** It is crucial to follow proper legal and compliance procedures to ensure that the bank has the necessary consent/permission to use and merge the targeted data. In our experience, most of the data available for the corporate segment can be leveraged for NPtB analysis; however, it is important to identify early in the design phase any business limitations that may delay implementation. Such limitations may include, for example, business involvement, change management challenges, and workers' council policies.

Develop the analytical solution

In the development phase, data scientists, data engineers, and business translators collaborate to build the analytical solution of the NPtB model. The goal is to identify the products a customer is likely to buy, prioritize recommendations, and determine the most effective channels for delivering an offer. The team works toward these goals by building algorithms to answer three main questions:

1. **Which products does a particular company need or is willing to acquire?** The NPtB analytical engine identifies opportunities for cross-selling. For each specific company doing business with the bank, the engine ranks commercial leads for each product according to two criteria: probability (which leads are most likely to result in a transaction?) and value (which will be most profitable?).

2. **Which companies need a particular product?** Next, the engine also prioritizes clients according to their potential value/business priorities, propensity to buy, and more. (This step is particularly helpful for relationship managers, who must decide how to prioritize follow-up calls and visits.)

3. **Which channels should be used to optimize the success rate of the commercial opportunities?** Leads are distributed to digital and traditional channels based on company behavior and preferences, contact policies, and relationship managers' commercial activities.

To answer these three questions precisely, banks can analyze customer data, such as payments transactions and digital interactions. Machine learning algorithms can identify patterns in past customer behavior to predict future customer purchases. Data scientists build NPtB engines leveraging modeling environments such as R, Python, or Spark.

From NBA to NPtB

Next-product-to-buy (NPtB) models represent a significant advance over those using first- and second-generation next-best-action (NBA) methodologies. In the past decade, product recommendations for corporate transaction banking were generated through statistical analysis of historically observed behaviors across diverse sub-segments (usually between 10 and 20 in number). Each recommendation called attention to a particular area (e.g., disbursement services, accounts receivable and working capital management, investments, trade finance, foreign exchange) and targeted broadly companies sharing general characteristics, as defined by a limited range of profile data, for example, company size, geography, industry, supply chain position, financial behaviors.

With the proliferation of data points, second-generation NBA recommendation engines required more intensive work to describe the data, that is, to teach machines the data features to look for when analyzing a data set. This enabled banks to identify specific products (e.g., receivables financing, FX hedging, corporate procurement cards) that a customer would likely be considering. However, the analysis behind these recommendations focused on behaviors within specific cash-management functions, reinforcing the narrow scope of product "siloes," with little opportunity to optimize financial performance across the full value chain.

The goal of feature engineering, therefore, is to find the best combination of variables to enable a learning algorithm to recognize meaningful patterns in a particular data set. It is a key process in developing a conventional NBA model and until recently was a time-consuming manual process drawing on deep knowledge of business practices.

Drawing upon a much broader set of structured and unstructured data and taking advantage of recent gains in processing capacity, today's NPtB engines generate recommendations that are much more granular and precise than is possible with NBA models. The improvement in analytical sophistication and processing power comes thanks to deep learning, which can develop highly accurate predictive algorithms. Neural networks make it possible to automate the discovery of data features and the identification of the best combination of features to produce the targeted prediction. These calculations produce remarkably precise predictions and deliver actionable results.

While deep-learning algorithms generate the most accurate predictions if data sources are complex and unstructured, gradient boosting machines, random forest algorithms, and even logistic regressions provide valuable commercial opportunities for the NPtB engine. In addition, gradient boosting machines have the advantage of generating reliable recommendations with smaller data sets or data sets where there are gaps. These models are trained with past data and statistically back-tested on an out-of-sample and out-of-time customer base to quantify model performance.

Deploy and pilot NPtB engine

The last phase is to embed recommendations in digital channels and relationship managers' interactions with clients. Banks test and refine the NPtB pilot in the field before rolling it out to the full market.

A European bank recently tested an NPtB engine in five branches to evaluate the precision of sales leads. A team of relationship managers, product specialists, and branch managers participated in the pilot, which included training on how to follow up on recommendations and testing the effectiveness of leads with clients. Over three months, the team tested the leads with companies and provided feedback. A sceptical relationship manager selected an offer for a letter of guarantee recommended for a particular client, and commented, “I don’t believe the customer will buy this, I know the company.” When the relationship manager asked the company owner if he needed a letter of guarantee this month, he answered, “How did you know? I am currently negotiating this product with another bank.”

In the course of a similar pilot with another financial institution, the model predicted that 4,500 companies, for which there was no indication in the data of previous international trade activity, would purchase international trade products in the coming month. As it happened one in five of these companies purchased an international trade product for the first time within 30 days. Based on the performance of the pilot, the analytics team updated the model before implementing it across the entire organization to target more than one million customers. The full launch included four weeks of coaching for more than 600 relationship managers. Within five months of starting the project, the NPtB was fully up and running, with the predictive model stabilized and relationship managers fully trained. Ultimately, this bank increased new sales by more than 30 percent, and relationship managers increased their interactions with commercial customers by more than 50 percent.

The pilot is an important opportunity to secure the endorsement of team members participating in the pilot, who then share information about the model with other colleagues. The pilot is also an opportunity to test metrics for evaluating the sales process, such as number of visits, percent of leads used by relationship managers, conversion rates, and the level of satisfaction among relationship managers participating in the pilot (versus control group). In addition, the pilot phase is the time to begin testing long-term performance metrics (in order to ensure sustainability in the front line), for example, hit rates for branch staff and relationship managers, customer profitability, and customer satisfaction.

In the transition from pilot to full roll-out, it is crucial to ensure that the organization is aligned around the NPtB use case and that a support team is assigned for the deployment.

Adapting an NPtB engine to serve large corporate clients

Banks have also been able to improve the relevance and timeliness of their recommendations to large corporate clients. At many institutions, relationship managers are thoroughly familiar with the general needs of their corporate customers, but sometimes they are at a loss to anticipate changes in these needs. This was once the case for a large European bank operating in diverse regions. It now draws on a broad range of data to understand general market trends and specific company behaviors. This requires not only applying advanced analytics to traditional types of information (annual reports, market conditions, competitor news) but also collecting publicly available data on social media (company-managed pages, customer comments, etc.). An NPtB engine extracts insights from available data to alert corporate treasurers to new opportunities, for example, to leverage complex banking relationships to improve cash flow and lower the cost of short-term financing. Identified leads include opportunities in various currencies, possibly triggering a change in cross-border pooling arrangements; letters of credit, domestic and international guarantees; or even investment banking products, e.g., debt capital management. NPtB engines can boost new sales among large corporate clients by as much as 15 percent.

Implementing NPtB for SMEs/mid-corporates

The lessons learned from banks that have implemented an NPtB engine can be summarized in five points:

1. **Design the NPtB engine according to the characteristics of the market segments served.** Consider first internal data, including company profiles, relationship characteristics, product granularity, and opportunities. Expand the data set to include external data, testing the relevance of the new variables in generating useful recommendations. In developing algorithms to generate predictions for the large corporate segment, it is important to test a broad variety of external data in order to build a robust data set that can produce insights with a new level of accuracy.

2. **Build the model around customer needs and interests.** One of the biggest impacts is shifting from a "product push" approach to interactions that address specific customer needs, as reflected in current transaction activity and financial performance. This shift enables relationship managers, service representatives, and product specialists to help customers weigh their options and choose the path that best serves the company's financial interests.

3. **Pilot the outcome of the NPtB engine to build confidence and secure buy-in.** Relationship managers must be confident in the opportunities identified by the NPtB engine; at the end of the process they will leverage leads to improve their sales effectiveness, but change management and internal buy-in are key for successful implementation.

4. **Focus on prototypes that create excitement.** Don't let IT and the complexity of legacy systems become the bottleneck, but start with a pragmatic "proof-of-concept" to demonstrate the model's potential. Quick test-and-learn prototypes have multiple purposes, including learning and improving but above all showing prompt impact to create enthusiasm.

5. **Ensure impact from multiple levers.** Better targeting based on analytics is crucial, but there are additional levers, including the timing of recommendations, framing recommendations within a broader value proposition, measuring the impact of recommendations (including the performance of relationship managers), which can also improve performance.

Ignacio Crespo is a partner in McKinsey's Madrid office, where **Carlos Fernandez Naveira** is an expert associate partner. **Huw Kwon** is a senior analytics expert in the London office.

Hidden figures: The quiet discipline of managing people using data

The "war for talent" in financial services has evolved to encompass new frontiers and unfamiliar battlefields. This evolution is fueled by the fundamental transformation of capabilities essential to payments organizations and more broadly, banks' future success: skills in digital technology, artificial intelligence, and automation alongside less tangible abilities such as problem-solving, emotional intelligence, resilience, and adaptability. Similar transformations are playing out across sectors however, leaving such capabilities in scarce supply. At the same time, banks' historical playbook for attracting, developing, and retaining talent is in need of update given erosion in the perceived advantages of a banking career relative to other sectors.

Carla Arellano
Vijay D'Silva
Malte Gabriel
Jaime Potter

Successfully addressing this challenge pays dividends. Banks that prevail in the renewed war for talent will place greater emphasis on employee attraction, management, development and retention, starting with revamping their employee value proposition and embracing evidence-based talent management by deploying new capabilities such as data, analytics, and organizational science. We believe that banks and payments companies that recognize their talent management opportunities, set bold aspirations, and embrace new capabilities to address these imperatives will increasingly distance themselves from their competition.

The new war for talent

The renewed war for talent is global, urgent and poses daunting challenges, for which the imperative is not only to attract and retain the highest performers, but also to enable leaders to better manage talent to deliver sustainable competitive advantage. A 2018 McKinsey analysis of European financial services firms on the future of work identified the critical talent segments these institutions need to fill over the coming years. Top of this list are new roles in software and application development (e.g., scrum master), analytics (e.g., data scientists), new risk management roles (e.g., cybersecurity analysts), and digital marketing (e.g., UX designer). However, research from technology-sourcing firm Catalant finds that while many companies have begun to address top technology and training challenges, most continue to rely on traditional recruiting models that show signs of erosion, leading to key roles often taking 90 or more days to fill. By 2021, McKinsey projects that demand for talent with digital capabilities will outstrip supply by a factor of four in areas like agile, and by 50 to 60 percent for big-data talent, according to a study conducted two years ago.

McKinsey framed the war for talent as a strategic business challenge in 1997, setting forth the notion that better talent leads to better corporate performance. Bank leaders embraced the concept of talent as their firm's most valuable asset; however, responsibility for hiring and development continued to be delegated to human resources or line managers. C-suite leaders focused on other priorities while battles for talent were won with monetary incentive packages that tech firms and companies in other sectors were unable to match.

In the years following the financial crisis, banks focused on cost reduction and risk management to battle margin and regulatory pressures. This left a blind spot for strategic talent questions, and many banks now find themselves with a significant gap in their perceived employee value proposition (EVP) compared to technology and other leading sectors (Exhibits 1,2). As a result, future leadership talent is turning away from careers in finance. In 2007, four times as many US MBA graduates chose to enter the finance field over technology. By 2017, these two groups were roughly at parity. IPOs and similar equity incentives made the tech field more lucrative, shifting the playing field just as banks trimmed their post-financial crisis bonuses. What once was a bank EVP selling point has now faded in comparison to other industries.

It's not only rigorous technical skills that are gaining importance for banks. An ironic aspect of the shift towards automation of many 20th-century jobs is the increasing focus on people skills—flexibility, problem-solving under uncertainty, collaboration, and emotional intelligence, to name a few. While these soft skills often get short shrift when sized up against measurable technical skills in Python and deep learning, their value

Exhibit 1

Technology has overtaken banking in perceived attractiveness of compensation and benefits over the last 4 years.

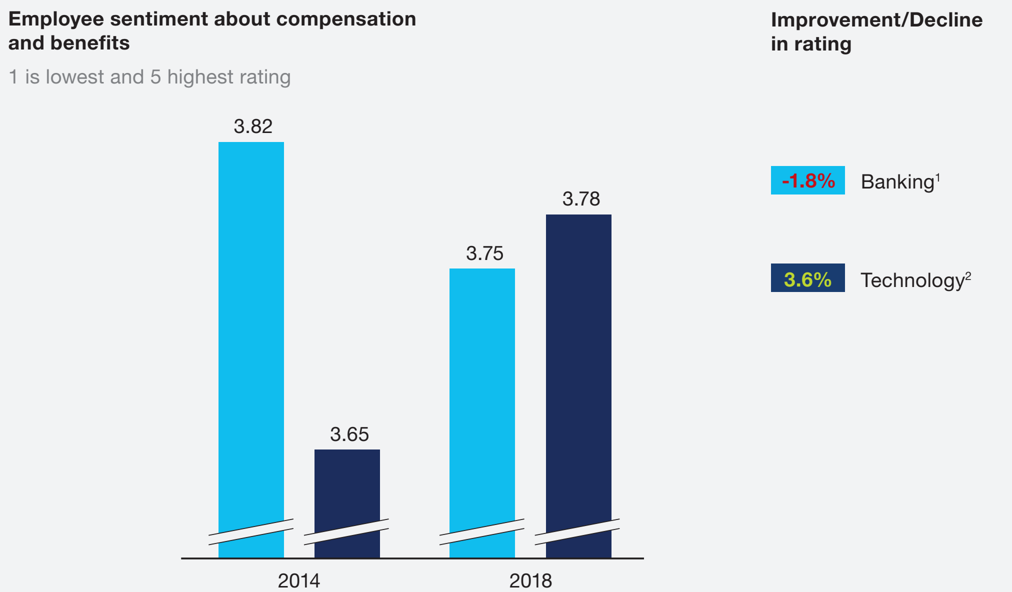

[1] Nine largest US banks by 2017 assets
[2] Ten largest US tech firms by 2017 revenue
Source: McKinsey analysis of publicly available data

should not be underestimated. A University of Michigan study showed that investing in training of soft skills yielded a whopping 250 percent return on investment in certain instances.

We have reason to believe that the war for talent is here to stay. Several parallel forces have fundamentally altered the global landscape, altering banks' roadmap to victory in the talent management space:

Massive amounts of workforce data: The explosion of data over recent years—combined with the power to store and dissect it—opens new avenues for talent management. Email and calendar data are now being used to benchmark the collaborative approaches of effective teams. Engagement surveys—with greater depth and quicker turnaround than in the past—offer valuable insights into the themes and sentiments diverting employee focus, and employee location data can improve operational efficiencies in sectors such as restaurants and delivery services.

Imperatives for speed and accuracy: In an uncertain competitive environment marked by ever-shorter technology lifespans, winners in the war for talent will be able to quickly identify talent gaps at a micro-granular level. Agile talent management will drastically reduce the cost associated with having the wrong or even no talent at all in critical roles.

Exhibit 2

Banks are perceived as more stressful and competitive than technology and other sectors.

Banks' culture is ...[1]

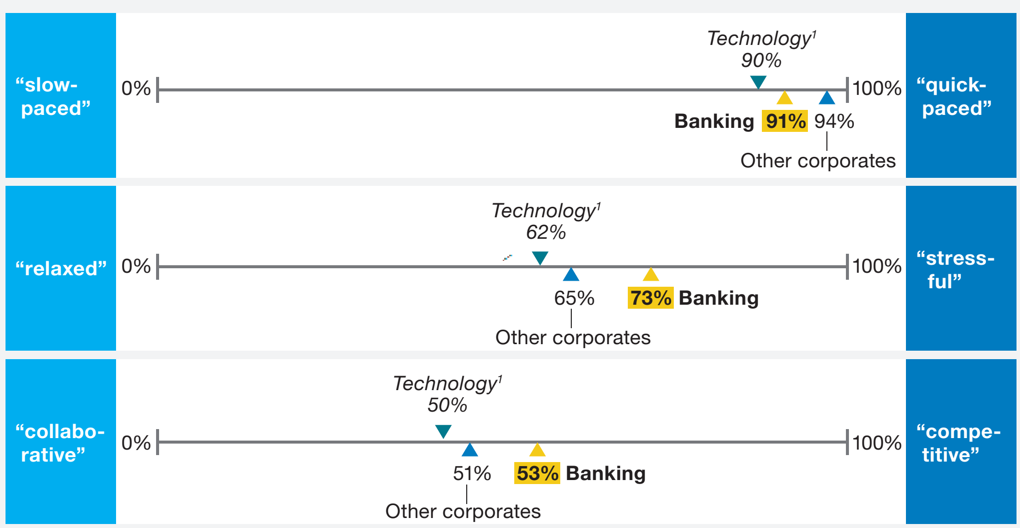

[1] Analysis on employee reviews from 2017-18 for 29 banks, corporate MNCs, and tech firms: largest US banks by assets 2017; largest US tech firms by revenue; top 10 non-financial, non-consulting, non-tech firms sought by MBAs.
Source: McKinsey analysis

Shortage of talent and a rapidly evolving workforce: Forty-six percent of employers have difficulty filling jobs, mainly due to a shortage of applicants. Finance staff, a core constituency of banks and payments firms, rose to the sixth-most-difficult job category to fill in 2016, up from ninth in 2015. Winners will be able to fill jobs quicker and with better people as microtargeting and evidence-based selection allow firms to identify and tailor hiring to target talent and assess candidates more consistently to hire high performers.

Many banks have taken note of how talent is impacted by these trends. An analysis of earnings call transcripts of the ten largest US banks reveals that talent-related terms are being used three to four times as often in recent quarters than they were in the 2012-15 period. With employee turnover rates at ten-year highs, CEOs need to find ways to not only secure talent for the future, but also to stem near-term attrition and its drag on profitability.

Winning with analytics

We have documented substantial performance differences between the leaders and laggards in this new war for talent. Research outlined in the book Talent Wins, co-authored by outgoing McKinsey Managing Partner Dominic Barton, demonstrates that companies that use data and advanced analytics to inform their talent decisions realize up to a 30 percent increase in profits through hiring focus alone—before accounting for the benefits of higher productivity and better retention. Furthermore, 2018 McKinsey research on performance management indicates that organizations with effective performance management are 77 percent more likely to outperform competitors and peers.

Winners will be those firms who can harness data, advanced analytics, and behavioral science to make sound people and organization decisions faster, better, and with a level of specificity previously unavailable. This will enable them to preserve advantages gained by better deploying and nurturing skills across the full talent lifecycle.

We see three areas in which firms must master data and analytics in order to win the war for talent:

Assess talent gaps and address accordingly: The adage "culture eats strategy for breakfast" may soon be replaced—or at least complemented with—"and capabilities take lunch." While many organizations invest significantly in strategy, the key is securing the capabilities needed to deliver that strategy. Leveraging internal and third-party data allow firms to quantify organizational skills deficits, target opportunities for re-skilling (through methods such as hierarchical clustering or cosine similarity), and identify the skills to source externally. A 2017 McKinsey Global Institute report on automation, employment and productivity showed that 43 percent of all finance and insurance activities can be automated through currently available technology. The aforementioned McKinsey study on the future of work in financial institutions found that one-third of existing talent gaps can be addressed by re-skilling current employees. One client established a best-practice adult-learning program, combining both in-house and external learning, and retrained more than 1,000 employees into new internet of things, analytics, and machine learning roles within the first ten months of the program. Winners in automation transformations pinpoint capability requirements and make the proper call on where to buy (source), build (re-skill), and rent (outsource or shift to contractors).

Attracting and retaining the best talent: Microtargeting allows a firm to tailor its EVP and its communication to critical talent segments to increase conversion rates. There is clear evidence that objective hiring powered by analytics and behavioral science (versus traditional interviews) leads to better hiring decisions and greater value creation (Exhibit 3). In this area banks can learn from each other, as well as from other sectors. For instance, firms like Aegis Worldwide conduct text message-based initial interviews to reduce bias and enable algorithmic analysis of answers. Unilever uses technology across its whole recruiting process; it begins by using LinkedIn profiles instead of résumés, deploying AI to select the best prospects. Next, it uses a series of online games to further narrow the field to a select few candidates for in-person interviews. The results are convincing: Using this approach Unilever tripled the roster of universities from which it recruits while reducing its average hiring cycle from four months to four weeks. McKinsey's own use of artificial intelligence to screen résumés not only delivered 30 to 50 percent increases in hiring efficiency (reflecting a 400 to 500

Exhibit 3

Bank x is missing data science capabilities, and can address the gap by tapping into existing backend database and programming expertise.

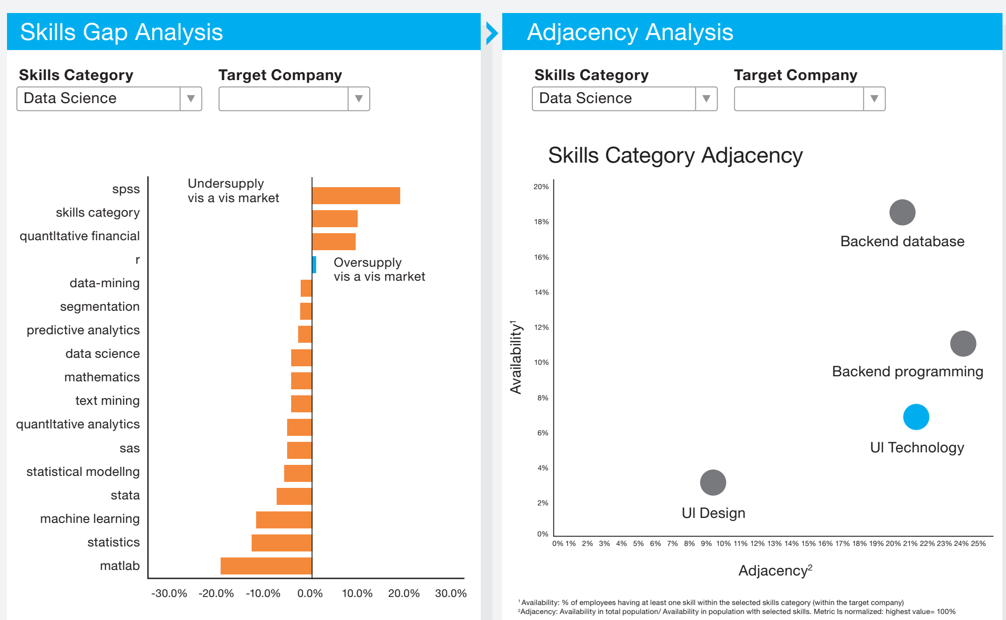

Source: McKinsey analysis

percent ROI), but also drove an increase in the share of women candidates passing initial screening.

One financial services firm found it could increase field effectiveness by systematically testing and recruiting for character traits such as curiosity and de-emphasizing humility, which was found to be underrepresented among its high performers. Another applied a predictive modeling technique called elastic net regularization to hire more employees with character traits similar to those of existing high performers, reallocating recruiting dollars to particular schools and majors and disintermediating headhunters in approaching high-potential candidates. A fast-food chain collected data on employees' character traits and behaviors in the workplace and generated more targeted recruiting, roles and expectations, culture-focused trainings, updated financial and non-financial incentives, and optimized shifts. The company soon documented customer satisfaction gains of 130 percent across stores, 30-second speed-of-service improvements, and per-store revenue increases of 5 percent.

We predict that winners will go beyond deploying "off the shelf" assessments to develop evidence-based models supporting the knowledge, skills, attributes, and experiences required to successfully deliver on a specific role in its unique environment. This can be accomplished through closed-loop machine learning to pinpoint what factors distinguish high performers from the rest, or science-based forensics on future work required, which informs objective screening criteria to be assessed through science-backed interviews and digital assessments, gamified or otherwise.

High performers are often at the highest risk of attrition, given their multitude of outside options. Data and analytics can serve as an early warning and mitigation system by predicting attrition risk at both individual and group levels and developing effective responses to address the root cause. For instance, using k-medoid and majority vote classification techniques, one financial institution found that attrition was elevated among three different groups—millennials seeking professional growth, employees working in larger teams, and those working for low-tenured managers. Leveraging this data-driven employee segmentation, the organization developed tailored preventive measures to reduce attrition for each of the clusters.

Better manage and deploy talent

A plan to grow and deploy talent starts with identification of what drives true performance—collecting data to create a 360-degree view of who your employees are, what they do, who they interact with, how they're deployed—linking this information to the relevant dependent variables and building optimization strategies. This typically starts with a data-driven assessment of the organizational context for employee performance. For instance, to what extent does a manager's span of control impact individual performance? What role does coaching play in performance? More ambitious initiatives might develop guidance on time allocation, collaboration patterns, meeting practices, and more, through behavioral data such as calendar and email metadata (with appropriate encryption methodologies to maintain employee privacy).

As an example, one financial institution built and analyzed a behavioral dataset of how leaders split their time across recruiting, coaching, clients, and other activities to identify a 30 percent growth opportunity in investments. The bank reallocated leaders' time from administrative and controls-oriented tasks to customer-centric coaching and fostering connectivity across lines of business. Another firm leveraged McKinsey's data-driven Talent to Value approach to identify a select number of critical roles driving the most economic value—some extending as far as four levels below the CEO. This client discovered that closer collaboration across four of 30 critical roles was critical to delivering more than 50 percent of the value at stake. Rather than merely encouraging general collaboration across the enterprise, the firm doubled down on tactical incentives for collaboration among these roles by, for example, implementing shared goals, with 30 to 50 percent of leaders' KPIs driven by factors beyond their direct control.

These steps create a "virtuous cycle" benefitting workers as well as employers. A better selection process leads to better organizational fits, which in turn fosters employee satisfaction and enhanced EVP. When organizations are redesigned to be more collaborative and agile, not only does employee time allocation change, but roles evolve too and traditional management hierarchies become redundant. At the end of the process, one European bank eliminated two entire layers of middle management while its employee engagement scores rose by over 20 percentage points.

What can be done today?

Humans do not change their behavior with the flip of a switch. It may take years to get a single individual to change behavior, which is only compounded when we consider the thousands of employees with unique values, goals, and aspirations working at modern-day organizations.

While many people-related changes take time to reach full potency, most organizations possess the building blocks in both capabilities and data to start with small changes today to pave the way for larger shifts tomorrow. A key first step is to identify the human component of business challenges and opportunities, and build an analytics engine to collect data and validate hypotheses on performance drivers. For example, in a corporate bank, analytics on calendar meta data may help pinpoint the interaction patterns related to deal success. While in payments, data and analytics can enable faster and more nuanced hiring of the right combination of technical and "soft" skills. Deploying analytics to create transparency into what matters—for leaders, managers, and employees—empowers them to cut through the noise and focus on what really matters. Financial institutions looking to upgrade their talent management practices can follow a few simple guidelines to get started:

Make talent the business's agenda: A firm's people analytics agenda must focus on critical business needs and originate from a strong hypothesis on which factors do and do not matter to business performance. Setting this agenda is a collaboration between business and HR leaders.

Don't underestimate what you already have: Relevant data is often already available and can be complemented with nominal effort. In our experience, three out of four banks already possess the necessary data—such as attrition rates, team structures, employee backgrounds, and average time to fill a position—to test the most pressing people analytics hypotheses.

Treat data with the care and rigor it deserves: Protecting data privacy and employee confidentiality are critical objectives, not only since the GDPR rules on data protection took effect earlier this year. Protecting data privacy is also core to preventing a situation where employees feel they are being unduly monitored or even manipulated. Create transparency on which data is sourced, how it is used and the tangible benefit that people analytics can provide. Establish protocols and encryption policies to appropriately anonymize and mask information.

Start small and build over time: Significant value can be gained by combining HR, financial, and operational data for basic "talent due diligence." The first step is to identify drivers of compensation growth, performance ratings, promotions, and variance across the organization. This will begin to infuse talent decisions with the rigor normally reserved for financial decisions. By running a talent due diligence, one European bank quickly identified an opportunity for retooling, finding that most new employees were not being hired into the most critical divisions and roles, and that base pay was more correlated with age and job grade than criticality of the role or performance. The people analytics journey is a transformation, comparable to robotic process automation. Start small, adhere to high standards when handling data, and quickly prove the value of the approach. A test-and-learn approach makes a difference, running trials to prove business value before scaling more broadly.

* * *

Twenty years ago, the war for talent was fought with major changes in employee environment and compensation systems, triggering a number of innovations and a new informality—down to casual Fridays. Today's changes are more nuanced and targeted. Instead of large-scale changes, the new war for talent will likely involve thousands of subtler microdecisions. This scope can seem daunting. Fortunately, embedding data and analytics into an organization's people function begins with a few simple changes today that will lay the groundwork for a more profitable organization and more fulfilled employees.

Carla Arellano is a partner in McKinsey's New York office, where **Vijay D'Silva** is a senior partner. **Malte Gabriel** is a data science specialist in the North America Knowledge Center, and **Jaime Potter** is a senior product analyst in the New Jersey office.

Using analytics to increase satisfaction, efficiency, and revenue in customer service

As payments providers around the globe cope with increasing pressure on revenues and margins, customer service is increasingly becoming an important asset for driving top- and bottom-line performance, and improving the customer experience. While most banks, card companies, and other payments providers have implemented various degrees of customer service transformation by using advanced analytics, the discipline has yet to be fully leveraged in this regard. To realize the full potential of today's analytical capabilities financial institutions will need to possess, acquire, or develop the relevant capabilities and use them to customize and enhance a wide range of customer interactions.

Krishna Bhattacharya
Greg Phalin
Abhilash Sridharan

Payments providers that adopt advanced analytics to develop broad integrated approaches are seeing significant improvement: customer satisfaction scores rose 5 to 10 percent and operating costs declined 15 to 20 percent when they used analytics to eliminate cross-channel leakage and migrate more customer interactions into self-serve channels. Analytics also enabled these firms to improve customer retention and revenues by 10 percent or more, by enhancing the customer journey and improving cross-selling.

The future of customer service

Customer service is shifting dramatically, from phone and branch-centric models to an omnichannel interaction dynamic in which customers move seamlessly among service channels, including mobile, phone, chat, and online. A McKinsey survey in 2015 showed digital channels accounted for 30 percent of customer interactions. We expect this share will approach 50 percent by 2020. And of this, 26 percent will be exclusively digital with no branch interaction.

Payments customers expect high-quality service across channels, similar to what they enjoy at other financial institutions and leading service providers, like Amazon and Zappos. To deliver this level of service, payments firms need to optimize customer and prospect telecommunications and deliver seamless omnichannel interactions.

Building an omnichannel customer service model

Traditionally, financial institutions have tried to optimize customer service within channel silos, including call centers, online, and mobile. The key to delivering a high-quality omnichannel experience is adopting a broad customer journey approach that integrates customer interactions across digital and traditional channels. Several institutions have already embarked on such a model. A global life insurer, for example, recently developed a five-year plan to migrate nearly half of its customer journeys into self-serve channels. However, too often such changes are viewed as one-time efforts rather than as a large-scale transformation. Designing a comprehensive, ongoing program is key to sustaining omnichannel service improvements.

Investing in the talent to transform

A key part of transforming the customer experience is migrating basic transactions to self-service channels, and complex transactions to agent-assisted channels. While most organizations invest in ongoing agent training and capability building, transforming the customer experience demands a more substantial investment in talent. It requires investing in technology that enables customer service professionals to have more effective interactions with customers. For example:

- Real-time coaching software, such as Cogito, provides live feedback about customers to agents during customer calls, so agents can tailor the discussion to customer needs.
- Applications such as Verint use speech analytics that foster more personalized interactions with customers.

To provide more personalized customer service, financial institutions must rethink how they interact with customers and prospects. Analytics can personalize customer experience by, for example, identifying the next-best action or product offering. (See "Using data to unlock the potential of an SME and mid-corporate franchise," page 28.)

Investments in technology are, of course, critical to transforming the customer experience. Two investment types in particular are key: developing the agility to rapidly build, pilot, and launch a broad transformation; and robotics or artificial intelligence (AI) to reduce manual workloads, improve cycle times, and minimize back office errors. McKinsey research shows that 65 percent of back-office tasks at contact centers, and 30 to 50 percent of front-line calls, can now be automated.

Six hallmarks of analytics success

Financial institutions that are successfully using advanced analytics to enhance the customer experience share six common hallmarks (Exhibit 1).

1. Migrating customers to digital channels

Given customers' preference for omnichannel service, there are two important questions financial institutions must address: First, how do they create seamless transactions for digital natives, who prefer digital-only service? Second, in serving less digitally inclined customers, how can financial institutions use tools like journey analytics to prevent the use of multiple channels for the same query? The main challenge for customer service organizations is to identify the most appropriate transactions for migration and ensuring they are completed satisfactorily in digital channels whenever possible. Payments leaders in digital migration are achieving 20 to 30 percent reductions in call volume and successfully enhancing the customer experience. Some industry leaders are also developing a 360-degree, multitouch, multichannel view of customer interactions using journey analytics; but this requires robust integrated datasets that can capture customer interactions across channels.

2. Improving behavioral routing and IVR containment

Financial institutions have been using interactive voice response (IVR) technology for several decades, but few have optimized these capabilities. Doing so requires more than investing in additional VR capabilities. Financial institutions can apply advanced analytics or AI-based technologies to improve behavioral routing and IVR containment:

- Using analytics to identify reasons for call transfers can help increase the number of interactions contained within the IVR environment. Deeper analysis of calls can classify customers into clusters based on value, behavior, and tenure, speeding up IVR service and streamlining unnecessary trees.

- Matching agents to callers based on personality (using technologies like Afiniti and Mattersight) can meaningfully improve customer experience and call efficiency.

- Directing calls from high-potential customers to agents trained to present tailored products (using algorithms based on the customer's needs) can boost productivity.

3. Strengthen identity validation and personalize product offerings

The layering of analytics on video and audio channels can improve identity validation and personalize the product offering. Examples include:

Exhibit 1

Analytics use cases in customer care center around six core imperatives.

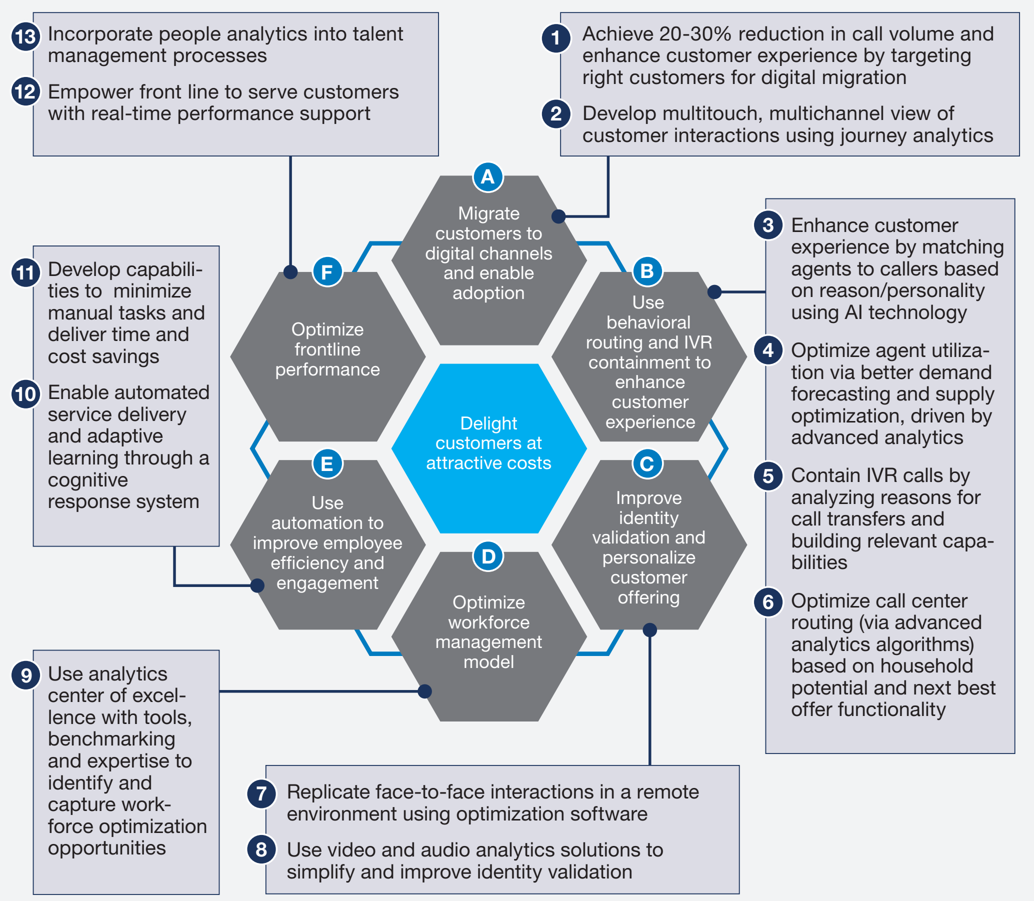

Source: McKinsey Analytics

- Replicating face-to-face interactions in a remote environment using optimization software enables more personalized and secure interaction.
- Identity validation can be simplified and improved with features like facial recognition (online identification) and voice recognition (in app account access).

4. Optimize the workforce management model

Most financial institutions have established internal analytics centers staffed with experts working to capture workforce optimization opportunities. Yet, most workforce management practices are rooted in backward-looking general demand-supply matching, assuming some average service level for a day. However, customer research reveals that assumptions of averages fall short. There are three important challenges for each financial institution:

- How can they effectively manage the tails that drive customer satisfaction or dissatisfaction?
- How can they use machine learning to manage resiliency and drive the next level of predictive modeling on demand (e.g., impact of hurricanes)?
- How can analytics centers use real-time simulation tools to create efficiencies in workforce management?

5. Automate to improve employee efficiency and engagement

Thus far, automation has not been systematically applied in the customer service environment. In customer care, AI can be used to automate services by supporting customers with virtual agents, and contact center agents through real-time interaction tools (e.g., automated knowledge management systems) and back-end automation (e.g., robotic process automation). Virtual agents can solve customer requests by using natural language processing technology, and get smarter over time through machine learning. For example, programs like IPSoft's Amelia can play the role of any customer service agent by rapidly absorbing call logs, recognizing emotional context, and interacting with customers, thereby saving costs and lifting both revenue and customer experience. With large tech players moving into the digital assistant arena, we expect things to evolve quickly in this area.

6. Optimize frontline performance through analytics in recruiting

Recruiting processes for customer service organizations are seldom informed by what makes agents successful. Leading firms take an approach called people analytics methodology, which reverse engineers the process, starting with the best customer service agents and identifying common traits that makes them successful. They then apply these insights at the top of the recruiting funnel in selecting candidates. By applying people analytics in this way, financial institutions can improve talent management in customer experience as well as in the wider organization.

Case example I: Improving digital channel experience and digital adoption

Recently, a North American bank used journey analytics to accelerate digital adoption across its customer base. Using analytics and design thinking to address digital adoption levers across customer journeys (rapid digitization, containment, signature moments, customer targeting), the bank achieved a gain of more than 20 percentage points in digital engagement. The initiative included the following elements:

- **Journey level scan:** Using interaction data and analytics from all channels (digital, call, branch, email/text, ATM), the bank prioritized about 15 core customer journeys and more than 40 sub-journeys for digitization,
- **Quantified journey redesign:** The bank then redesigned each core journey using analytics-based Quantified Experience Design (QED),[1] leading to an increase in digital engagement of 10 to 15 percentage points, and similar improvements in customer experience measures. Analytics drills targeted key drivers of customer experience and other cross-cutting themes.
- **Real-time customer nudging:** The bank introduced a customer targeting process based on customer behaviors and journeys to accelerate digital adoption, which generated a 5- to 10-percentage-point increase in product adoption
- **Journey tracking:** The bank transitioned from an overall customer experience-based performance measurement system to one based on operating drivers for each journey and channel, to track improvements and re-orient program
- **Capability building:** Using journey analytics and QED, the bank designed and launched a capability-building program for more than 800 contact center agents.

Case example II: Enhanced contact management

A credit card company was struggling to migrate customers to its self-serve channels despite having invested in natural-language speech IVR. Consequently, it devised a three-pronged approach to accelerate migration, which focused on resolving (and containing) a higher percentage of calls within their IVR, and delivered a differentiated experience along the customer journey:

1. To better understand its customers' behavior, the company analyzed five million customer calls. With these findings, they classified customers into eight archetypes based their value, behavior, and length of time as customers.
2. Management also used brainstorming techniques to develop and refine several initiatives based on feasibility, potential economic impact, and customer experience improvement. This generated 48 prioritized initiatives that spanned VR (e.g., capture additional information and make it less easy for customers to "rep out"), routing (e.g., adapt service standards to match expectations of different customers), and post-VR (e.g., focus on education and self-service awareness for disengaged customers).
3. The company also surveyed 1,500 employees, conducted focus groups that engaged managers, and surveyed more than 1,000 customers to explore tactics for increasing IVR containment and digital engagement.

Through these efforts, the credit card provider identified 200 to 500 bps in potential improvement in the containment rate (Exhibit 2). The VR enhancements and post-VR agent initiatives also led to a 5 to 10 percent reduction in costs or incremental annualized savings.

Case example III: Demand forecasting

The call center head of a large UK-based bank turned to analytics to optimize agent utilization by automating demand forecasting, as part of a larger analytics-driven transformation at the institution. The approach incorporated the following elements:

[1] QED is a data and design methodology that helps executives prioritize and implement customer journey designs, and helps designers focus on features that will maximize value for customers. For more insight visit: www.mckinsey.com/business-functions/digital-mckinsey/our-insights/digital-blog/design-meet-data-unlocking-design-roi.

Exhibit 2

VR containment and differentiated customer experience led to 200-500 bps containment increases and annual cost savings of 5-10%.

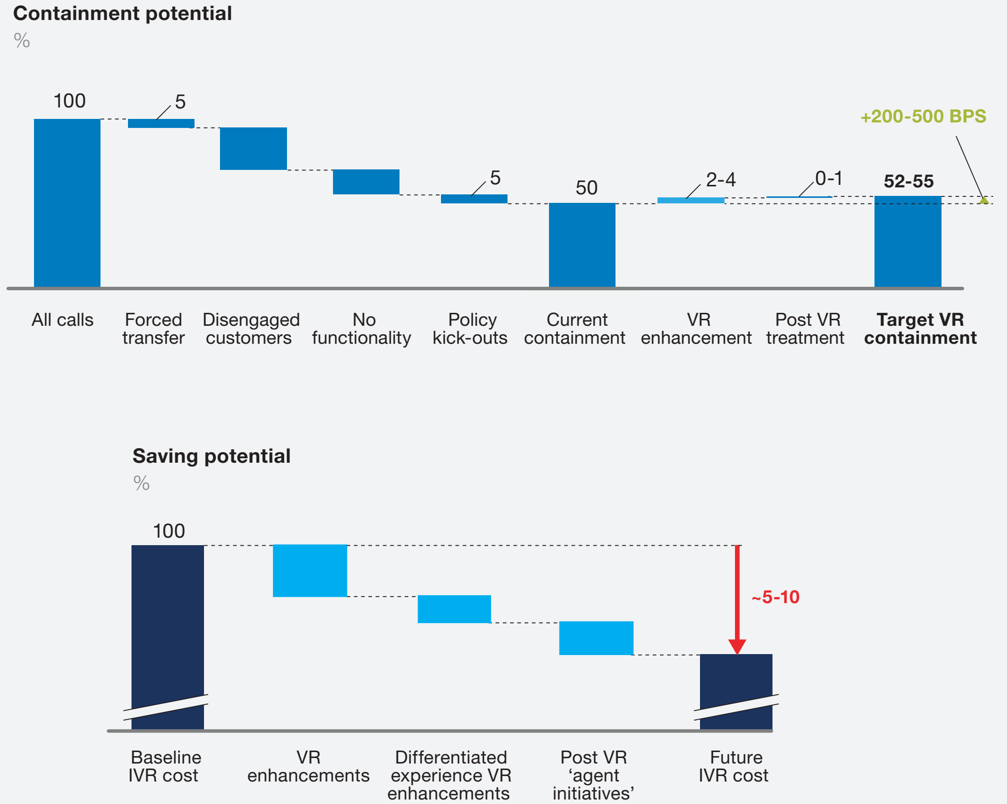

NOTE: 100M calls into the client VR per year

Source: McKinsey Analytics

- **Creation of a robust integrated dataset** that is foundational for the analytics exercise, by combining five different data sources—data for more than ten million customers, call data, agent data, bank data related to IT outages, and other external data (e.g., weather)
- **Development of two sets of random forest machine learning models** to continuously learn thresholds and forecast both number of calls and average handling time, on a monthly basis (4 to 16 months ahead and updated monthly) and a 30-minute level basis (8 to 10 weeks ahead and updated daily)
- **Bayesian techniques** to capture most recent dynamics for extrapolation, non-linear regression models for forecasting, and more than a hundred features to capture different levels of seasonality.

The bank achieved a 20 to 40 percent error reduction in forecasting for a subset of population and are rolling it out across all FTEs.

Starting the journey on analytics to customer service

When introducing advanced analytics, a critical first step is clearly understanding the organization's current position in terms of one of three horizons (Exhibit 3):

Those on *Horizon 1* generally have low levels of awareness regarding recent developments in advanced analytics for customer service. These organizations need to begin their transformation by building a business case, educating their leadership, and obtaining organizational buy-in. Once these initiatives are underway, quick, tangible wins should be pursued to reinforce the organization's commitment to a full transformation. Additionally, another challenge faced by these organizations is lack of in-house knowledge on relevant frameworks and solutions, to diagnose and prioritize initiatives.

Enterprises at *Horizon 2* have a better understanding of recent advances in the field, and have started to experiment with or

Exhibit 3

Most financial institutions are situated on one of three horizons in terms of their customer service analytics journey.

	Horizon 1	Horizon 2	Horizon 3
	Financial institutions mostly in the preliminary stages of digital and analytics evolution	Financial institutions that have an analytics plan but an unstructured approach, leaving value on the table	Firms that are already in the midst of next-gen digital and analytics transformation, but need support to accelerate value delivery
Typical challenges	Need to build awareness among senior leaders about analytics use cases and develop buy-in to start transformation Lack of framework and relevant solutions to diagnose and prioritize initiatives Need early tangible wins to reinforce senior leadership's confidence in transformation	Lack structured roadmap, benchmarked with external best practices, to drive milestones Lack well articulated plan for how to create and capture value Lack capabilities, tools and talent to execute	Business-as-usual activities crowd out dedicated resourcing to implement the project Lack knowledge about best practices in analytics solutions across industries, to help set benchmarks and accelerate value delivery

Source: McKinsey Analytics

adopt them. However, they have done so largely on an ad hoc, unstructured basis. Unfortunately, informal approaches are likely to leave significant value on the table. The key challenge for Horizon 2 organizations is to identify the most efficient path for delivering the desired results. This might be accomplished, for instance, by shaping their perspectives through a sharing of external best practices, and then setting challenging timelines.

Horizon 3 firms are well ahead of the curve, applying next-generation analytics solutions to transform the customer service model. At this stage, the key challenge is finding ways to advance to even higher levels, and to continue to invest in next-generation solutions.

* * *

The use of new analytical tools and capabilities are transforming customer service in financial services. The following questions can help firms shape their strategy discussions:

- Where do we stand currently in terms of the three advanced analytics/customer service horizons?
- What challenges are preventing us from advancing to the next horizon?
- What immediate steps can we take to address these challenges?

Krishna Bhattacharya is a partner, **Greg Phalin** is a senior partner, and **Abhilash Sridharan** is a consultant, all in McKinsey's New York office.

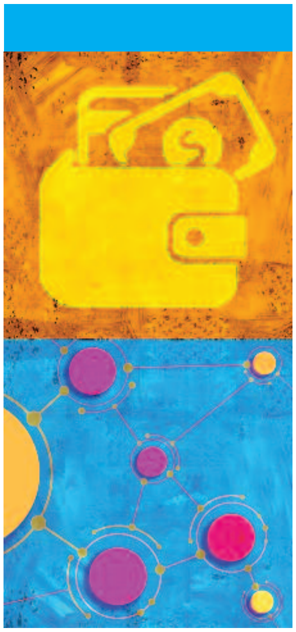

Designing a data transformation that delivers value right from the start

The CEOs of most financial institutions have had data on their agenda for at least a decade. However, the explosion in data availability over the past few years—coupled with the dramatic fall in storage and processing costs and an increasing regulatory focus on data quality, policy, governance, models, aggregation, metrics, reporting, and monitoring—has prompted a change in focus. Most financial institutions are now engaged in transformation programs designed to reshape their business models by harnessing the immense potential of data.

Chiara Brocchi
Davide Grande
Kayvaun Rowshankish
Tamim Saleh
Allen Weinberg

Leading financial institutions that once used descriptive analytics to inform decision-making are now embedding analytics in products, processes, services, and multiple front-line activities. And where they once built relational data warehouses to store structured data from specific sources, they are now operating data lakes with large-scale distributed file systems that capture, store, and instantly update structured and unstructured data from a vast range of sources to support faster and easier data access. At the same time, they are taking advantage of cloud technology to make their business more agile and innovative, and their operations leaner and more efficient. Many have set up a new unit under a chief data officer to run their data transformation and ensure disciplined data governance.

Successful data transformations can yield enormous benefits. One US bank expects to see more than $400 million in savings from rationalizing its IT data assets and $2 billion in gains from additional revenues, lower capital requirements, and operational efficiencies. Another institution expects to grow its bottom line by 25 percent in target segments and products thanks to data-driven business initiatives. Yet many other organizations are struggling to capture real value from their data programs, with some seeing scant returns from investments totaling hundreds of millions of dollars.

A 2016 global McKinsey survey found that a number of common obstacles are holding financial institutions back: a lack of front-office controls that leads to poor data input and limited validation; inefficient data architecture with multiple legacy IT systems; a lack of business support for the value of a data transformation; and a lack of attention at executive level that prevents the organization committing itself fully (Exhibit 1). To tackle these obstacles, smart institutions follow a systematic five-step process to data transformation.

1. Define a clear data strategy

Obvious though this step may seem, only about 30 percent of the banks in our survey had a data strategy in place. Others had embarked on ambitious programs to develop a new enterprise data warehouse or data lake without an explicit data strategy, with predictably disappointing results. Any successful data transformation begins by setting a clear ambition for the value it expects to create.

In setting this ambition, institutions should take note of the scale of improvement other organizations have achieved. In our experience, most of the value of a data transformation flows from improved regulatory compliance, lower costs, and higher revenues. Reducing the time it takes to respond to data requests from the supervisor can generate cost savings in the order of 30 to 40 percent, for instance. Organizations that simplify their data architecture, minimize data fragmentation, and decommission redundant systems can reduce their IT costs and investments by 20 to 30 percent. Banks that have captured benefits across risk, costs, and revenues have been able to boost their bottom line by 15 to 20 percent. However, the greatest value is unlocked when a bank uses its data transformation to transform its entire business model and become a data-driven digital bank.

Actions: Define the guiding vision for your data transformation journey; design a strategy to transform the organization; establish clear and measurable milestones.

2. Translate the data strategy into tangible use cases

Identifying use cases that create value for the business is key to getting everyone in the organization aligned behind and committed to the transformation journey. This process typically comprises four steps.

In the first step, the institution breaks down its data strategy into the main goals it wants to achieve, both as a whole and within individual functions and businesses.

Exhibit 1

Typical challenges in data transformations at banks.

Challenges faced in improving data quality at the enterprise level, ranked by perceived importance

Number of participants ranking the challenge in 1st or 2nd place, from a total of 43 respondents

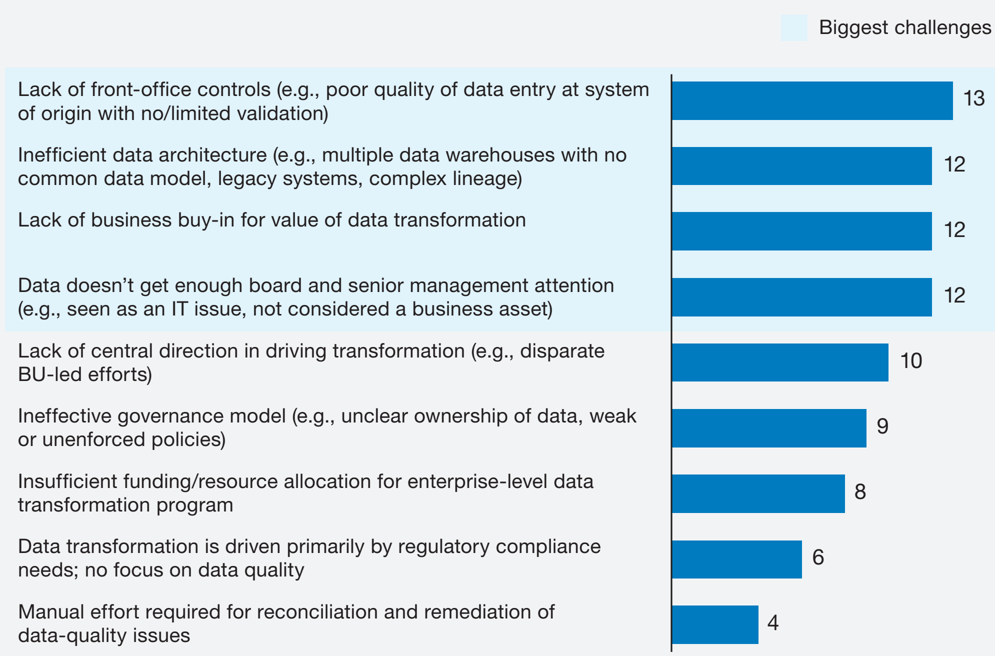

Source: McKinsey analysis

Next it draws up a shortlist of use cases with the greatest potential for impact, ensures they are aligned with broader corporate strategy, and assesses their feasibility in terms of commercial, risk, operational efficiency, and financial control. These use cases can range from innovations such as new reporting services to more basic data opportunities, like the successful effort by one European bank to fix quality issues with pricing data for customer campaigns, which boosted revenues by 5 percent.

Third, the institution prioritizes the use cases, taking into account the scale of impact they could achieve, the maturity of any technical solutions they rely on, the availability of the data needed, and the organization's capabilities. It then launches pilots of the top-priority use cases to generate quick wins, drive change, and provide input into the creation of a comprehensive business case to support the overall data transformation. This business case includes the investments that will be needed for data technologies, infrastructure, and governance.

The final step is to mobilize data capabilities and implement the operating model and data architecture to deploy the use cases through agile sprints, facilitate scaling up, and deliver tangible business value at each step (Exhibit 2). At one large European bank, this exercise identified almost $1 billion in expected bottom-line impact.

Actions: Select a range of use cases and prioritize them in line with your goals; use top-prior-

Exhibit 2

Banks can deliver end-to-end use cases at speed via agile sprints.

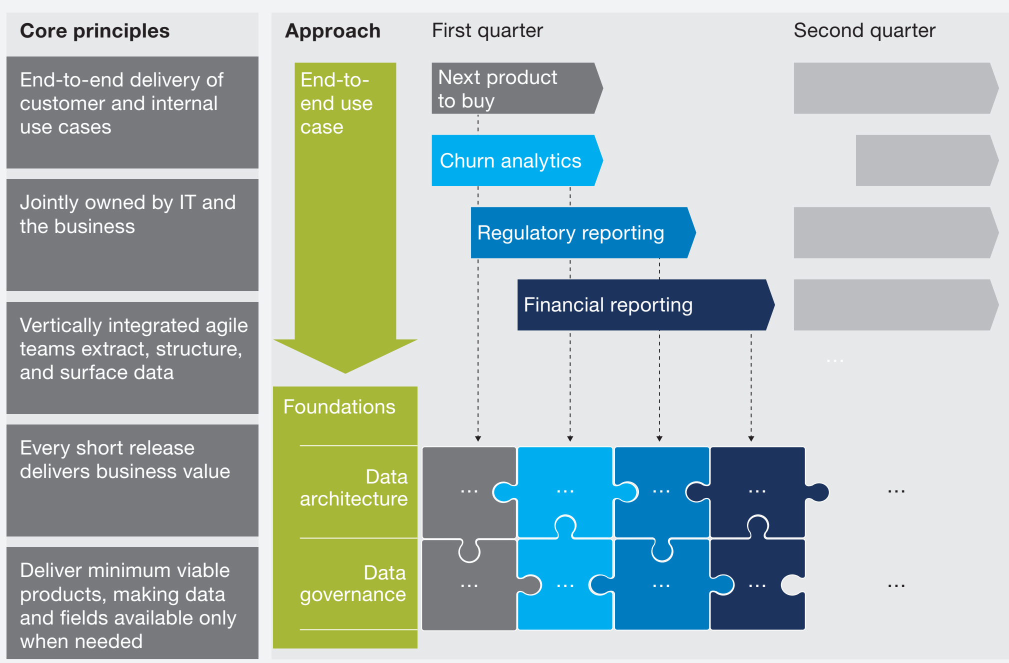

Source: McKinsey analysis

ity use cases to boost internal capabilities and start laying solid data foundations.

3. Design innovative data architecture to support the use cases

Leading organizations radically remodel their data architecture to meet the needs of different functions and users and enable the business to pursue data-monetization opportunities. Many institutions are creating data lakes: large, inexpensive repositories that keep data in its raw and granular state to enable fast and easy storage and access by multiple users, with no need for pre-processing or formatting. One bank with data fragmented across more than 600 IT systems managed to consolidate more than half of this data into a new data lake, capturing enormous gains in the speed and efficiency of data access and storage. Similarly, Goldman Sachs has reportedly consolidated 13 petabytes of data into a single data lake that will enable it to develop entirely new data-science capabilities.

Choosing an appropriate approach to data ingestion is essential if institutions are to avoid creating a "data swamp": dumping raw data into data lakes without appropriate ownership or a clear view of business needs, and then having to undertake costly data-cleaning processes. By contrast, successful banks build into their architecture a data-governance system with a data dictionary and a full list of metadata. They ingest into their lakes only the data needed for specific use cases, and clean it only if the business case proves positive, thereby ensuring that investments are always linked to value creation and deliver impact throughout the data transformation.

However, data lakes are not a replacement for traditional technologies such as data warehouses, which will still be required to support tasks such as financial and regulatory reporting. And data-visualization tools, data marts, and other analytic methods and techniques will also be needed to support the business in extracting actionable insights from data. Legacy and new technologies will coexist side by side serving different purposes.

The benefits of new use-based data architecture include a 360-degree view of consumers; faster and more efficient data access; synchronous data exchange via APIs with suppliers, retailers, and customers; and dramatic cost savings as the price per unit of storage (down from $10 per gigabyte in 2000 to just 3 cents by 2015) continues to fall.

In addition, the vast range of services offered by the hundreds of cloud and specialist providers—including IaaS (infrastructure as a service), GPU (graphics-processing unit) services for heavy-duty computation, and the extension of PaaS (platform as a service) computing into data management and analytics—has inspired many organizations to delegate their infrastructure management to third parties and use the resulting savings to reinvest in higher-value initiatives.

Consider ANZ's recently announced partnership with Data Republic to create secure data-sharing environments to accelerate innovation. The bank's CDO, Emma Grey, noted that "Through the cloud-based platform we will now be able to access trusted experts and other partners to develop useful insights for our customers in hours rather than months."

Actions: Define the technical support needed for your roadmap of use cases; design a modular, open data architecture that makes it easy to add new components later.

4. Set up robust data governance to ensure data quality

The common belief that problems with data quality usually stem from technology issues is mistaken. When one bank diagnosed its data quality, it found that only about 20 to 30 percent of issues were attributable to systems faults. The rest stemmed from human error, such as creating multiple different versions of the same data.

Robust data governance is essential in improving data quality. Some successful financial institutions have adopted a federal-style framework in which data is grouped into 40 to 50 "data domains," such as demographic data or pricing data. The ownership of each domain is assigned to a business unit or function that knows the data, possesses the levers to manage it, and is accountable for data quality, with metadata management (such as mapping data lineage) typically carried out by "data stewards." A central unit, typically led by a chief data officer, is responsible for setting up common data-management policies, processes, and tools across domains. It also monitors data quality, ensures regulatory compliance (and in some cases data security), supports data remediation, and provides services for the business in areas such as data reporting, access, and analytics.

Best-in-class institutions develop their own tools to widen data access and support self-service data sourcing, like the search tool one bank created to provide users with key infor-

Exhibit 3

A custom-designed search tool provides users with key information on data elements.

Source: McKinsey analysis

mation about the definition, owner, lineage, quality, and golden source of any given piece of data (Exhibit 3). Organizations with readily accessible information and reliable data quality can deliver solutions much more quickly and with greater precision. They can also create enormous efficiencies along the whole data lifecycle from sourcing and extraction to aggregation, reconciliation, and controls, yielding cost savings that can run as high as 30 to 40 percent.

Actions: Assess data quality; establish robust data governance with clear accountability for data quality; provide self-service tools to facilitate data access across the whole organization.

5. Mobilize the organization to deliver value

Successful data transformations happen when a company follows an approach driven by use cases, promotes new ways of working, and mobilizes its whole organization from the beginning. Adopting a use-case-driven

Exhibit 4

Data governance is rolled out domain by domain.

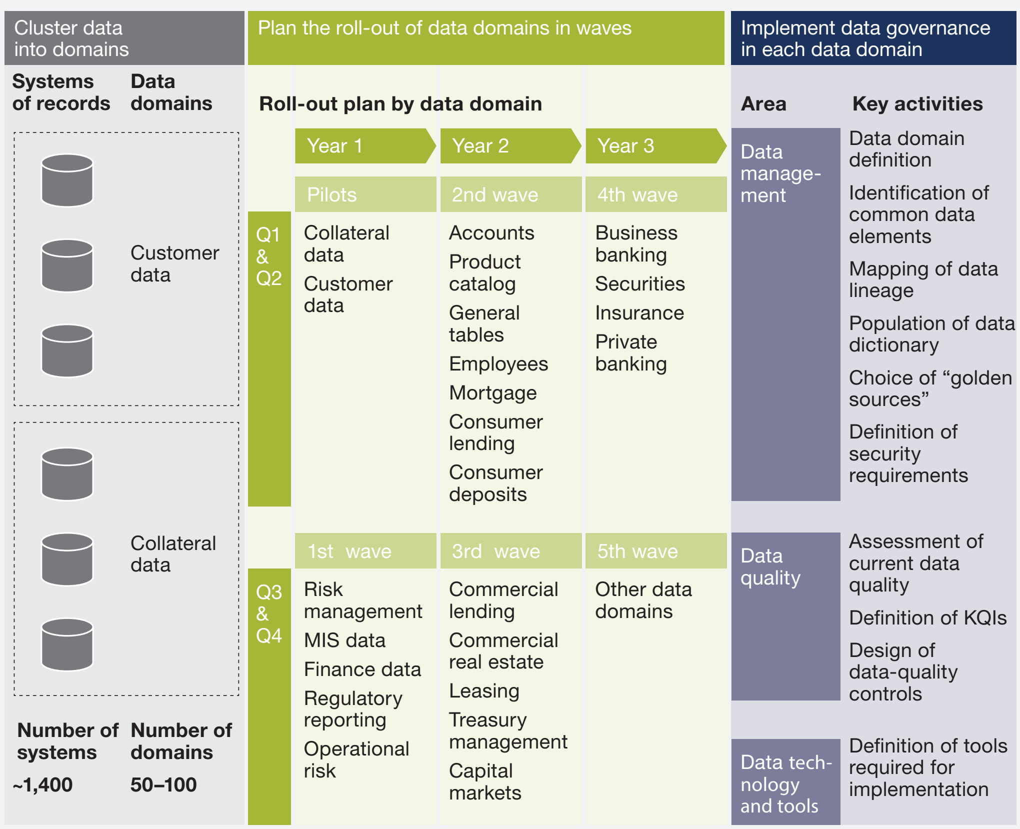

Source: McKinsey analysis

approach means developing target data architecture and data governance only when it is needed for a specific use case. One European bank implemented this approach in three steps (Exhibit 4):

First, it identified the data it needed for key use cases and prioritized those data domains that included it. Typically, 20 percent of data enables 80 percent of use cases. Second, the bank developed a rollout plan for implementing data architecture and governance in three to four data domains per quarter.

Third, the bank set up a cross-functional team for each data domain, comprising data stewards, metadata experts, data-quality experts, data architects, data engineers, and

Exhibit 5

Banks need new roles to compete effectively in a data-driven market.

Senior executive	Translator	Data scientist	Data owner	Head of data governance	Data quality manager	Data technology manager
Digital culture						
Design and agile thinking	Design and agile thinking	Design and agile thinking	Fundamentals of data management	Fundamentals of data governance	Fundamentals of data quality	Fundamentals of data technology tools
Use-case reflections	Source of value	Source of value	Data culture	Data management	Data management	Data management
	Best practices in data engineering	Best practices in testing and piloting	Data quality		Data quality tools	Data modeling
	Best practices in data management	Technical leadership program				Data design
	Best practices in data modeling	Data science				Data software ecosystem
	Technical leadership program	Advanced analytics				
	"Train the trainer“ approach					

Source: McKinsey analysis

platform engineers. Before data was ingested into the data lake, these teams worked to identify key data elements, select golden sources, assess data quality, carry out data cleansing, populate the data dictionary, and map data lineage. Each team worked in agile sprints in a startup-like environment for three to four months. A central team took care of value assurance and defined common standards, tools, and policies.

This approach delivered numerous benefits for the bank, including rapid implementation, capability building, and the creation of tangible business value at every stage in the journey. During any transformation, calling out and celebrating such achievements is critical. As the CDO of JPMorgan Chase, Rob Casper, observed, "The thing that achieves buy-in and builds momentum better than anything is success . . . trying to deliver in small chunks incrementally and giving people a taste of that success [is] a very powerful motivator."

More broadly, senior executives need to champion their data transformation to encourage widespread buy-in, as well as role-modeling the cultural and mindset changes they wish to see. Formal governance and performance-management systems, mechanisms, and incentives will need to be rethought to support new ways of working. At the same time, most organizations will need to develop new capabilities; only 20 percent of the banks we surveyed believe they already have adequate capabilities in place. Given the scarcity of external talent, in particular for key roles such as business translators, organizations will need to provide on-the-job training for employees involved in the transformation, and complement this effort with a data and analytics academy to build deep expertise in specialist roles (Exhibit 5, page 58).

Actions: Adopt a use-case approach to the whole journey; establish central governance to ensure cross-functional working, the use of standard methods, and clear role definition; build new data capabilities through hiring and in-house training.

* * *

In the past few years data has been established as a fundamental source of business value. Every financial institution now competes in a world characterized by enormous data sets, stringent regulation, and frequent business disruptions as innovative ecosystems emerge to break down the barriers between and across industries. In this context, a data transformation is a means not only to achieve short-term results, but also to embed data in the organization for long-term success.

Chiara Brocchi is an expert in McKinsey's Milan office, where **Davide Grande** is a partner. **Kayvaun Rowshankish** is a partner and **Allen Weinberg** is a senior partner, both in the New York office. **Tamim Saleh** is a senior partner in the London office.

Building an effective analytics organization

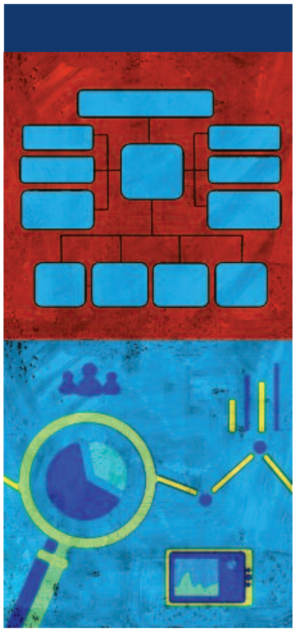

As companies recognize the predictive power of advanced analytics, many are hoping to use AA to drive their business decisions and strategies. While most companies understand the importance of analytics and have adopted common best practices, fewer than 20 percent, according to a recent McKinsey survey, have maximized the potential and achieved AA at scale.

Gloria Macias-Lizaso

In working with a wide range of organizations, McKinsey has seen many companies start their analytics journey eagerly, but without a clear strategy. As a result, their efforts often end up as small pilots that fail to scale or have significant impact. Some of these pilots have been mere exercises in "intellectual curiosity" rather than a serious effort to change the business. Consequently, they are not designed with an end-to-end approach that incorporates the necessary conditions for implementation. Instead, the pilots are carried out in small labs with limited connection to the business, and fail to provide the answers the business needs to move forward. Even if a pilot does answer the right questions, it may not address the cultural aspects that would, for example, make a sales representative trust a model more than her own experience.

These companies quickly become frustrated when they see their efforts falling short while more analytically driven companies are leveraging their data. Democratization of data is blurring sector boundaries; businesses will increasingly find themselves disrupted not by the company they have been monitoring for the last several years, but by a newcomer from another industry. Being the best in an industry is no longer enough; now companies must aspire to be at least at par across industries to compete effectively. Functional expertise, beyond specific sector expertise, will become more and more relevant.

With this in mind, McKinsey conducted an extensive, primary research survey of over 1,000 organizations across industries and geographies to understand how organizations convert AA insights into impact, and how companies have been able to scale analytics across their enterprise (see sidebar on page 61). In this article, we will discuss how to design, implement, and develop the right organization and talent for an AA transformation. An AA transformation usually requires new skills, new roles, and new organizational structures.

Building an AA-driven organization

Top-performing organizations in AA are enabled by deep functional expertise, strategic partnerships, and a clear center of gravity for organizing analytics talent. These companies' organizations usually include an ecosystem of partners that enables access to data and technology and fosters the co-development of analytics capabilities, as well as the breadth and depth of talent required for a robust program of AA.

For a company aspiring to an AA transformation, these elements can be incorporated into any of several organizational models, each of which is effective as long as there is clear governance, and the company encourages an an-

McKinsey's Insights to Outcome Survey

In the fall of 2017, McKinsey performed quantitative research (using a survey-based approach) of approximately 1,000 organizations across industries and geographies. The survey contained 36 questions, most of which measured respondents' degree of agreement or asked respondents to choose their top three responses. The 1,000 responses encompassed more than 60 responses per geography and over 50 responses per industry, which ensured statistical relevance in various cuts of the data. The responding companies represent more than $1 billion in revenues.

The survey targeted analytics leaders and C-level executives with a broad perspective on their organization's analytics capabilities across the enterprise. These respondents included 530 individuals in analytics roles and 470 in business roles.

The industries covered by the survey included: A&D, automotive, banking, insurance, energy (including oil and gas), resources (including mining and utilities), telecom, high tech, consumer, retail, healthcare, pharmaceuticals, transportation, and travel. The geographies covered included: US, UK, France, Germany, Spain, Brazil, India, Australia, New Zealand, Singapore, China, Japan, and the Nordics.

alytical culture across business units to learn and develop together. Answering a few key questions can help to identify the best model.

1. **Centralized, decentralized, or a hybrid:** First, the company should decide whether to create one centralized AA organization, in which AA stands alone in a center of excellence (COE) that supports the various business units; a decentralized organization, in which analytics is embedded in individual businesses; or a hybrid, which combines a centralized analytics unit with embedded analytics areas in some units.

   Our benchmark of several organizations indicates that any of these models can work effectively, as long as governance is established to prevent the various units from becoming islands. The proposed organization depends somewhat on how advanced the company and the business units are in their use of analytics.

   It is important to note that any organization will change over time as the AA transformation evolves. Some companies start out decentralized and eventually move AA into a centralized function, while others that are centralized later move into a hybrid model of hubs and spokes. Top-performing companies prepare for these eventual changes.

One organizational example

A large financial and industrial conglomerate created a separate COE that reports directly to the CEO and supports the organization with AA expertise, AA resources (on "loan"), use case delivery, infrastructure to execute use cases, and technical interviewing. The center also manages data partnerships, develops new businesses by designing and deploying cross-company and ecosystem use cases on the company's own infrastructure, facilitates aggregated AA impact calculation, reports progress to the executive committee, and executes the data committee's mandates. The center started out as a small cost center but aspires to transform into a self-standing profit center within two years.

Two successful strategies

One industry conglomerate addressed this scale requirement by starting with a centralized COE serving all business units. As the use and understanding of analytics grew across the organization's companies, they demanded more support, and the COE was split into sub-groups that were fully dedicated to the largest companies. Over time, ownership of these groups was transferred to the "client" company—but not until they had built a sense of community and common methodology across the entire conglomerate. This sense of community was further reinforced by requiring all new recruits to spend six months at the COE and to go through specific AA training and networking events. Since fragmentation of the analytical talent across functions is almost inevitable over time, it is critical to start out with the appropriate processes and mechanisms to ensure consistency and community across these new profiles.

A leading pharmaceutical company developed an integrated talent strategy that merged business and analytics functions. The company recruited technology and analytics executives in key management roles and developed analytics career paths for them. Placing analytics professionals in key business roles enabled the company to identify and operationalize new analytics opportunities before their competitors could. The organization successfully embedded analytics in key elements of the business—for example, analytics on clinical trial data to enable more cost-effective data.

The choice between centralization and decentralization is not an all-or-nothing decision but should be decided per sub-function. Data governance, however, should be centralized, even if data ownership is not. For data architecture, top-performing companies often have data centralized within business units. This data typically includes data from marketing, sales, operations, and so on. Most top-performing companies centralize partnership management; otherwise, competing or redundant partnerships could inadvertently be set up in various parts of the organization, and intellectual property could be at risk.

2. **To outsource or not to outsource:** Another decision is whether AA talent should be partially outsourced, and if so, how. Should outsourcing be limited to low-level data analytics activities? Or should the company establish several tactical partnerships for selected tasks? Or would a strategic partnership with an external vendor be the best approach? AA will effectively become the "brain" of the organization, so companies should be careful not to outsource too much. Top-performing companies often keep analytics that provide a competitive advantage—such as pricing analytics—within the organization. A central, internal unit can oversee all AA outsourcing, and partnerships can be established for specific AA solutions or to bring in particular assets, such as unique sources of data or advanced solutions.

3. **Locating the AA unit:** Yet another important decision is where to locate the AA unit. AA is most effective when it is cross-functional, accessible enterprise-wide, and integrated with the business. Various levels and functions can host it, but the final location should have enough visibility and access to the C-suite to break through inertia and enable transformation. It is helpful if the unit has an enterprise-wide view, given its transformational potential for all functions.

The AA unit is often most effective when it is a sub-unit of business intelligence—as long as this area has an enterprise-wide perspective—or of strategy or digital. Some companies locate their AA units in IT, but this arrangement can be challenging. IT staff—who are used to managing longer-term projects that are often disconnected from the business—may not be prepared to manage short-term, agile AA projects. AA projects can end up last on their list of priorities. Including AA within marketing or operations, meanwhile, can limit its potential to transform the remaining parts of the organization.

Staffing the AA center of excellence

Sixty percent of top-performing companies in AA have a "center of gravity" for their analytics efforts, according to our survey. They typically include a specific set of roles, skills, and capabilities within the COE (Exhibit 1),

Exhibit 1

An organizational blueprint of the advanced analytics COE.

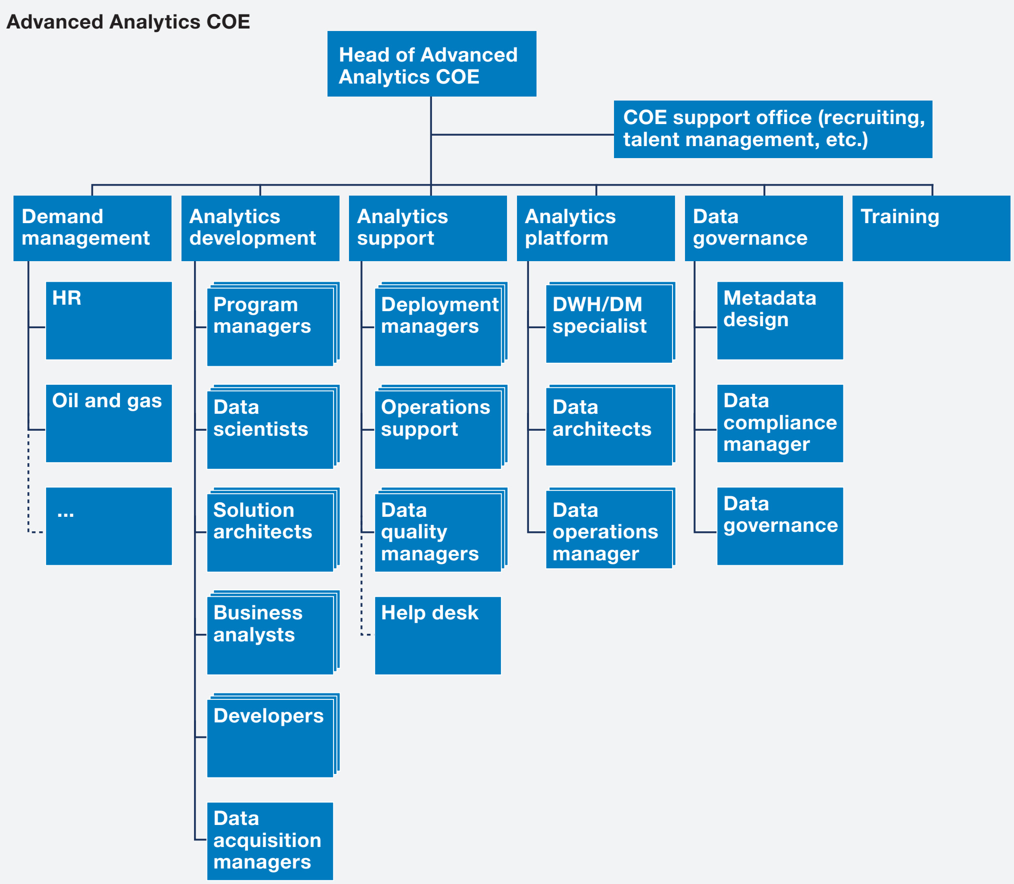

Source: McKinsey analysis

Lessons learned from COE failures

Companies that have rolled out full-scale COEs during an AA transformation have encountered some pitfalls. Some of the most common include:

1. **A failure to focus on business value.** Successful roll-out of a COE requires a clear vision as to where the COE's biggest business impact can be captured. The COE should be linked to the organization's overall business strategy and should commit to achieving a measurable impact during a defined time frame, with its goals clearly prioritized.

2. **Over- and under-thinking technology.** Some companies attempt to replicate an entire data history when building their AA COE, resulting in data "scope creep." It is vital to sufficiently think through data integrity/architecture; failing to do so may result in missing data and missing data connections. Some companies try to do too much at once by replacing their hardware, software, and analytics stack simultaneously rather than tackling one at a time. These companies may buy the "best of breed" in each category but then find that none of them "talks" to each other. Instead, companies should build systems and functionality as needed—especially since technologies tend to become obsolete within just a couple of years. New innovations can be integrated later if the system is built gradually.

3. **Taking more than 18 months to deliver value.** The COE's benefits should begin to come online well before the entire roll-out is complete. If the COE does not deliver benefits sooner, it is often because it depends too heavily on insights-delivery FTEs instead of automation. The delivery of insights should be staged to capture value sooner.

4. **Insufficient skill-building and change management.** Top management and the internal team must be 100 percent committed to the COE if it is to succeed. Internal stakeholders must be engaged in development or accountable for delivery. The effort cannot focus exclusively on technology instead of the process and people; in particular, the organization must build the requisite front-line skills needed for an effective AA COE. Companies should hew closely to the business case and avoid the functional scope creep that can occur when mid-flight changes not included in the business case suddenly become priorities.

5. **Operating analytics as an island.** One large US insurance company interviewed by McKinsey hired a sizeable number of data scientists and launched more than 50 pilot projects to test its new capabilities. Despite a real commitment and considerable investment, the analytics team was isolated from the rest of the company, with no connection to the overall business strategy—a critical mistake. Not surprisingly, this company's ad hoc analytics projects had no real impact.

At the other end of the spectrum, successful AA-driven companies are building centralized AA capabilities and then creating end-to-end agile teams ("use case factories") that integrate profiles from IT, sales, marketing, finance, and other functions. This approach ensures that use cases are immediately integrated into business processes and thus create value.

including data scientists ("quants"), data engineers, workflow integrators, data architects, delivery managers, visualization analysts, and, most critically, translators from the business who act as a bridge between the COE and business units. The translators usually have a combination of business, analytics, and technology skills and are found in the business partner role in data analytics leadership.

Many COE roles are filled with highly specialized analytical resources recruited from advanced degree programs in computer science or math. But these individuals must also be able to translate sophisticated models into simple, visual decision support tools for front-line employees.

They also need to have a collaborative mindset, given the interdependencies among data, systems, and models. With translators bridging any communication gaps, team members from analytics and the business work together in two- to three-month agile "sprints" as they identify problems; find out whether relevant data exists and, if not, whether that data can be acquired; test their models; determine how those models will be put into production; and learn from the results.

The COE can be built in about 18 months, typically in incremental steps. It may start with five to ten data professionals, including data engineers, data scientists, and translators. In its end-state, it likely will require significantly more. The number of translators needed will vary by business unit but is generally about 10 percent of business unit staff. Most companies source their translators from "client" business units and then train them, since these employees will have deep knowledge of the processes that AA is trying to optimize. These individuals are usually analytical, critical thinkers who are well respected in the company.

Exhibit 2

Beyond the COE, employees at all levels need to be trained.

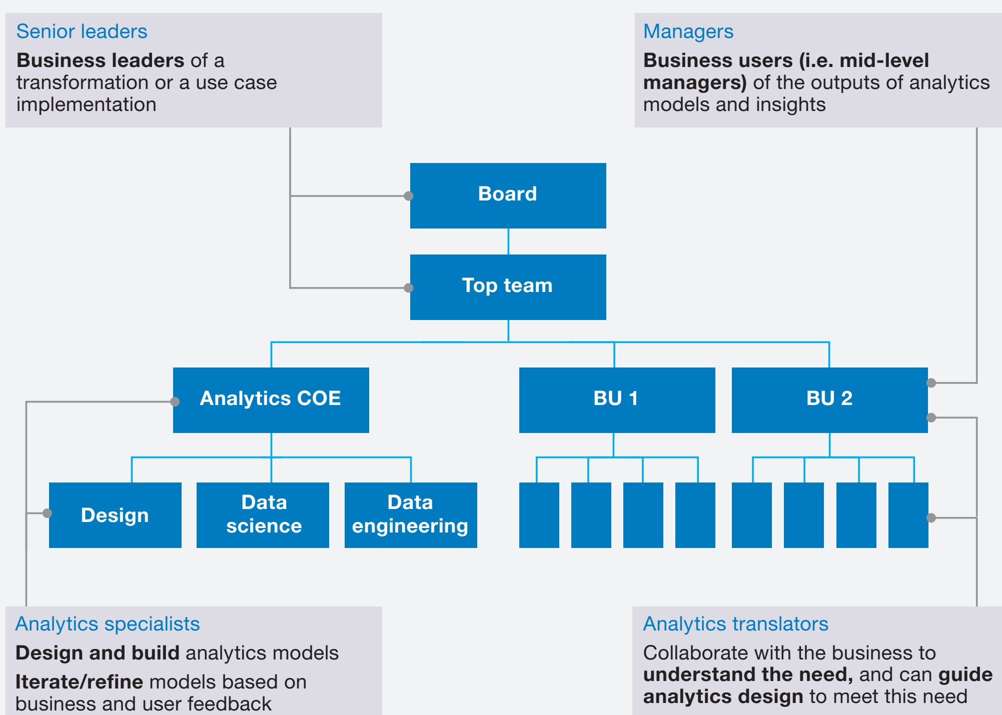

Source: McKinsey analysis

To illustrate how the various key skills and roles come together in the COE, here is an example description of these roles' working together to fulfill a business request:

- The translator and business owner identify and prioritize the business request.
- The data scientist works with the translator to develop an analytics use case, including an algorithm and analyses to test.
- A data engineer from the COE works with the relevant business division to understand the data requirements of the use case and to identify data sources.
- The data engineer works with IT/the business to ensure data availability, identify gaps, and develop ETL (extract, transform, load) to load data into analytics sandbox.
- A data scientist programs the algorithm and analyzes the data in the sandbox to generate insights.
- A visualization analyst develops reports and dashboards for business users.
- A COE workflow integrator works with the business owner to develop a prototype for models and tools.
- The COE ensures that key business and IT stakeholders test the prototype tools and solutions.
- A delivery manager pilots the prototype and dashboard and works to obtain a go/no-go decision.
- The delivery manager and COE workflow integrator work with IT to scale the prototype to the enterprise level.
- The COE delivery team and translator work with the business and IT to ensure adoption and ongoing model maintenance.

In this process, feedback would be gathered between steps nine and ten.

While the COE and some of its roles may emerge gradually, it is best to have the data, platform, and career paths needed for an AA transformation in place from the beginning. If the platform is still under development, adding more people may only make that development more complicated. And without a clear career path, attracting this scarce talent will be difficult. As much as possible, roles should be clearly delineated to prevent squandering valuable talent on functions for which they are over-qualified, which can undermine retention.

Career development and strategic partnerships

Gaining an edge in analytics requires attracting, retaining, and sourcing the right talent. In McKinsey's survey, 58 percent of respondents at top-performing companies say that their organization has deep functional expertise across data science, data engineering, data architecture, and analytics transformation. Top-performing organizations have four times as many analytics professionals and one and a half times more functional experts than other companies.

These companies also retain three times more talent—primarily by creating strong career development opportunities. People with superior analytics talent usually have many potential opportunities and thus need to see a clear career path and opportunities for growth within a company if they are to join or stay with it. Several career tracks should be available, as some analytics staff may wish to pursue a more technical profile, others may move into translator or integrator roles with the business, and some will likely move into managerial positions.

In all cases, these individuals tend to stay motivated if they are learning on the job and from one another. Achieving this goal requires a minimum scale for each analytics group. Having only one or two data scientists in each function will not help them learn, and they may have difficulty making themselves understood.

To fill any gaps in talent, 62 percent of survey respondents at top-performing companies say that they strategically partner with others to gain access to skill, capacity, and innovation. For example, a large, multinational retailer developed a strategic partnership with a start-up incubator that focuses on identifying cutting-edge technologies—such as drones—to transform the retail industry. The retailer found that employing a mix of in-house talent and smart, strategic partnerships with other organizations enabled it to get the best out of both, thus affording access to skills, capacity, and innovation on a much larger scale. Through the incubator, the retailer formed partnerships with start-ups and venture capital investors. The company also created a compelling value proposition for attracting top analytics talent.

Beyond the COE: training employees for cultural change

As detailed in "Hidden figures: The quiet discipline of managing people using data," on page 36, an AA transformation requires a profound cultural change, as the entire organization must change the way it operates. Employees need to learn to trust in AA, to understand what they can ask of it, and to know that AA can answer far more complex questions than traditional analytics ever could. Outside of the COE, then, employees at all levels—senior leaders, managers, analytics specialists, and analytics translators—need to be trained to be AA-proficient and to drive the transformation forward (Exhibit 2, page 65).

A sweeping—but feasible—transformation

Transforming a company to be AA-driven is a monumental task that should not be undertaken in one fell swoop, but instead incrementally, based on use cases. Since AA can and will transform a company, the effort to cultivate an AA-driven organization is most effective when it comes from the top, from senior executives. If a company focuses on the value of advanced analytics and builds AA capabilities as needed—while still having the data, platform, and talent strategy in place from the beginning—its AA transformation will succeed.

Gloria Macias-Lizaso Miranda is a partner in McKinsey's Madrid office.

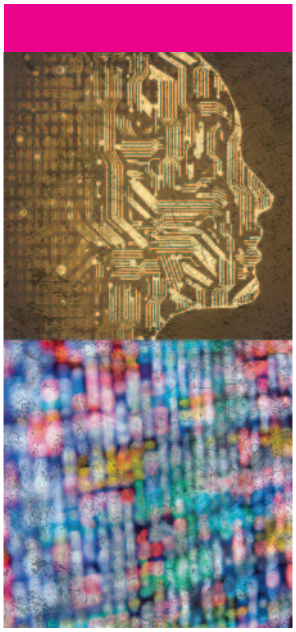

"All in the mind": Harnessing psychology and analytics to counter bias and reduce risk

The management of risk in financial services is about to be transformed. A recent McKinsey paper identified six structural trends that will reshape the function in the next decade. Five are familiar—they concern regulation, costs, customer expectations, analytics, and digitization—but one is less so: debiasing. That means using insights from psychology and behavioral economics, combined with advanced analytical methods, to take the bias out of risk decisions. The institutions pioneering this approach have seen tremendous benefits: for instance, banks adopting psychological interventions in consumer collections have achieved a 20 to 30 percent increase in the amount collected.[1]

Tobias Baer
Vijay D'Silva

The interest in debiasing is growing as psychological research uncovers more and more subconscious effects that influence our decision making. Meanwhile, an explosion in data availability is providing businesses with an abundant flow of information for their analytic engines. Not all the data theoretically available can be exploited, for legal and privacy as well as technical reasons. But institutions still have a massive amount of underused data that they can mine, using an increasingly sophisticated array of advanced analytics techniques, to develop behavioral segmentations and predictive models. With these foundations in place, they can go on to design powerful interventions to tackle bias.

Take the example of a bank using a recursive neural network to extract customer profiles from credit-card transaction data. One profile that emerges is of a cardholder who clocks up dozens of low-value transactions at a convenience store every week. The customer's habit of making multiple repeat visits at odd hours—seemingly for only one or two items at a time—suggests a lack of forward planning. Seen through a psychometric lens, the customer seems to be exhibiting poor impulse control and a lack of conscientiousness, traits that are likely to determine which types of decision bias this customer can be expected to manifest.

Compare this profile with that of a cardholder who completes one big supermarket transaction at more or less the same time every Friday evening, with little or no evidence of convenience-store shopping in between. That profile is indicative of a well-organized person who plans ahead. It's likely that the first customer would benefit from financial products designed to help customers who struggle to meet their financial obligations— such as a credit card with weekly rather than monthly payment installments—whereas the second customer would probably have no need of them. And if, say, the bank is considering ways to motivate cardholders to pay off delinquent credit, its knowledge that customers with the first cardholder's profile are likely to prioritize immediate consumption over clearing their debts will help it design suitable incentives to counter this tendency.

Analytics-driven psychological insights like these can be a spur to tremendous value cre-

[1] For a comprehensive discussion of the psychological levers that can be used to improve performance in consumer debt collection, see Tobias Baer, "Behavioral insights and innovative treatments in collections," *McKinsey on Risk*, Number 5, March 2018.

ation. This article considers some of the most common biases in business decision making and looks in detail at three areas where debiasing can reap rich rewards: credit underwriting, consumer debt collection, and asset management.

Uncovering biases in business

Biases are predispositions of a psychological, sociological, or even physiological nature that can influence our decision making (see sidebar, "A quick guide to common biases"). They often operate subconsciously, outside the logical processes that we like to believe govern our decisions. They are frequently regarded as flaws, but this is both wrong and unfortunate. It's wrong because biases are an inevitable side-effect of the mechanics our brains need to achieve their astonishing speed and efficiency in making tens of thousands of decisions a day. And it's unfortunate because the negative perception of biases leads us to believe we are immune to them—a bias in itself, known as overconfidence, exhibited by the 93 percent of US drivers who believe themselves to be among the nation's top 50 percent.

Even if we accept that biases may influence our decisions, we might assume that successful organizations have developed processes to keep them in check. But experience indicates otherwise. For example, academic research has found that ego depletion materially affects the work of judges, doctors, and crime investigators, and our own research has revealed how it affects credit officers' decisions, manifesting itself in tangible business metrics such as credit approval rates. When financial institutions work to counter bias in judgmental underwriting—in small business credit, for example—they can typically cut

A quick guide to common biases

Heuristic biases are computational shortcuts taken by the brain to achieve lightning-fast, almost effortless decision making. Thanks to the Nobel Prize–winning work of Daniel Kahneman and Richard Thaler, these biases have become more widely understood in recent years. More than a hundred have been identified, ranging from the relatively familiar loss aversion to the less well-known Hawthorne effect. For practical business purposes, five groups of biases are key:

- **Action-oriented biases** prompt us to act with less forethought than is logically necessary or prudent. They include *excessive optimism* about outcomes and the tendency to underestimate the likelihood of negative results; *overconfidence* in our own or our group's ability to affect the future; and competitor neglect, the tendency to disregard or underestimate the response of competitors.

- **Interest biases** arise where incentives within an organization or project come into conflict. They include misaligned individual incentives, unwarranted emotional attachments to elements of the business (such as legacy products), and differing perceptions of corporate goals, such as how much weight to assign to particular objectives.

- **Pattern-recognition biases** cause us to see nonexistent patterns in information. They include *confirmation bias*, in which we overvalue evidence that supports a favored belief and discount evidence to the contrary; availability bias, in which we misperceive likelihoods of events because we recall one type of event much more easily (and hence frequently) than others; management by example, in which we rely unduly on our own experiences when making decisions; and *false analogies*, faulty thinking based on incorrect perceptions and the treatment of dissimilar things as similar.

- **Stability biases** predispose us toward inertia in an uncertain environment. They include *anchoring without sufficient adjustment*, in which we tie actions to an initial value but fail to adjust when new information becomes available; loss aversion, the fear that makes us more risk-averse than logic would dictate; the *sunk-cost fallacy*, where our future course of action is influenced by the unrecoverable costs of the past; and *status-quo bias*, the preference for keeping things as they are when there is no immediate pressure to change.

- **Social biases** arise from our preference for harmony over conflict, or even constructive challenge. They include *groupthink*, in which the desire for consensus prevents us making a realistic appraisal of alternative courses of action, and *sunflower management*, the tendency for group members to fall into line with their leaders' views.

For all their importance, however, heuristic biases represent only the tip of the iceberg as far as subconscious influences on our decisions are concerned. Exhibit A illustrates other factors that lie deep below the surface. Somatic and emotional effects tinker with the parameterization of our brain, and can be triggered by factors as diverse as blood-sugar level, smells, or mood: for instance, if our blood sugar is low, we (quite reasonably) estimate that completing a given task, such as climbing a mountain, will take us longer. Ego depletion, a form of mental fatigue, leads us to move from logical thinking to unconscious short-cuts that favor easy default decisions. And group psychological effects override rational decision making out of a deep-seated fear of ostracism.

Exhibit A

Psychologists and neuroscientists have discovered many forces that cause decisions to gravitate away from logical considerations.

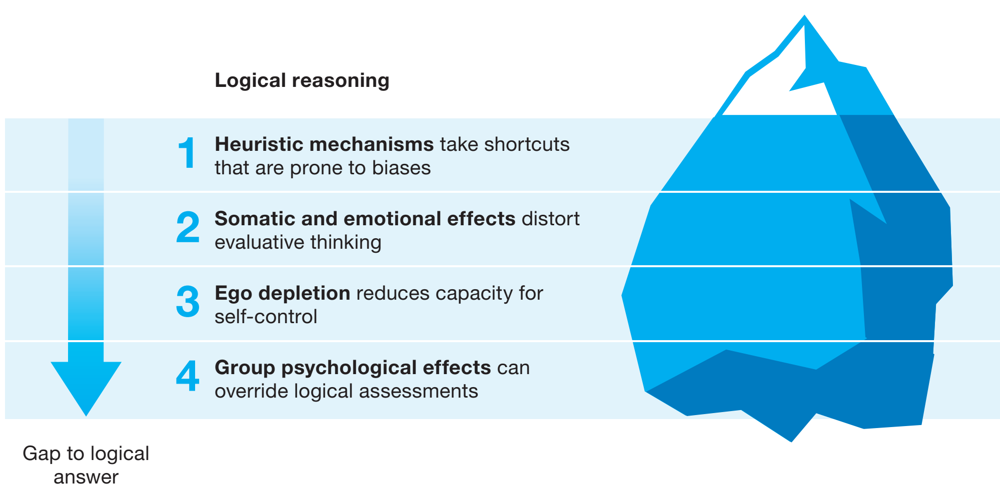

Source: McKinsey analysis

credit losses by at least 25 percent, and even as much as 57 percent in one case.

For lenders, an area particularly ripe for debiasing is debt collections, where biases can shape the behavior of collectors and customers alike. Consider how collectors handle calls with recalcitrant customers. Over the course of a call, they need to make numerous split-second decisions that expose them to the full gamut of biases, such as anchoring and over-optimism, as well as somatic effects and ego depletion. Whether they persist in trying to elicit a promise to pay or give up and move on to the next delinquent account may partly depend on the time of day. The effectiveness of collectors' calls dwindles over the course of the working day as ego depletion sets in (Exhibit 1). The good news is that companies aware of this phenomenon can make adjustments in collectors' working environment to help counter it.

And when it comes to customers with overdue accounts, leading financial institutions are harnessing a plethora of psychological insights to encourage payment. This often means making targeted interventions that increase customers' motivation to pay, help those with low self-control to keep their commitments, and respect individuals' need for agency (and thereby avoid triggering what psychologists call "reactance"). A credit-card provider could, for instance, present high-

Exhibit 1

Call effectiveness dwindles over the course of the day through ego depletion as collectors tire.

Case example; percent of calls eliciting promise to pay, by time of day

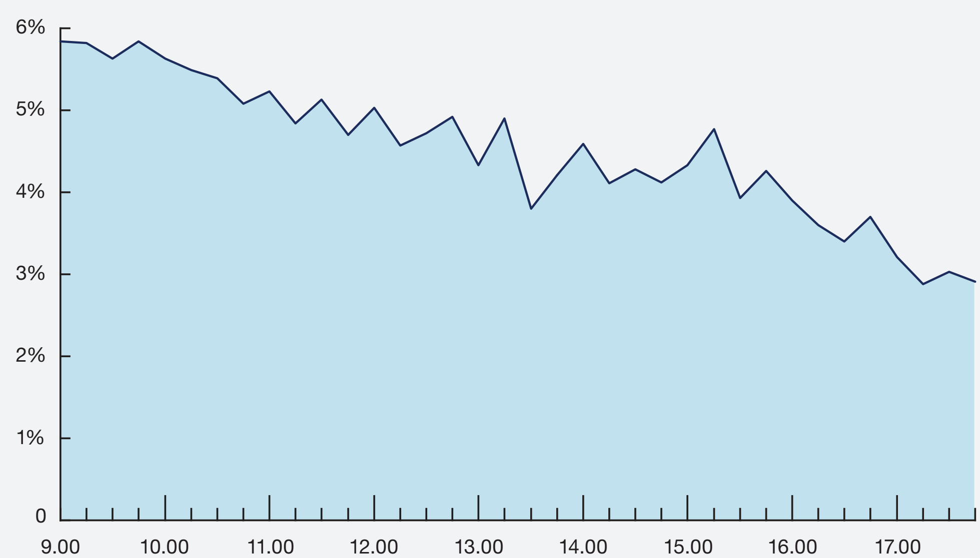

Source: McKinsey analysis

How can financial institutions tackle biases?

The questions organizations need to consider include:

The decision type: High or low frequency, formal or informal

Formal high-frequency decisions, such as credit underwriting or standard manual fraud checks, lend themselves to analytical solutions coupled with "industrial strength" psychological interventions. For example, a bank that usually asks about the frequency of CFO changes in the past three years—a question that may be susceptible to the availability bias—could instead design a simple table prompting credit officers to construct a timeline for pertinent data points.

Formal low-frequency decisions—such as approvals of new lending products or a credit committee's quarterly recalibration of the PD rating model that drives underwriting and risk-based pricing—call for decision processes to be redesigned to support logical thinking and ensure adequate challenge. Analytical modeling is often helpful here. One US bank used four different econometric models to produce four distinct default-rate forecasts in an elegant effort to counteract groupthink and introduce automated "devil's advocates" into the discussion.

Informal decisions, such as a supervisor's override on a policy violation, may first have to be formalized before any intervention can be deployed. A review of historical losses may shed light on a few decision types that warrant such an investment, such as debt collectors' decisions to give up on difficult accounts. If a bank wants a collector to spend longer than usual on a call to a particular customer, for instance, it could flag up an above-average incentive payment in a pop-up on the collector's screen.

Who to target and how

Institutions need to use behavioral segmentation to distinguish which groups are affected by which primary biases, and which personality traits determine the choice of countermeasure. In consumer debt collection, for instance, the psychological need for agency can cause customers to resist resolution if they feel they have been put on the spot by a call from an assertive collector. An invitation to restructure the debt on a self-service website could effectively overcome this bias. However, this same approach could be disastrous if used to deal with a customer who is biased toward avoidance.

The role of automation

Carefully designed algorithms can not only speed up decisions and take out costs, but also remove biases from a growing range of decision types. But financial institutions must beware of a major trap: building past biases into the algorithm.

risk customers with a late-fee waiver or a gift card from a favorite shop that they would lose if they didn't make a payment. Framing the offer as a loss for a payment missed, rather than a reward for a payment made, enlists the help of the loss aversion bias and can double the effectiveness of the offer.

Before deciding where and how to use behavioral levers, financial institutions need to consider a range of factors (see sidebar, "How can financial institutions tackle biases?", page 72).

To give a sense of what can be achieved when these techniques are applied in practice, let's now examine what leading institutions have been doing to take bias out of credit underwriting and consumer debt collection. And looking beyond lending, the sidebar "Debiasing asset management" (page 76) describes how firms in an adjacent industry uncovered bias in their investment decisions.

Commercial credit underwriting

Most credit officers possess a strong professional ethic and have honed their skills over years, if not decades. Yet evidence indicates they are just as susceptible as anyone else to decision bias.

One bank with poor performance in its commercial credit underwriting made a retrospective assessment of the predictive value of its judgmental credit ratings using Gini coefficient measures on a scale from 0 (no predictive power) to 100 (perfect prediction). The analysis examined 20 dimensions stipulated by the bank's credit policy, such as management quality and account conduct, and compared judgments made by credit officers with actual defaults observed over the following 12 months. One dimension (account conduct) stood out with a relatively high Gini of 45, but most dimensions had much lower scores (Exhibit 2). By way of comparison, comprehensive best-practice models for rating small businesses can achieve a Gini of 60–75.

In fact, half of the dimensions in the bank's rating model achieved a Gini score of 7 or lower—little better than a roll of the dice—yet the bank was paying them just as much attention as it gave to dimensions with genuine predictive power. For instance, despite scoring a Gini of just 1 in back-testing, shareholder composition was usually discussed in depth in credit memos, and relationship managers were even prompted to ask customers follow-up questions about it. Factoring in such irrelevant dimensions anchored credit officers' overall rating in randomness, dragging it down to a Gini of just 22.

In order to debias its commercial underwriting, the bank had to separate the wheat from the chaff—a systematic process combining analytics with psychological insights. First, the bank replaced fuzzy concepts with carefully chosen sets of proxies for which more objective assessments could be developed. Eliminating factors that were irrelevant, or impossible to assess without crippling bias, would substantially improve the overall credit rating. Second, explicit psychological "guard rails," such as the use of tables to prompt credit officers to plot data along a timeline rather than relying on a customer's spontaneous recall of events, were put in place to safeguard qualitative assessment processes from biases.

Finally, the bank used statistical techniques to validate each redesigned factor and calibrate its weight. As is common in commercial credit portfolios, the bank ran up against the problem of a small sample size. This was compounded by the need to compile additional data manually for the sample used in developing the new assessment, which comprised just 30 to 50 defaulters and the same number of performing debtors. Although such constraints ruled out the statistical techniques most commonly used in credit scoring, such as logistic regression, the bank was able to deploy powerful statistical concepts from social science instead, such as Cohen's d and t-test.

The bank has now been using its qualitative credit rating, with minimal annual adjustments, for more than a decade, scoring an overall Gini between 60 and 80 every year, even during the financial crisis.

Consumer collections

A recent McKinsey survey of 420 US consumers with credit delinquencies sheds light on some of the decision biases that contribute to non-payment. For instance, many consumers are unable to resist the temptation of immediate consumption—an example of what's known as "hyperbolic discounting"—and so they struggle to manage money through a monthly cycle. A third of those surveyed expressed a preference for a schedule that would allow payment every week or every other week, either because it would fit better with their paydays or because smaller, more

Exhibit 2

One bank's credit-rating model contained factors that were subject to bias or had little or no predictive power.

Case example; predictive power of judgmental ratings assigned by credit officers, measured by Gini coefficient: 0 = useless (random), 100 = perfect prediction

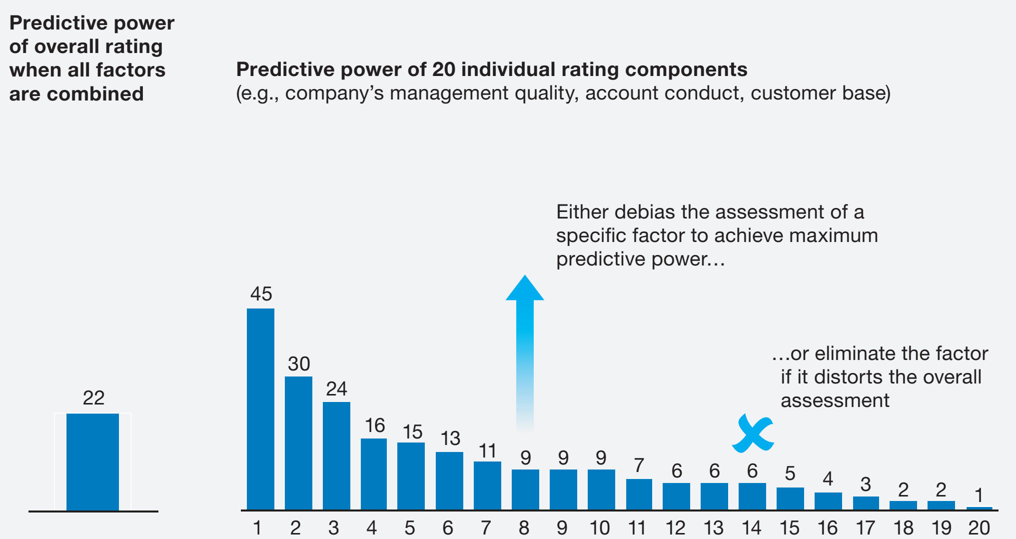

To debias credit memos, institutions can replace lengthy prose with concise questions, multiple-choice options, and simple tables—which will also streamline assessment, cut costs, and speed up turnaround

Source: McKinsey analysis

frequent payments would be less painful and easier to manage than monthly bills.

Understanding how consumers decide what to pay and when is particularly important when they owe money to more than one lender. Only a third of survey respondents prioritized payments rationally by, say, tackling debts with the highest interest rate first, or seeking to retain their most useful credit card. The remaining two-thirds followed less rational patterns: some apportioned payments equally, others showed loyalty to a particular bank, and yet others paid off the smallest balance first (Exhibit 3). Banks that are aware of such motivations can either reinforce them with tailored payment plans or help customers adjust their rationales—for instance, by breaking down large balances into smaller chunks or milestones.

Some leading banks are putting behavioral targeting into practice by applying psychometrics: the factual scoring of a customer's personality profile according to a framework such as the widely used OCEAN Big Five. Such a profile allows banks to micro-target marketing messages not only in origination—

Exhibit 3

Research into consumers with credit delinquencies yielded valuable behavioral insights.

Survey of 420 US consumers who have been at least one month overdue

What payment frequency do you prefer?

%

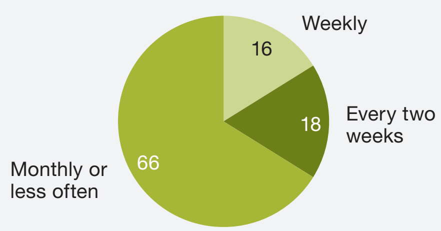

Comments and insights

"I always prefer to pay smaller payments more frequently because it takes the sting out of making a payment. Making a large payment always feels like a punch."

When several accounts are overdue, which do you pay first?[1]

%

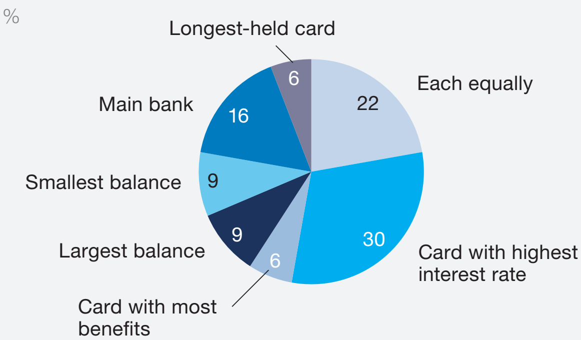

Comments and insights

20% of respondents said they have **withheld a planned payment** because of an upsetting call from a collector

38% of respondents had a very positive experience with at least one collector who was **empathetic and genuinely helpful**

[1] Figures do not sum to 100% because of rounding

Source: McKinsey survey

Debiasing asset management

Few industries have subjected their investment decision-making processes to more scrutiny than asset management, yet biases still affect many high-value decisions throughout the lifecycle of individual funds. In the early stages of structuring a new fund's strategy and processes, for instance, stability biases can influence whether an index or some other means is chosen for assessing performance. Interest biases, such as misaligned incentives, need to be monitored to ensure that the long-term interests of unit holders and asset owners are taken into account when funds are managed and promoted.

Leading asset management organizations are becoming increasingly alert to the impact of decision-making biases on fund performance too. A few have adopted an innovative approach to diagnosing bias and its drivers. Working with analytics experts and behavioral scientists, they have applied machine-learning algorithms to their own historical data and discovered clusters of suboptimal investment decisions. Having examined these decisions more closely, they have detected signs of consistent bias in the processes by which the decisions were reached.

When one such organization analyzed its trades, processes, and associated emotions for signs of bias, it found that more than 35 percent of fund managers' decisions were influenced by biases such as loss aversion, anchoring, and what's known as the "endowment effect," in which we attach more value to items that we own. Dan Ariely, a behavioral economist and the best-selling author of Predictably Irrational, notes that this effect kicks in when individuals fall in love with what they already have and focus on what they may lose rather than what they may gain. Such a sentiment can drive fund managers to hold on to stocks for too long and ignore better investment opportunities elsewhere—a trap into which many seasoned investors have fallen.

In one of the funds that this organization examined, the endowment effect had led one fund manager to hold on to 20 percent of positions for too long. The stocks affected had underperformed the relevant index by an average of 25 percent in the 12 months prior to exit. The fund manager acknowledged that he had paid insufficient attention to these stocks, had not rated them as performing badly enough in absolute terms to divest, and could have tried harder to identify better investment opportunities. He admitted that if he had asked himself from time to time whether he would still buy the stocks today, he would have been unlikely to hold on to them for so long. In this case, the value left on the table as a result of the endowment effect was equivalent to 250 to 300 basis points per year.

And this fund manager is not alone. According to Cabot research, institutional investors lose an average of 100 basis points in performance a year as a result of the endowment effect—or 250 basis points in the case of the 10 percent of most-affected funds.

A typical debiasing process is a learning exercise for an asset management fund. By exposing patterns of bias with the help of analytics and then selecting and applying debiasing methods in its investment decisions, the fund will be able to target the specific biases and situations that affect its own investment decisions. From the many interventions available to address every type of bias, it will need to select and customize measures that suit its fund mandate, investment philosophy, team process, culture, and individual personalities.[1]

Magdalena Smith is an expert in McKinsey's London office.

[1] For more on this topic, see Nick Hoffman, Martin Huber and Magdalena Smith, "An analytics approach to debiasing asset-management decisions," McKinsey & Company, December 2017.

choosing the visuals, tag line, and highlighted features to use in a product pitch—but also in debt collection. When applying such an approach, banks often find it helpful to break down a collections episode into four distinct "moments":

1. *Opening*. When the phone rings, customers must decide whether or not to engage with the bank or card provider. If they pick up, they then have to decide whether to take a defensive or evasive stance or to collaborate in problem solving (for example, by disclosing financial difficulties).

2. *Commitment*. Once collaboration has been established, the collector needs to move the customer toward a promise to pay.

3. *Negotiation*. A major part of the conversation will be a negotiation over the customer's financial limitations and the payment to which he or she is willing to commit.

4. *Follow-through*. Finally, the customer needs to keep the promise to pay—a complex decision with ample opportunities for derailment.

At each of these moments, the customer must decide whether or not to cooperate with the lender, and the lender must try to understand the customer's behavior and identify opportunities to increase the likelihood of repayment, using psychological interventions carefully calibrated to each customer's profile.

In the opening moment, a collector who puts a customer in the right mood (or "positive affect") will increase that person's receptiveness to exploring solutions and self-confidence in resolving the situation. Conversely, creating the opposite mood—negative affect—will impede resolution. One approach that institutions have found effective is to use collectors with profiles similar to those of customers, matching regional dialect, gender, and age. Similarly, requesting a call back via email, text message, or app alert instead of calling the customer directly shows respect for an individual's need for agency. Customers too ashamed or anxious to speak on the phone can sometimes be steered to self-service channels through advertisements on social media.

By telling a customer that the solution being offered has been popular with other clients, collectors can trigger the "herd effect"— one of several techniques proven to move a customer towards a commitment. Anchoring negotiations in a full repayment within a short time-frame will help a customer commit to making the biggest payment they can manage. This not only maximizes recovery for the bank but also protects the customer from unnecessary interest charges and bankruptcy that could result from falling victim to hyperbolic discounting.

Ensuring that customers keep their promise to pay is arguably the hardest part of collections. Again, behavioral segmentation sheds light on the intricate factors determining the decision to follow through—or not—on a promised payment. One justification customers frequently use to rationalize broken promises is the hassle (actual or perceived) involved in making a payment. A quarter of respondents in our survey of US consumers with credit delinquencies said that making payments was a hassle, and a third of this group said they would be more likely to pay if more convenient payment methods were available (Exhibit 4).

Exhibit 4

For some consumers, payment can be a hassle.

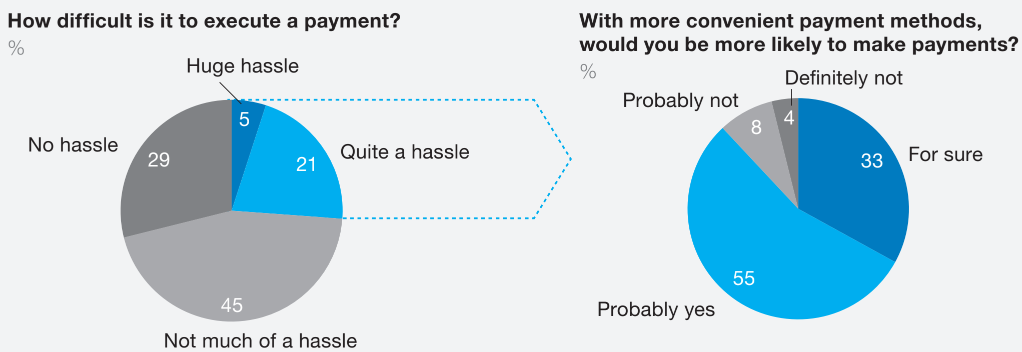

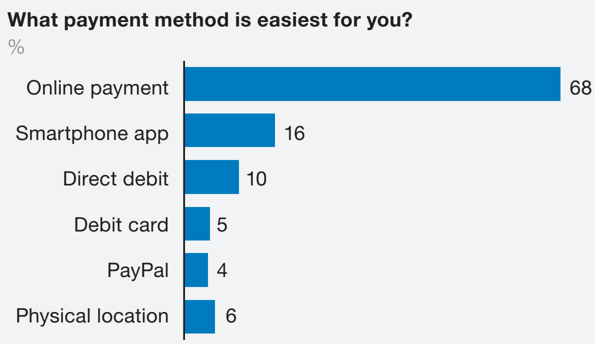

Comments

"An option to use a prepaid card or something like that would help. Nine times out of ten if the money gets put in the bank account, it will be taken out by another bill."

"I wish there was an easier way to send payments from my debit account. I hate finding out all the account numbers."

Source: McKinsey survey

One bank that piloted innovative treatments saw multiple benefits: a 30 percent increase in collections, a 20 percent reduction in write-offs on delinquent debt, a 33 percent fall in delinquencies remaining for late-stage collection, and a 20 percent reduction in the number of customers subsequently relapsing into default (Exhibit 5, page 79). First, the bank used K-means clustering to create an initial segmentation of five behavioral clusters. Next, it used a range of tools including closed-file reviews, psychometric surveys, and interviews to compile an ethnographic profile for each cluster. Finally, it drew on the growing body of psychological research and real-life experience with nudges and other psychological interventions in other industries to design effective treatments..

* * *

Being aware of bias and taking deliberate steps to counter it has already proved effective in areas such as gender bias in hiring. A few pioneering financial institutions have adopted a similar approach to debiasing their business decisions and have seen impressive results, such as a 25 to 35 percent reduction in credit losses from improved underwriting and collections. Yet for most institutions, the big prizes have yet to be captured.

Exhibit 5

Tailored treatments based on behavioral segmentation can deliver multiple benefits.

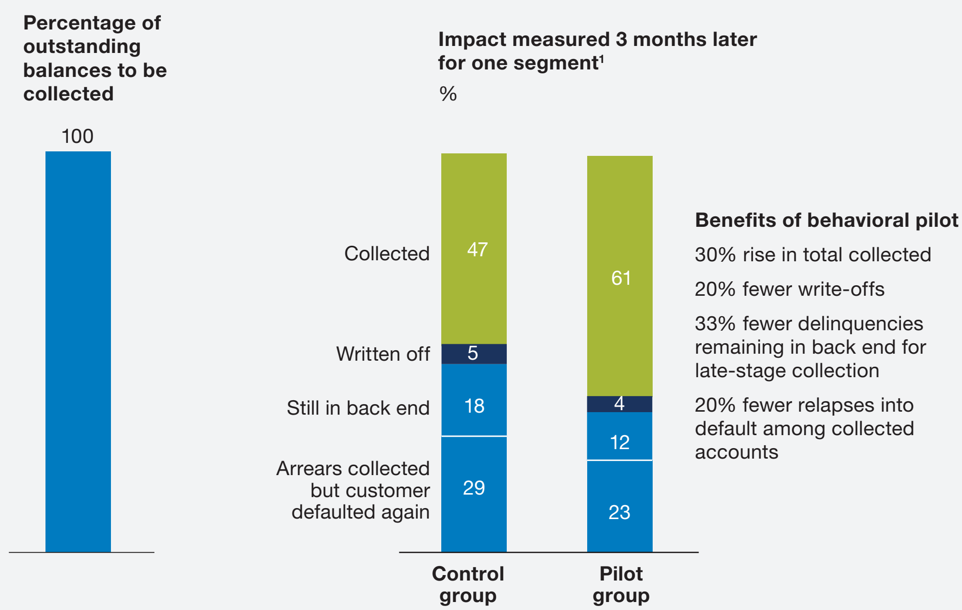

Behavioral-based prescriptive treatment for this segment

High call priority; no delay in efforts because of messages being left

Thorough inquiry with detailed questions about the customer's situation

Assertive script with no inappropriate "customer service" mindset

Questions about how, where, and when customer will pay help form an "implementation intention" that makes them more likely to keep their promise

[1] Figures may not sum to 100% because of rounding

Source: McKinsey analysis

The secret lies in combining psychological insights with advanced statistical methods to develop a pragmatic but powerful behavioral segmentation linked to targeted treatments. By introducing creative workarounds into their existing infrastructure, especially in IT implementation, providers can have a new approach up and running in as little as three months, with dramatic effects. Given the impact that early efforts have achieved, it can be only a matter of time before such innovative treatments become the norm.

Tobias Baer conducts psychological research at the University of Cambridge; he is a a former partner at McKinsey and a member of our Behavioral Insights Group. **Vijay D'Silva** is a senior partner in McKinsey's New York office.

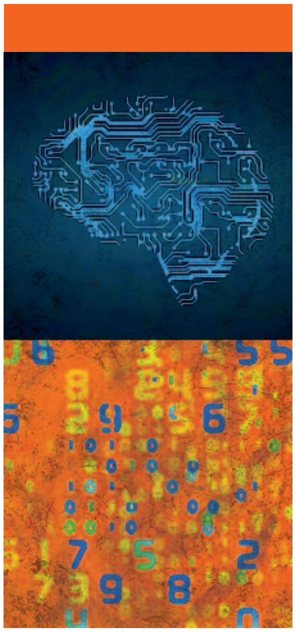

Mapping AI techniques to problem types

As artificial intelligence technologies advance, so does the definition of which techniques constitute AI (see sidebar, “Deep learning’s origins and pioneers”).[1] For the purposes of this paper, we use AI as shorthand specifically to refer to deep learning techniques that use artificial neural networks. In this section, we define a range of AI and advanced analytics techniques as well as key problem types to which these techniques can be applied.

Michael Chui
James Manyika
Mehdi Miremadi
Nicolaus Henke
Rita Chung
Pieter Nel
Sankalp Malhotra

Neural networks and other machine learning techniques

We looked at the value potential of a range of analytics techniques. The focus of our research was on methods using artificial neural networks for deep learning, which we collectively refer to as AI in this paper, understanding that in different times and contexts, other techniques can and have been included in AI. We also examined other machine learning techniques and traditional analytics techniques (Exhibit 1, page 81). We focused on specific potential applications of AI in business and the public sector (sometimes described as “artificial narrow AI”) rather than the longer-term possibility of an “artificial general intelligence” that could potentially perform any intellectual task a human being is capable of.

Neural networks are a subset of machine learning techniques. Essentially, they are AI systems based on simulating connected “neural units,” loosely modeling the way that neurons interact in the brain. Computational models inspired by neural connections have been studied since the 1940s and have returned to prominence as computer processing power has increased and large training data sets have been used to successfully analyze input data such as images, video, and speech. AI practitioners refer to these techniques as “deep learning,” since neural networks have many (“deep”) layers of simulated interconnected neurons. Before deep learning, neural networks often had only three to five layers and dozens of neurons; deep learning networks can have seven to ten or more layers, with simulated neurons numbering into the millions.

In this paper, we analyzed the applications and value of three neural network techniques:

- Feed forward neural networks. One of the most common types of artificial neural network. In this architecture, information moves in only one direction, forward, from the input layer, through the “hidden” layers, to the output layer. There are no loops in the network. The first single-neuron network was proposed in 1958 by AI pio-

Editor’s note: This article is a reprint of a chapter from the April 2018 McKinsey Global Institute discussion paper, “Notes from the AI frontier: Insights from hundreds of use cases.” The full paper can be downloaded here: https://www.mckinsey.com/featured-insights/artificial-intelligence/notes-from-the-ai-frontier-applications-and-value-of-deep-learning.

[1] For a detailed look at AI techniques, see An executive’s guide to AI, McKinsey Analytics, January 2018. https://www.mckinsey.com/business-functions/mckinsey-analytics/our-insights/an-executives-guide-to-ai.

Exhibit 1

Artificial intelligence, machine learning, and other analytics techniques examined for this research.

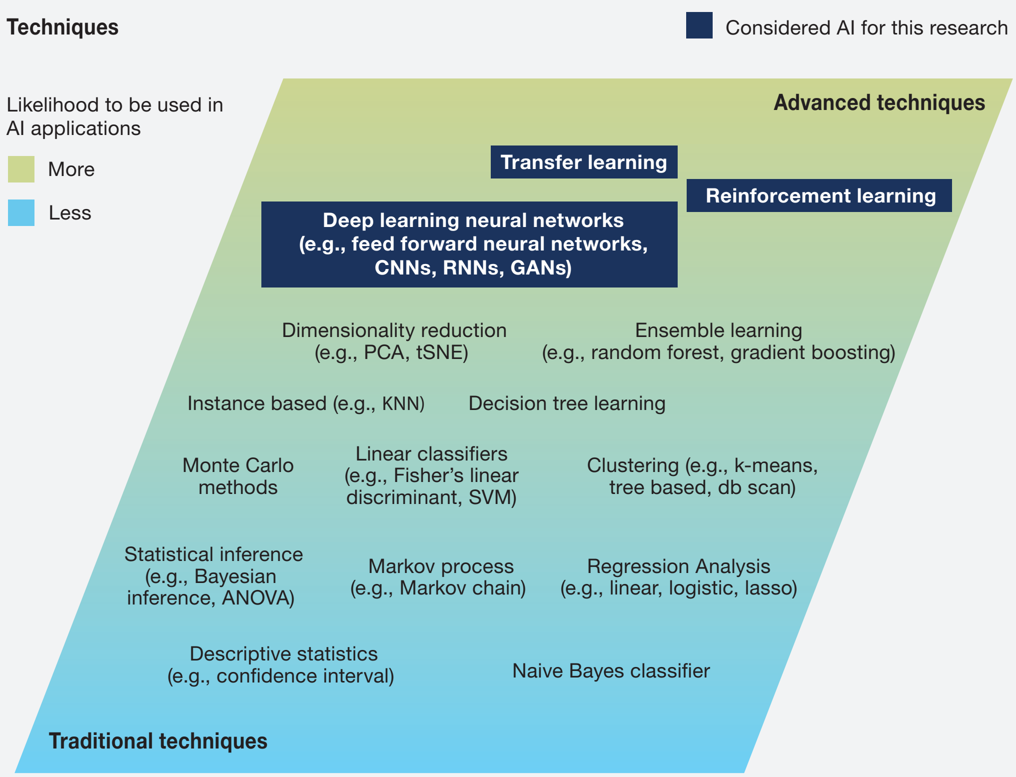

Source: McKinsey Global Institute analysis

neer Frank Rosenblatt. While the idea is not new, advances in computing power, training algorithms, and available data led to higher levels of performance than previously possible.

- Recurrent neural networks (RNNs). Artificial neural networks whose connections between neurons include loops, well-suited for processing sequences of inputs, which makes them highly effective in a wide range of applications, from handwriting, to texts, to speech recognition. In November 2016, Oxford University researchers reported that a system based on recurrent neural networks (and convolutional neural networks) had achieved 95 percent accuracy in reading lips, outperforming experienced human lip readers, who tested at 52 percent accuracy.

- Convolutional neural networks (CNNs). Artificial neural networks in which the connections between neural layers are inspired by the organization of the animal visual cortex, the portion of the brain that

Deep learning's origins and pioneers

It is too early to write a full history of deep learning—and some of the details are contested—but we can already trace an admittedly incomplete outline of its origins and identify some of the pioneers. They include Warren McCulloch and Walter Pitts, who as early as 1943 proposed an artificial neuron, a computational model of the "nerve net" in the brain.[2] Bernard Widrow and Ted Hoff at Stanford University developed a neural network application by reducing noise in phone lines in the late 1950s.[3] Around the same time, Frank Rosenblatt, an American psychologist, introduced the idea of a device called the Perceptron, which mimicked the neural structure of the brain and showed an ability to learn.[4] MIT's Marvin Minsky and Seymour Papert then put a damper on this research in their 1969 book "Perceptrons," by showing mathematically that the Perceptron could only perform very basic tasks.[5] Their book also discussed the difficulty of training multi-layer neural networks. In 1986, Geoffrey Hinton at the University of Toronto, along with colleagues David Rumelhart and Ronald Williams, solved this training problem with the publication of a now famous back propagation training algorithm—although some practitioners point to a Finnish mathematician, Seppo Linnainmaa, as having invented back propagation already in the 1960s.[6] Yann LeCun at NYU pioneered the use of neural networks on image recognition tasks and his 1998 paper defined the concept of convolutional neural networks, which mimic the human visual cortex.[7] In parallel, John Hopfield popularized the "Hopfield" network which was the first recurrent neural network.[8] This was subsequently expanded upon by Jurgen Schmidhuber and Sepp Hochreiter in 1997 with the introduction of the long short-term memory (LSTM), greatly improving the efficiency and practicality of recurrent neural networks.[9] Hinton and two of his students in 2012 highlighted the power of deep learning when they obtained significant results in the well-known ImageNet competition, based on a dataset collated by Fei-Fei Li and others.[10] At the same time, Jeffrey Dean and Andrew Ng were doing breakthrough work on large scale image recognition at Google Brain.[11] Deep learning also enhanced the existing field of reinforcement learning, led by researchers such as Richard Sutton, leading to the game-playing successes of systems developed by DeepMind.[12] In 2014, Ian Goodfellow published his paper on generative adversarial networks, which along with reinforcement learning has become the focus of much of the recent research in the field.[13] Continuing advances in AI capabilities have led to Stanford University's One Hundred Year Study on Artificial Intelligence, founded by Eric Horvitz, building on the long-standing research he and his colleagues have led at Microsoft Research. We have benefited from the input and guidance of many of these pioneers in our research over the past few years.

---

[2] Warren McCulloch and Walter Pitts, "A logical calculus of the ideas immanent in nervous activity," B*ulletin of Mathematical Biophysics*, volume 5, 1943.

[3] Andrew Goldstein, "Bernard Widrow oral history," *IEEE Global History Network*, 1997.

[4] Frank Rosenblatt, "The Perceptron: A probabilistic model for information storage and organization in the brain," *Psychological review*, volume 65, number 6, 1958.

[5] Marvin Minsky and Seymour A. Papert, *Perceptrons: An introduction to computational geometry*, MIT Press, January 1969.

[6] David E. Rumelhart, Geoffrey E. Hinton, and Ronald J. Williams, "Learning representations by back-propagating errors," Nature, volume 323, October 1986; for a discussion of Linnainmaa's role see Juergen Schmidhuber, Who invented backpropagation?, Blog post http://people.idsia.ch/~juergen/who-invented-backpropagation.html, 2014.

[7] Yann LeCun, Patrick Haffner, Leon Botton, and Yoshua Bengio, Object recognition with gradient-based learning, Proceedings of the IEEE, November 1998.

[8] John Hopfield, Neural networkds and physical systems with emergent collective computational abilities, PNAS, April 1982.

[9] Sepp Hochreiter and Juergen Schmidhuber, "Long short-term memory," Neural Computation, volume 9, number 8, December 1997.

[10] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton, ImageNet classification with deep convolutional neural networks, NIPS 12 proceedings of the 25th International Conference on Neural Information Processing Systems, December 2012.

[11] Jeffrey Dean et al., Large scale distributed deep networks, NIPS 2012.

[12] Richard S. Sutton and Andrew G. Barto, Reinforcement learning: An introduction, MIT Press, 1998.

[13] Ian J. Goodfellow, Generative adversarial networks, ArXiv, June 2014.

processes images, well suited for visual perception tasks.

We estimated the potential of those three deep neural network techniques to create value, as well as other machine learning techniques such as tree-based ensemble learning, classifiers, and clustering, and traditional analytics such as dimensionality reduction and regression.

For our use cases, we also considered two other techniques—generative adversarial networks (GANs) and reinforcement learning—but did not include them in our potential value assessment of AI, since they remain nascent techniques that are not yet widely applied in business contexts. However, as we note in the concluding section of this paper, they may have considerable relevance in the future.

- Generative adversarial networks (GANs). These usually use two neural networks contesting each other in a zero-sum game framework (thus "adversarial"). GANs can learn to mimic various distributions of data (for example text, speech, and images) and are therefore valuable in generating test datasets when these are not readily available.
- Reinforcement learning. This is a subfield of machine learning in which systems are trained by receiving virtual "rewards" or "punishments," essentially learning by trial and error. Google DeepMind has used reinforcement learning to develop systems that can play games, including video games and board games such as Go, better than human champions.

Problem types and the analytic techniques that can be applied to solve them

In a business setting, those analytic techniques can be applied to solve real-life problems. For this research, we created a taxonomy of high-level problem types, characterized by the inputs, outputs, and purpose of each. A corresponding set of analytic techniques can be applied to solve these problems. These problem types include:

- **Classification.** Based on a set of training data, categorize new inputs as belonging to one of a set of categories. An example of classification is identifying whether an image contains a specific type of object, such as a truck or a car, or a product of acceptable quality coming from a manufacturing line.
- **Continuous estimation**. Based on a set of training data, estimate the next numeric value in a sequence. This type of problem is sometimes described as "prediction," particularly when it is applied to time series data. One example of continuous estimation is forecasting the sales demand for a product, based on a set of input data such as previous sales figures, consumer sentiment, and weather. Another example is predicting the price of real estate, such as a building, using data describing the property combined with photos of it.
- **Clustering.** These problems require a system to create a set of categories, for which individual data instances have a set of common or similar characteristics. An example of clustering is creating a set of consumer segments based on data about individual consumers, including demographics, preferences, and buyer behavior.
- **All other optimization.** These problems require a system to generate a set of outputs that optimize outcomes for a specific

Exhibit 2

Problem types and sample techniques.

Essential Relevant

Total AI value potential that could be unlocked by problem types as essential versus relevant to use cases
%

Problem types	**Sample techniques**
Classification	CNNs, logistic regression
Continuous estimation	Feed forward neural networks, linear regression
Clustering	K-means, affinity propagation
All other optimization	Genetic algorithms
Anomaly detection	One-class support vector machines, k-nearest neighbors, neural networks
Ranking	Ranking support vector machines, neural networks
Recommender systems	Collaborative filtering
Data generation	Generative adversarial networks (GANs), hidden Markov models

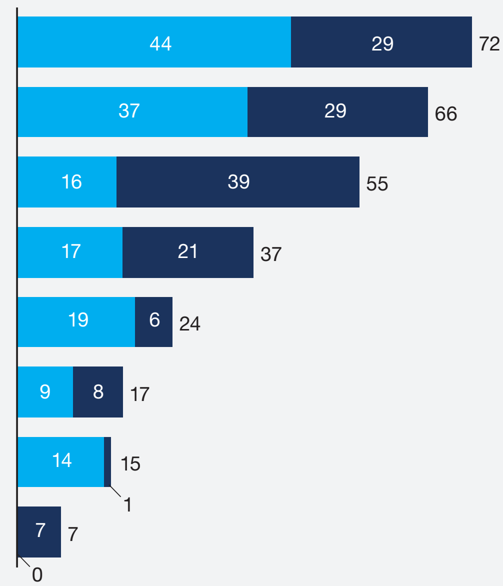

NOTE: Sample techniques include traditional analytical techniques, machine learning, and the deep learning techniques we describe in this paper as AI. Numbers may not sum due to rounding.
Source: McKinsey Global Institute analysis

objective function (some of the other problem types can be considered types of optimization, so we describe these as "all other" optimization). Generating a route for a vehicle that creates the optimum combination of time and fuel use is an example of optimization.

- **Anomaly detection.** Given a training set of data, determine whether specific inputs are out of the ordinary. For instance, a system could be trained on a set of historical vibration data associated with the performance of an operating piece of machinery, and then determine whether a new vibration reading suggests that the machine is not operating normally. Note that anomaly detection can be considered a subcategory of classification.
- **Ranking.** Ranking algorithms are used most often in information retrieval problems in which the results of a query or request needs to be ordered by some

criterion. Recommendation systems suggesting next product to buy use these types of algorithms as a final step, sorting suggestions by relevance, before presenting the results to the user.

- **Recommendations.** These systems provide recommendations, based on a set of training data. A common example of recommendations are systems that suggest the "next product to buy" for a customer, based on the buying patterns of similar individuals, and the observed behavior of the specific person.

- **Data generation.** These problems require a system to generate appropriately novel data based on training data. For instance, a music composition system might be used to generate new pieces of music in a particular style, after having been trained on pieces of music in that style.

Exhibit 2 (page 84) illustrates the relative total value of these problem types across our database of use cases, along with some of the sample analytics techniques that can be used to solve each problem type. The most prevalent problem types are classification, continuous estimation, and clustering, suggesting that meeting the requirements and developing the capabilities in associated techniques could have the widest benefit. Some of the problem types that rank lower can be viewed as subcategories of other problem types—for example, anomaly detection is a special case of classification, while recommendations can be considered a type of optimization problem—and thus their associated capabilities could be even more relevant.

Michael Chui is a partner in McKinsey's San Francisco office, where **James Manyika** is a senior partner. **Mehdi Miremadi** is a partner in the Chicago office, **Nicolaus Henke** is a senior partner in the London office, and **Rita Chung** is a consultant in the Silicon Valley office. **Pieter Nel** is a digital expert in the New York office, where **Sankalp Malhotra** is a consultant.

Data sheet: Advanced analytics

The size of the prize in banking as data and analytical tools and use cases proliferate is significant, and leaders are pulling ahead.

Advanced analytics could be worth $0.5-1 trillion to banks globally, representing 8-14% of revenue.

$ annually

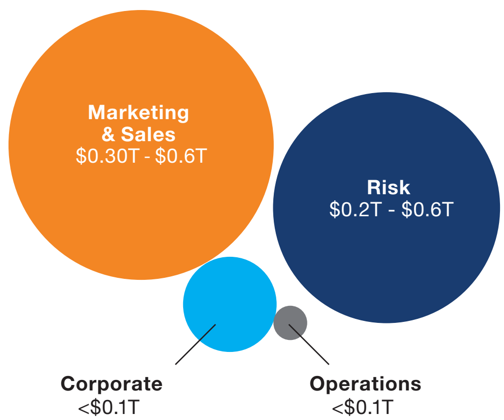

Source: McKinsey Global Institute

Big data and artificial intelligence/machine learning ventures receive most VC investment outside of SaaS.

$ million, 2017

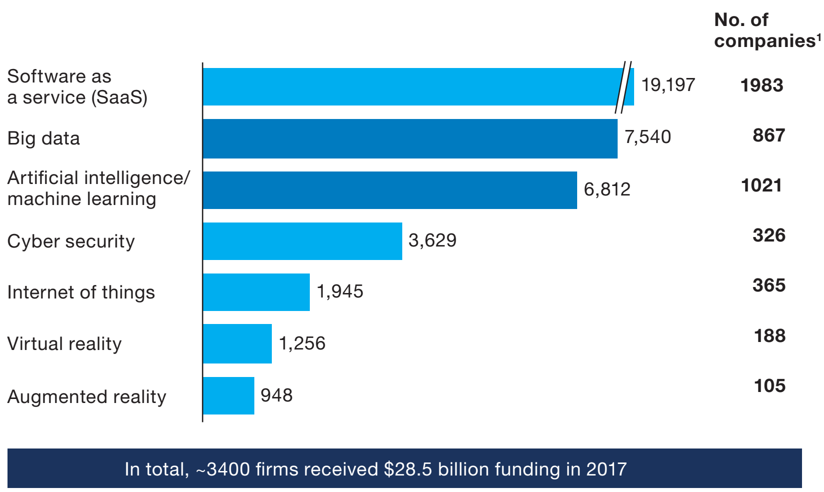

[1] Number of companies which have received funding in 2017. Some of the startups are classified in more than one vertical.

Source: PitchBook; SILA (Startup and Investment Landscape Analytics) leverages McKinsey databases covering more than 1.7 million companies, to significantly accelerate M&A target scans in a wide variety of industries and sub sectors globally.

Banks are increasingly deploying more sophisticated analytics use cases in early stages, but not yet at scale.

Percent of banking respondents deploying use case at scale

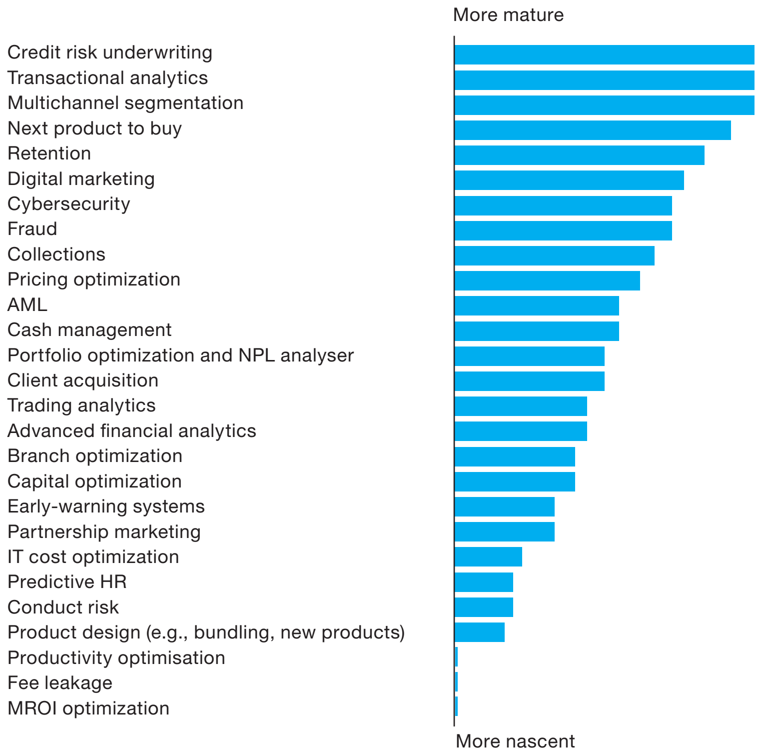

Source: McKinsey Analytics Quotient

Analytics leaders[1] exhibit stronger financial performance than other companies

Analytics maturity correlation to financial metrics

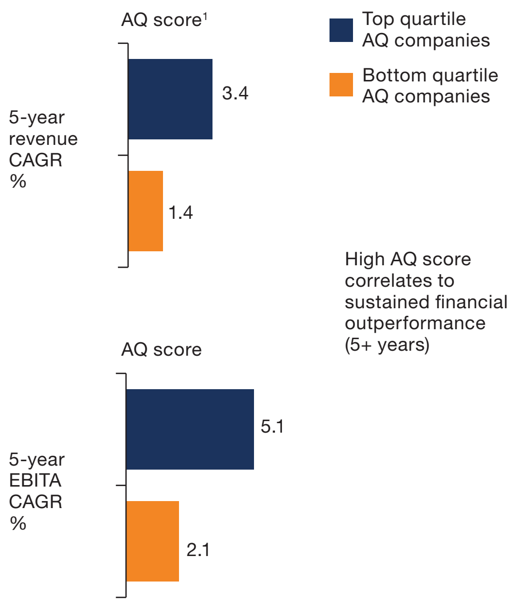

[1] Analytics Quotient (AQ) is a McKinsey solution developed in 2017 to standardize the measure of analytics maturity across sectors. The AQ survey has been taken by over 120 companies, across industries worldwide. Analytics leaders are defined as the top quartile of companies as determined by overall AQ score.

Source: McKinsey Analytics Quotient

Geospatial analytics, now indispensable in the retail industry, will become one of the critical drivers in consumer banking.

Tokyo residents with similar purchasing behaviors shop in a few specific areas of the city

Demographic and income data were inferred by analyzing where mobile devices are stationary at night

McKinsey Geospatial team analysis using Factual's point of interest and mobile device location data

Some of the smallest transportation corridors by traffic volume have most potential for high-end transportation services, given commuters socio-economic profile

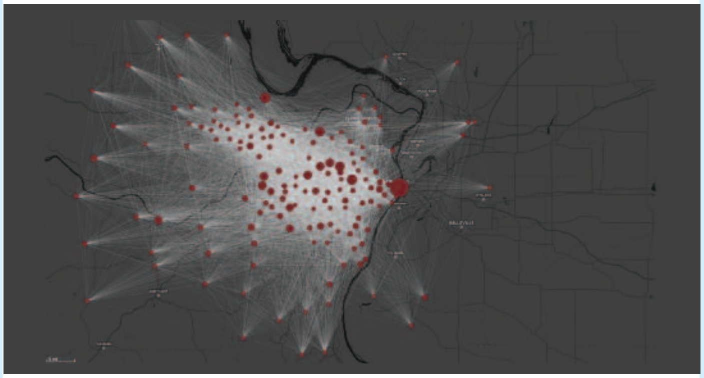

McKinsey team analysis based on commuting flow and demographic data from the US Census

Foot traffic, dispersed across the city during work hours, becomes much more concentrated in one "special ward" (municipality) of Tokyo in the evening

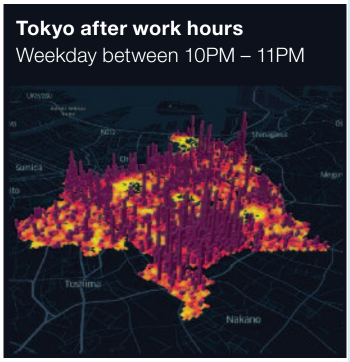

McKinsey Geospatial team analysis using Skyhook`s mobile location data and visualized with Kepler.gl

Luxury retailer identified store specific halo effects on web sales based on store and online transaction data, by considering the effects of proximity to target demographic, competitor stores and foot traffic

McKinsey's OMNI Solution blends online activity and in-store sales along with demographics and mobile device location data.

Banks improving intelligence of their omnichannel coverage models

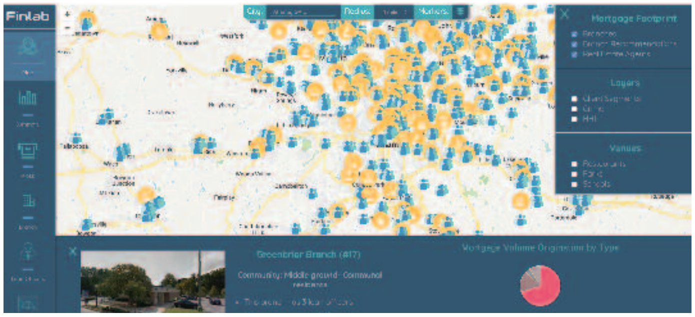

Many banks have a sub-optimal branch footprint, as a result of legacy business decisions

There are substantial behavioral differences between customer segments, even within the same neighborhood, which drive their choices of branch visits and online channels

Source: McKinsey Retail Branch Geospatial Optimizer Tool provides an advanced analytics solution to retail banks focused on transforming the branch network into a high-value, high-functioning operation, by utilizing digital and offline data.

Global Banking Practice
August 2018

www.mckinsey.com/clientservice/financial_services