

Straight talk about big data

Transforming analytics from a “science-fair project” to the core of a business model starts with leadership from the top. Here are five questions CEOs should be asking their executive teams.

by Nicolaus Henke, Ari Libarikian, and Bill Wiseman

The revolution isn’t coming—it’s already under way. In the science of management, the revolution in big data analytics is starting to transform how companies organize, operate, manage talent, and create value. Changes of this magnitude require leadership from the top, and CEOs who embrace this opportunity will increase their companies’ odds of long-term success. Those who ignore or underestimate the eventual impact of this radical shift—and fail to prepare their organizations for the transition—do so at their peril.

It’s easy to see how analytics could get delegated or deprioritized: CEOs are on the hook for performance, and for all of the potential associated with analytics, many leaders operating in the here and now are reporting underwhelming results. In fact, when we surveyed a group of leaders from companies that are committed to big data–analytics initiatives, three-quarters of them reported that their revenue or cost improvements were less than 1 percent. Some of the disconnect between promise and payoff may be attributed to undercounting—the sum of the parts is not always immediately apparent. Ironically, the results of “big data” analytics are often thousands— or more— of incrementally *small* improvements realized system-wide. Individually, any one of these gains may appear insignificant, but when considered in the aggregate they can pack a major punch.

The shortfalls, however, are more than just a matter of perception, and the pitfalls are real. Critically, an analytics-enabled transformation is as much about a cultural change as it is about parsing the data and putting in place advanced tools. “This is something I got wrong,” admits Jeff Immelt, the CEO of GE. “I thought it was all about technology. I thought if we hired a couple thousand technology people, if we upgraded our software, things like that, that was it. I was wrong. Product managers have to be different; salespeople have to be different; on-site support has to be different.”

CEOs who are committed to a shift of this order, yet wonder how far the organization has truly advanced in its data-analytics journey to date, should start by stimulating a frank discussion with their top team. That includes a clear-eyed assessment of the fundamentals, including your company’s key value drivers, your organization’s existing analytics capabilities, and, perhaps most important, your purpose for committing to analytics in the first place. (See “Making data analytics work for you—instead of the other way around,” on McKinsey.com.) This article poses questions—but not shortcuts—to help a company’s senior leaders determine where they are and what needs to change for their organization to deliver on the promise of advanced analytics.

TWO SCENES FROM THE FRONT LINES OF THE REVOLUTION

Immelt reached his conclusions from witnessing—and, in many respects, leading—the revolution. GE’s CEO is keenly aware that so far in the 21st century, the digitization of commerce and media has allowed a handful of US Internet stalwarts to capture almost all the market value created in the consumer sector. To avoid a similar disruption as the industrial world goes online over the coming decade, Immelt is driving a radical shift in the culture and business model of his 124-year-old company. GE is spending \$1 billion this year alone to analyze data from sensors on gas turbines, jet engines, oil pipelines, and other machines and aims to triple sales of software products by 2020 to roughly \$15 billion. To make sense of those new streams of data, the company is also building a cloud-based platform called Predix, which combines its own information flows with customer data and submits them to analytics software that can lower costs and increase uptime through vastly improved predictive maintenance. Getting this right will require hiring several thousand new software engineers and data scientists, retraining tens of thousands of salespeople and support staff, and fundamentally shifting GE’s business model from product sales coupled with service licenses to outcomes-based subscription pricing. “We want to treat analytics like it’s as core to the company over the next 20 years as material science has been over the past 50 years,” says Immelt.

To understand further the growing power of advanced analytics, consider as well how a consumer-electronics OEM is picking up more speed in an inherently slow-growth market. The company started with a Herculean effort to pull together information on more than 1,000 variables previously collected in silos across millions of devices and sources—product sales and usage data, channel data, online transactions, and service tickets, plus external consumer data from third-party suppliers such as Acxiom. Mining this integrated big data set allowed the company to home in on a dozen or so unrealized opportunities where a shift in investment patterns or processes would really pay off. Armed with a host of new, fine-grained insights on which moves offered the greatest odds to increase sales, reduce churn, and improve product features, the company went on to realize \$400 million in incremental revenue increases in year one. As success builds, the leadership has begun to fundamentally rethink how it goes about new-business development and what future capabilities its top managers will require.

BIG CHALLENGES, BIGGER OPPORTUNITIES

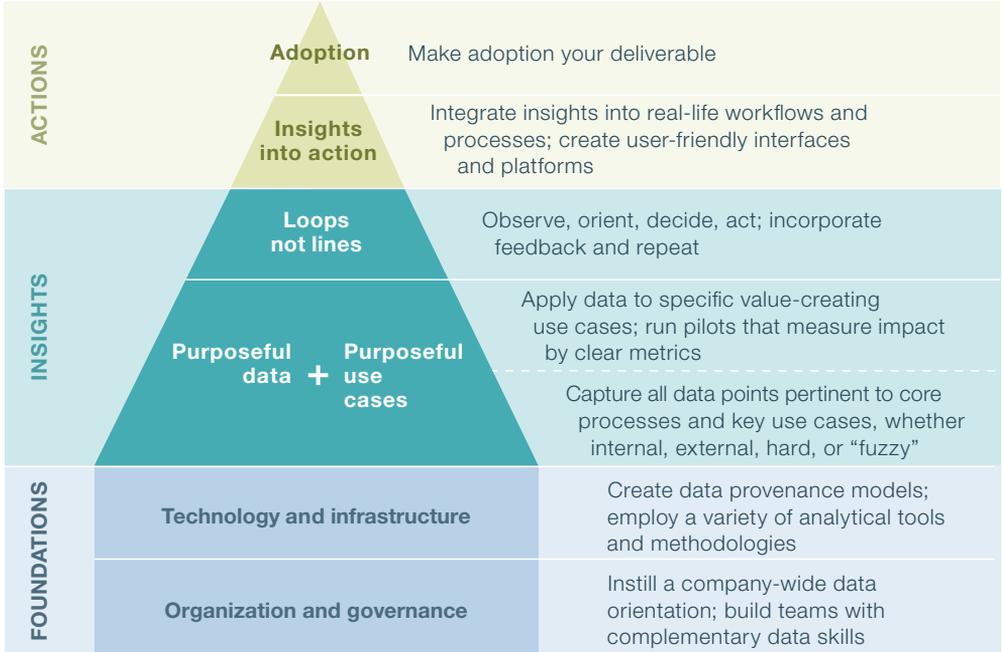
But for all the enormous promise, most companies—outside of a few digital natives such as Amazon, Facebook, Google, Netflix, and Uber—have so far struggled to realize anything more than modest gains from their investments in big data, advanced analytics, and machine learning. Many organizations remain preoccupied with classic large-scale IT-infrastructure programs and have not yet mastered the foundational task of creating clean, powerful, linked data assets; building the capabilities they need to extract insight from them; and creating the management capacity required to lead and direct all this toward purposeful action (exhibit).

Still, similar birthing pains have marked every previous major technology transition as well. And we're still in the early days of this one: though about 90 percent of the digital data ever created in the world has been generated in just the past two years, only 1 percent of that data has been analyzed. Often, those analyses are conducted as discrete one-offs—nifty experiments, but not much more. Indeed, in many companies, analytics initiatives still seem more like sideline science-fair projects than the core of a cutting-edge business model.

But the potential for significant breakthroughs demands an overhaul of that model, and the speed at which these breakthroughs advance will only accelerate. As computer-processing power and cloud-storage capacity swell, the world's current data flood becomes a tidal wave. By 2020, some 50 billion smart devices will be connected, along with additional billions of smart sensors, ensuring that the global supply of data will continue to more than double every two years.

Exhibit

Data analytics should have a purpose, be grounded in the right foundation, and always be conducted with adoption in mind.



LEADING QUESTIONS

All of these developments ensure that there will be a lot of data to analyze. Almost by definition, big data analytics means going deep into the information weeds and crunching the numbers with a technical sophistication that can appear so esoteric that senior leadership may be tempted simply to “leave it to the experts” and disengage from the conversation. But the conversation is worth having. The real power of analytics-enabled insights comes when they become so fully embedded in the culture that their predictions and prescriptions drive a company’s strategy and operations and reshape how the organization delivers on them. Extending analytics from the realm of tactical insights into the heart of the business requires hard work, but the benefits can be profound. Consider, for example:

- A global airline stitched together data from multiple operational systems (including those related to aircraft location and aerobridge position) to identify more precisely when and why flights were delayed as they pushed back or arrived at a gate. Its advanced prediction algorithms were able to quantify the knock-on impact of events such as mishandled luggage and helped build a system to alert supervisors in real time so that they

could react before potential problems developed. Impact: a reduction in delayed arrivals of about 25 percent over the past 12 months.

- A global consumer-packaged-goods company seeking to drive growth across categories integrated a wide range of information (including financial, promotional, and even weather-related factors) into a single data source and then developed sophisticated algorithms to understand the incremental effects that changes based on this source could have at even granular levels. Sifting through disparate data and building up from the ground level enabled the company to identify valuable insights about its competitive landscape as a whole, such as optimal price points and opportunities for new products. Impact: a gross-profit increase in the tens of millions of dollars within one year.
- A pharmaceuticals company is using analytics to stem the rising cost of clinical trials. After spending billions of dollars conducting hundreds of trials over the past five years, the company began integrating information on more than 100,000 patient participants with operational data from finance and HR. Out of those tens of millions of data points, it has started to pinpoint which locations are most efficient, which patient-screening techniques increase “pass rates,” and how best to configure its own teams. Analysis of email and calendar data, for example, underscored that improving collaboration between a team leader and two specific roles within clinical operations was among the most significant predictors of delays. The anticipated result: cost savings of more than 10 percent and better-quality outcomes.

And the list goes on: case after case of reduced churn, less fraud, improved collections, better return on investment from marketing and customer acquisition, and enhanced predictive maintenance. Right now, only a few leaders outside the tech sector are truly transforming their organizations with data. But more could be. To that end, we suggest five questions that company leaders should be prepared to explore in depth.

1. Do we have a value-driven analytics strategy?

Businesses can waste a lot of energy collecting data and mining them for insights if their efforts aren't focused on the areas that matter most for the company's chosen direction. Successful big data and advanced-analytics transformations begin with assessing your own value drivers and capabilities versus those of the competition and developing a picture of the ideal future state, one aligned with the broad business strategy and key use cases. Asking the right questions is the critical first step. These should start big: “What is

the size of this opportunity? If I had the additional insights possible through advanced analytics, how much could I save? How much additional revenue could I achieve?” And they should quickly get granular. To frame and develop the right hypotheses, frontline managers must engage alongside the analytics experts throughout the process.

That consumer-electronics OEM’s planning exercise, for example, led the team to ask several questions: “Who are our highest-value customers, how do we reach them, and what do we talk about? How can we drive more cross-sell of our broader portfolio of products and services? Which product features drive the highest usage or engagement, and how do we promote higher adoption of them?” At a leading private bank, the questions from a similar exercise included these: “How can we set optimal price points, day in and day out, and by the thousands each day? Which customers are most at risk of leaving, which are most likely to respond favorably to retention efforts, and what types of retention efforts work best?”

2. Do we have the right ‘domain data’ to support our strategy?

In answering such questions, companies typically identify 10 to 20 key use cases in areas such as revenue growth, customer experience, risk management, and operations where advanced analytics could produce clear-cut improvements. On the basis of that self-assessment and the anticipated impact on earnings, the use cases are ranked and pilot projects are sequenced. Measuring the impact of each use case, with specific indicators and benchmarks, highlights what data are needed and keeps things on track.

A critical foundational step is to overcome obstacles to using existing data. This work could include cleaning up historical data, integrating data from multiple sources, breaking down silos between business units and functions, setting data-governance standards, and deciding where the most important opportunities may lie to generate new internal data—for example, by adding sensors, or, in the case of, say, casinos, by installing webcams to assess high-roller betting behavior.

Most companies, even those with rich internal data, will also conclude they need to mine the far-larger universe of structured and unstructured external sources. When one emerging-markets insurer decided to launch a new peer-to-peer-lending start-up, it realized it could make even better credit decisions by analyzing potential customers’ data and movements on its various platforms, including social networks.¹

¹ For more on opportunities to use public information and shared data from private sources, see *Open data: Unlocking innovation and performance with liquid information*, McKinsey Global Institute, October 2013, on McKinsey.com.

All these data eventually can be pooled into more shareable, accessible assets, such as new “data lakes.” Getting the foundation sorted is not the work of weeks or even months, as anyone who has wrestled with the shortcomings of legacy IT systems knows. And the cost can eventually run into hundreds of millions of dollars, while the full impact of those investments will not always be obvious in one quarter, or two, or three. But that doesn’t mean you should wait for years to capture value. Which makes it all the more important to ask the next question.

3. Where are we in our journey?

Like any transition, the data-and-analytics journey takes place in stages. It’s crucial both to start laying the foundation and to start building analytics capabilities even before the foundation is set. Or, as one of our clients recently recalled as he thought about his company’s successful analytics transformation: “We needed to walk before we could run. And then we ran like hell.”

To step smartly in fast-forward mode, the consumer-electronics OEM created an interim data architecture focused on building and staffing three “insights factories” that could generate actionable recommendations for its highest-priority use cases. While further foundational investments continued in parallel, those factories enabled the early pilots to deliver quick results that made them largely self-funding. The key is to move quickly from data collection to “doing the math,” with an iterative, hypothesis-driven modeling cycle. Such rapid successes help break down silos and build enthusiasm and buy-in among often skeptical frontline managers. Even if it works, a “black box” developed by data scientists working in isolation will usually prove a recipe for rejection. End users need to understand the basic assumptions and how to apply the model’s output: Are its recommendations binding, or is there flexibility to deviate? Will it be integrated directly into core tools such as customer relationship management, or will it be an additional overlay? What visual display will be most useful for the front line—in general, simpler is better—once the data are produced? Pilots should be designed to answer these questions even before the data are collected and the model is built.

Once proof of concept is established and points start going on the board, it’s critical to go big as quickly as possible, which can require an infusion of talent. Best-practice companies rarely cherry-pick one or two specialist profiles to recruit to address isolated challenges. In our recent survey of more than 700 companies, we found that 15 percent of operating-profit increases were linked to the hiring of data-and-analytics experts at scale.

4. Are we modeling the change personally?

In a recent survey of more than 500 executives, we turned up a distressing finding: while 38 percent of CEOs self-reported that they were leading their companies' analytics agendas, only 9 percent of the other C-suite executives agreed. They instead identified the chief information officer or some other executive as the true point person. What we've got here, to paraphrase the warden in *Cool Hand Luke*, is more than a failure to communicate; it's about not walking the talk.

While CEOs and other members of the executive team don't need to be the experts on data science, they must at least become conversant with a jungle of new jargon and buzzwords (Hadoop, genetic algorithms, in-memory analytics, deep learning, and the like) and understand at a high level the limits of the various kinds of algorithmic models. In addition to constant communication from the top that analytics is a priority and public celebration of successes, small signals such as recommending and showing up for learning opportunities also resonate.

The most important role modeling a CEO can deliver, of course, is to ensure that the right kind of conversations are taking place among the company's top management. That starts with ensuring that the right people are both in the room and empowered, and then continues with direct intervention and questioning to ensure the transition from experience-based decision making to data-based decision making: Was a conclusion A/B tested? What have we done to build up our capability to conduct rapid prototyping, to test and learn and experiment, to constantly engage in what Google chief economist Hal Varian calls "product kaizen"?²

5. Are we organizing and leading for analytics?

The most important shift, which only the CEO can lead, is to reorganize to put advanced analytics at the center of every core process. The aspiration, in fact, should be to eventually eliminate the distinct term "analytics" from the company lexicon. Data flow through the whole organization, and the analytics should organically follow. "I just think it's infecting everything we do, in a positive way," says GE's Immelt.

² See Hal R. Varian, "Kaizen, that continuous improvement strategy, finds its ideal environment," *New York Times*, February 8, 2007, nytimes.com; and Hal R. Varian, *Computer mediated transactions*, UC Berkeley and Google presentation, Berkeley, CA, January 3, 2010, people.ischool.berkeley.edu.

Still, even a central nervous system requires a brain—a central analytics hub, or center of excellence. Without a dedicated team and leader, whether a chief analytics officer or a chief data officer or a senior C-level executive clearly tasked with the role, companies struggle to create a distinctive culture that can attract and nurture the best talent. But at the same time, as with a function like finance, individuals connected with the central team should also be embedded in the separate business units. We’ve found that executives from companies with a hybrid model reported a greater impact from analytics on revenue and costs than other respondents did.

What additional roles, skills, and structures are necessary? Clearly, scaling up analytics requires recruiting and retaining a sufficient number of world-class data scientists and model builders. Buying such talent on an outsourced model is only an option for those still in the exploratory phase of their journey. But, to take one example, most banks in the post-stress-test world have created separate, in-house units with their own reporting lines, charged with constantly testing and validating those models to minimize the risk of spurious correlations. We believe this approach makes sense for nonbanks as well.

To turn modeling outputs, however robust, into tangible business actions, companies also need a sufficient supply of “translators,” people able to connect the needs of the business units with the technical skills of the modelers. Don’t assume such “two sport” leaders are easy to find. In our experience, executives often report that attracting and retaining business users with analytics skill sets is actually slightly harder than recruiting those in-demand data scientists themselves. Alongside aggressive recruiting, winning this war for talent requires doubling down on training and improved HR analytics.

In general, most organizations are underinvesting in creating intuitive tools with easy-to-use interfaces that can help frontline managers integrate data into day-to-day processes. Our rule of thumb: for the highest payoff, split your analytics investments roughly 50–50 between spending on building better models and spending on tools and training to ensure that the front line uses the new insights being generated. In many companies, that ratio is still closer to 80–20, or worse.

Beyond big data and analytics, an even broader shift is under way, as robots, machine-learning algorithms, and “soft AI” systems, such as IBM’s Watson, take on more and more of the tasks that human labor used to conduct. Early in 2016, AlphaGo, a system developed by DeepMind, a British company owned

by Google, unexpectedly rolled over a celebrated human champion in the ancient game of Go.³ To prepare for a contest in which, unlike chess, there are more possible positions than grains of sand in the universe, AlphaGo trained itself by playing endless rounds of games, which enabled the path-optimization strategies.⁴

As the use of data and analytics incorporates machine learning, and artificial intelligence continues to blur, humans can take comfort from one near certainty: as proved true in chess after 1997, when IBM's Deep Blue defeated Garry Kasparov, the new "best players" of Go will turn out to be neither humans nor machines alone, but rather humans working in tandem with machines. Mastering how to leverage that combination may be the ultimate CEO management challenge.

Science-fiction writer Arthur C. Clarke once said that "any sufficiently advanced technology is indistinguishable from magic." We haven't advanced to that level—yet. But as the age of big data gives way to the age of advanced analytics and machine learning, we are entering an era where the ability to analyze data will deliver a predictive capability that feels almost like magic.

As in other historic shifts, such as when modern firearms "disrupted" the crossbow, the competition between those who master the new technology and those who don't will be fierce. But the upside of adaptation is as inspiring as the downside is stark. In the years ahead, the companies and institutions that address these challenges frankly, transform their organizations accordingly, and apply these near-magical abilities seamlessly to the world's most complex and critical issues will deliver a level of value creation that today we can barely imagine. 

³ For more details about the match, see Choe Sang-Hun, "Google's computer program beats Lee Se-dol in Go tournament," *New York Times*, March 15, 2016, nytimes.com. For more on Google's acquisition of AlphaGo, see Rolfe Winkler, "Google acquires artificial-intelligence company DeepMind," *Wall Street Journal*, January 26, 2014, wsj.com.

⁴ For more on AlphaGo's learning process, see *Google Research Blog*, "AlphaGo: Mastering the game of Go with machine learning," blog entry by Demis Hassabis and David Silver, January 27, 2016, research.googleblog.com.

Nicolaus Henke is a senior partner in McKinsey's London office, **Ari Libarikian** is a senior partner in the New York office, and **Bill Wiseman** is a senior partner in the Taipei office.