


MARCH 2016



SEMICONDUCTORS

Improving the semiconductor industry through advanced analytics

New techniques can help companies make better decisions by using accurate, reliable, and scientific information to analyze risk, optimize processes, and predict failure.

Gaurav Batra, Zach Jacobson, and Nick Santhanam

When used effectively, advanced analytics can not only significantly improve operations and margins but also spur growth. Yet many companies, including several semiconductor players, have been slow to embrace these techniques. According to the International Data Corporation, the global pool of data is more than 2.8 zettabytes and growing, but companies generally use only about 0.5 percent of that ocean of information to make decisions. Businesses—usually consumer-facing ones—that do collect and analyze a broad range of data achieve many benefits. Banks, insurers, and retailers, for example, have used insights from advanced analytics to build sustained competitive advantages, including stronger customer relationships and greater operational efficiency.

Semiconductor companies have been leaders in generating and analyzing data. But few have effectively applied advanced analytics to fab operations, where they could improve predictive maintenance and yield, or to R&D and sales, for enhanced pricing, market-entry strategies, sales-force effectiveness, cross-selling, portfolio optimization, and other tasks.

But we may soon see the more widespread adoption of advanced analytics in semiconductors. First, computing power and storage infrastructure have become markedly easier to deploy with the advent of cloud computing. Second, there has been a step change in the power of the tools used to extract, aggregate, manage, analyze, and display data, as well as in techniques to incorporate data into actionable

models. Hadoop applications, for example, have made it possible to extract insights from unstructured data by simplifying the integration of disparate data sources.¹ Finally, slower industry growth is prompting semiconductor companies to look for operational efficiencies, including some that may be more easily identified and deployed with advanced analytics.

In this article, we provide a high-level overview of advanced-analytics strategies to illustrate how they can help companies transform their manufacturing, R&D, and sales functions. We include several examples of companies that have used these methods to optimize productivity, time to market, and other important metrics. Our objective is not to assess, compare, or evaluate different analytical tools or techniques or specific capabilities in detail. Rather, we wish to show the overall value of advanced analytics and the importance of improving these capabilities.

A new mind-set

Advanced analytics can provide a framework to unlock insights about where to invest and how to improve the performance of manufacturing, R&D, and sales. The engineers who handle these tasks should increasingly consider themselves data scientists—people mining information that can be used to improve business functions.

The methodologies of advanced analytics are quite different from those used in traditional empirical analysis (Exhibit 1). To put it simply, traditional data analysis is based on what information you have, advanced analytics on what information you need. The traditional approach usually starts with available data and focuses on the information they reveal and the insights they provide. Project and executive teams then determine how these insights might help them make specific decisions.

With advanced analytics, by contrast, teams begin by asking what business problems they are trying to solve and which critical decisions a company must

make. They then identify the insights that will help clarify those decisions, the type of information that might produce the required insights, and, finally, the data sources the organization needs to obtain this information. In a properly designed program, advanced analytics offers not only accurate, reliable, and timely information on past and present operations but also invaluable predictive insights to guide decision making (see sidebar, “Getting more from R&D and sales,” for an example of how advanced analytics can assist with these functions). Companies can use advanced analytics, for example, to create models anticipating future developments, such as R&D bottlenecks that could delay production. With this information, they can make better decisions to direct the business.

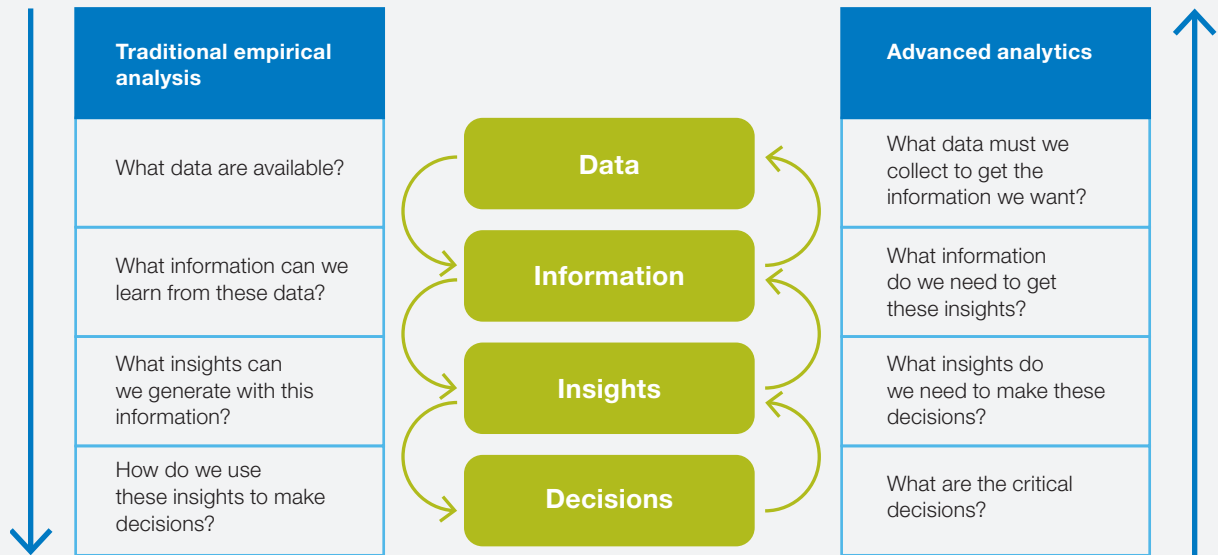
Applying advanced analytics in manufacturing

In chip manufacturing, the volume of data generated on the fab floor has continued to expand exponentially with each new node dimension. Leading-edge tools have so many measuring instruments that each one routinely identifies and gathers over 300 sensor inputs. In consequence, all information collected throughout the fab—including metrics for processes, products, and machine state—will quickly exceed terabytes of data. Fabs also gather extensive in-line, end-of-line inspection, and metrology data. Few, however, combine and apply advanced analytics to all these production data, even though that could improve many important manufacturing dimensions, including yield, throughput, equipment availability, and operating costs.

Consider, for example, a fab that wants to decrease equipment downtime. The fab could conduct a multi-variate analysis to enhance condition-based monitoring—a maintenance strategy that involves examining certain indicators to determine if equipment performance is decreasing. Among other benefits, the analysis would help the fab to predict more accurately when parts or consumables will fail.

Exhibit 1

Advanced analytics requires an approach completely different from that of traditional empirical analysis.



Source: McKinsey analysis

With this information, the fab can optimize the planned maintenance schedule, which will reduce downtime, as well as costs for parts and labor.

In addition to preventing equipment failures, fabs can use advanced analytics for more complex purposes. For instance, they could link equipment and process-level data to inspection and metrology data to make more accurate predictions about yield failures or yield degradation. Predictive modeling is difficult, since it requires multiple steps. Fabs must first gather complete data sets and then apply algorithmic approaches to identify patterns in the data before building any models. However, the payoffs can be great. Take the case of a company that recently used advanced analytics to predict process failure in a production step that involved depositing material on a wafer. The company was able to make the prediction with a confidence interval of about 70 percent—a level that might seem low but is comparable to the results obtained when oil, gas, or mining companies apply

advanced analytics to their processes. By identifying the factors responsible for failure, the analysis helped prevent significant yield loss early in the production process.

Applying advanced analytics to R&D

Our experience suggests that most semiconductor companies have inefficient R&D operations. Some 80 percent of development projects do not meet their initial schedules, often because teams overestimate their productivity and underestimate the complexity of their projects—the level of effort and resources required.² Our study of more than 2,000 integrated-circuit (IC) projects, for instance, showed that companies often drastically underestimated staffing requirements in the early and late stages but over-staffed in the middle of the cycle in reaction to the previous dearth of resources. Partly because of such inaccurate estimates, IC projects frequently struggle to meet their budget and time-to-market targets.

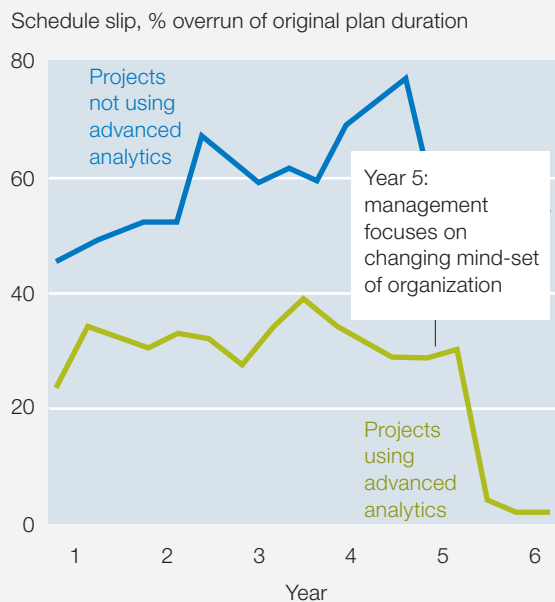
R&D's inefficiency may also result from an absence of rigor when companies assess their performance or measure their success. Semiconductor companies, for example, have few reliable metrics for gauging complexity or evaluating productivity (of individual project teams or the research effort as a whole). Similarly, they lack an adequate framework for assessing new market developments and estimating the needs of projects, so they base their allocation of R&D resources on instinct and history. But our research shows that this method is problematic: companies usually underestimate the time needed for completion by at least 15 percent and sometimes by as much as 400 percent.

Advanced analytics can make R&D more efficient by replacing instinct and guesswork with a fact base for decision making, thus ensuring that resources are deployed to the right projects and used optimally throughout the project life cycle. Companies can, for example, improve R&D's effectiveness and efficiency by statistically modeling the complexity of projects (such as the impact of adding a certain type of resource) and determining the best staffing levels. In fact, our research shows that semiconductor companies can reduce the time to market of their IC projects by up to 10 percent (Exhibit 2). Advanced analytics also plays an important role in streamlining R&D processes, optimizing product portfolios, and helping business leaders reduce costs.

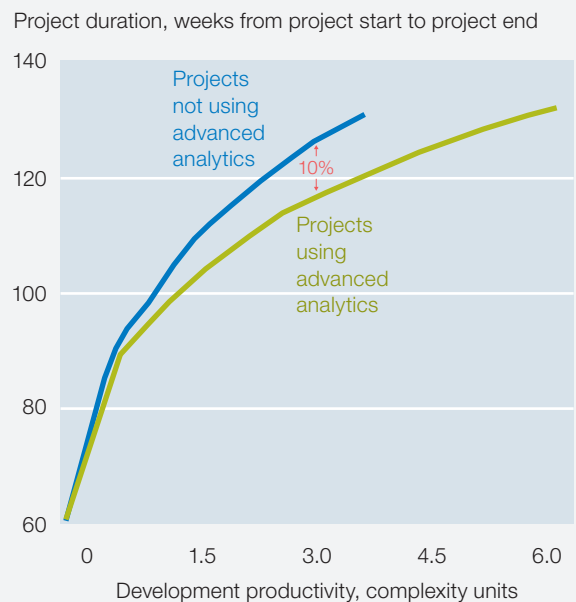
Exhibit 2 Using advanced analytics in R&D decision making can lead to many improvements.

Number of projects before using advanced analytics: **129**
 Number of projects after using advanced analytics: **80**

Schedule overrun of integrated-circuit (IC) projects reduced to <10%



Duration of IC projects reduced by up to 10%



Note: All 209 projects finished and released to volume production.
 Source: McKinsey analysis

Moving from seat-of-the-pants decision making to rigorous analytics

In the absence of analytics, R&D staffers at semiconductor companies sometimes make poor investment decisions. Design engineers, for example, might focus on incremental improvements to existing products, since they are inclined to believe that such changes will increase sales, rather than on developing new solutions. But in many cases, relatively small upgrades produce few returns. To counter such personal biases, advanced analytics offers a range of data-based tools that can examine customer and product segments at a detailed level, sharply improving the cost-effectiveness and returns of investments. The resulting insights also help companies to assess the risks of various investments and to balance their research portfolios.

In one case, a semiconductor company used advanced-analytics tools to conduct a segment-specific analysis of two important areas: market opportunity (such as potential market size, projected growth rate, and margin) and competitiveness (including the number of companies in the market, the level of differentiation among product offerings, and the customer's willingness to change suppliers). In addition to uncovering hidden opportunities for growth, the assessment transformed the company's R&D portfolio strategy from a reactive, seat-of-the-pants process into a far more insightful, objective, and predictive one based on solid numbers.

Improving R&D processes to increase efficiency

Advanced analytics can also improve many R&D processes, including some cumbersome, time-consuming, or error-prone tasks. A large chip manufacturer, for example, wanted to improve its time to market by at least three months and to stabilize development costs, which had been increasing at over 25 percent annually. Problems often arose because the company had not automated the process for comparing the results of large simulations.

It also had difficulty assessing the quality of its test inputs. To address these issues, the manufacturer used advanced-analytics tools to automate its design-verification process, with the goal of reducing the number of iterations for regression testing and improving the quality of the test. Automation helped the company shorten the product-development cycle by one to three months, thus generating an additional \$100 million in revenue. The company also eliminated many development costs, saving over \$3 million annually.

Optimizing resources throughout the project life cycle

Companies can create predictive models for R&D projects by employing proprietary advanced-analytics tools that use chip- and block-level parameters (for instance, node, power, transistor count, and memory) as proxies for design complexity. The models can determine how each parameter correlates with the completion of projects on time and on budget. They also allow companies to gain an objective view of their R&D performance compared with best-in-class benchmarks—a straightforward definition of what outstanding productivity looks like. Such root-cause analyses can help to explain the gap between current and best-in-class performance and to identify specific drivers for efficiency and productivity.

Predictive R&D models provide rapid insights that let companies adjust a project's staffing or make other changes in real time, thereby increasing the efficiency of R&D investments. For instance, while projects are under way, companies can use advanced-analytics tools to simulate realistic scenarios—such as altering the number of sites or the size of teams—and thus predict their impact on time to market and other important variables. Again, advanced-analytics tools help minimize risk, since they allow companies to see the potential implications (including the costs) of new strategies before they are implemented.

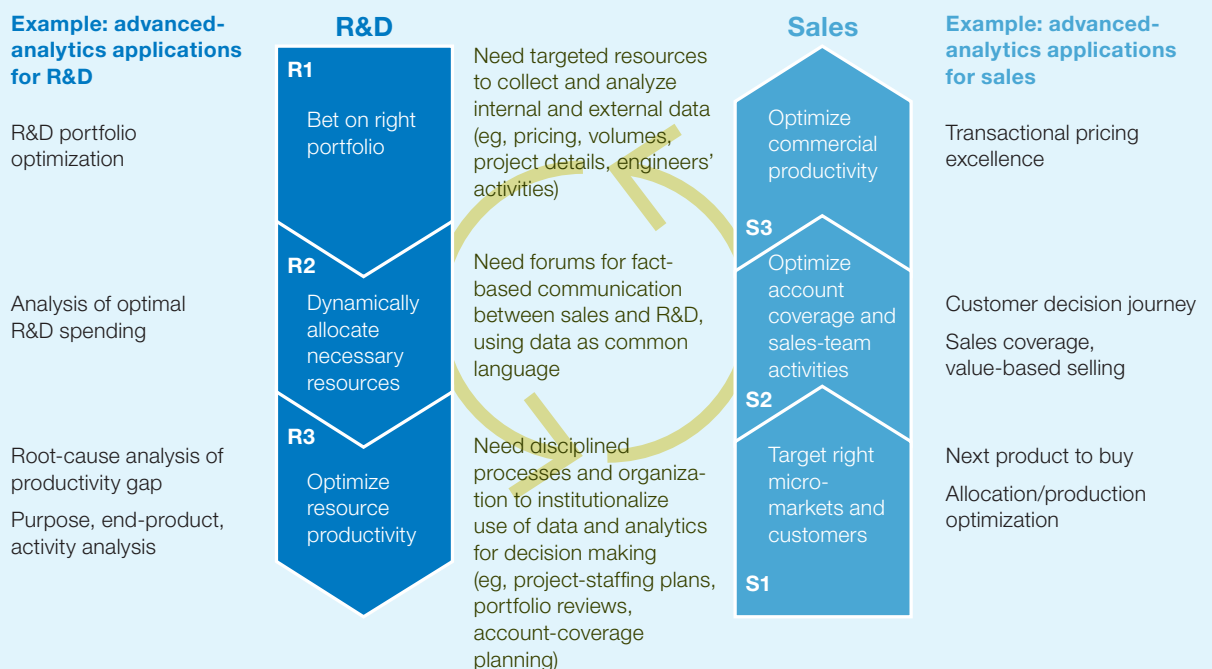
Consider the results obtained when a company in the advanced-manufacturing industry used an advanced-analytics tool to improve the productivity of over 7,000 engineering personnel responsible for designing and implementing complex electro-mechanical projects. The tool analyzed internal data at all levels—for instance, information on the project itself, team collaboration, and personnel—and determined which factors led to the best outcomes for the business in multiple areas, from team

resourcing to communication frequency to project management. Among other findings, the tool revealed that productivity was lagging because the company often pulled engineers away from their projects when unexpected problems arose. This practice alone created a 7 percent net productivity drag (as measured by the number of engineering hours needed to complete a project) for the entire engineering department. In addition, the tool showed that large teams (of more than seven people)

Getting more from R&D and sales

In our experience, most semiconductor companies struggle to make R&D and sales truly complementary functions and coordinate them only about half of the time. Better integration can keep companies ahead of evolving markets and the competition, as well as provide a 7 to 14 percent revenue uptick within 12 to 18 months—potentially with more to follow. The exhibit shows advanced-analytics strategies that can not only improve R&D and sales but also assist with integration by providing objective data. It also outlines the requirements for optimal implementation.

Exhibit Some best-in-class semiconductor companies leverage advanced analytics to drive organic growth.



Source: McKinsey analysis

were associated with diminishing returns and that global teams working around the clock could actually be counterproductive. By making some organizational and operational changes, the company identified improvement opportunities that could increase productivity by 27 percent.

Applying advanced analytics in sales

As with R&D, semiconductor companies frequently lack analytical rigor when they make investment decisions in two other important areas: pricing and sales coverage.

Pricing: From intuition to data

Sales teams often ignore or don't have critical pricing data, including historical information on each customer and market segment. They typically treat markets as whole entities rather than examining specific customer relationships or trends related to product demand. Sales teams may, for example, pursue opportunities in the automotive or consumer-electronics industries instead of identifying potential high-margin microsegments within these markets or the portfolios of individual customers. The lack of a numbers-based framework also encourages companies to rely solely on their executives' instincts to set prices. That often leads to inconsistent or ineffective discount policies and to an overreliance on poorly targeted cost-plus strategies.

Advanced analytics can help semiconductor companies bring new rigor to pricing. For instance, one player discovered through advanced analytics that it had set its list prices so low that customers rarely attempted to negotiate. By adopting a detailed analytics-based approach, combined with new statistical-analysis tools that can clean up and analyze volumes of transaction data, it optimized its pricing and its customers' willingness to pay. In this way, it captured more revenue with no loss in sales volumes. In another case, advanced analytics showed a semiconductor company that it was offering

different discounts for deals of similar size and setting minimum prices regardless of sales volumes. The company then created a more rational and profitable discount program—again, with no loss in sales volumes.

A number of advanced-analytics tools, some shown in Exhibit 3, can help semiconductor makers analyze pricing. In addition to proprietary analytics programs, they can use algorithms, heuristics, web crawlers, and scientific data-collection surveys to amass, sort, clean up, and normalize different kinds of information into actionable insights. To fill in gaps in data on the competition, for example, they can take advantage of web-crawling tools that collect publicly available pricing data for more than 100,000 components. Clustering algorithms can segment customers and markets in objectively meaningful ways. The systematic use of these and other analytical approaches can increase a company's return on sales by up to seven percentage points.

Account management: A detailed look at the customer base

Many industry players base their key-account-management strategies on current revenues. That sometimes leads them to overinvest in existing customer relationships and underfund high-potential prospects. Advanced analytics could help these companies make more rational, well-considered decisions about sales coverage. For instance, advanced analytics showed one chip maker that it was devoting only 45 percent of its sales team's resources to a customer segment that produced 50 percent of its revenues. Meanwhile, 30 percent of the resources went to a less promising segment that generated only 20 percent of its sales. The company also determined that 70 to 80 percent of its average salesperson's time was devoted to tending accounts it had already penetrated and only 20 to 30 percent to winning new business.

Exhibit 3

Advanced-analytics tools can assist with various sales tasks, including transaction pricing.

Barriers to value capture

Mitigating advanced-analytics tools/techniques

Unclean data

Cleansing heuristics to eliminate fuzzy logic
Clean-up algorithms for transaction data

Lack of updated competitor data

Web-crawl tool to collect publicly available competitive pricing data for >100,000 parts

Limited transaction data with respect to large, complex part portfolio

Multivariate regression algorithms to identify and quantify statistically significant influencers of price

Subjective customer and part segments

Clustering algorithms to segment customers and parts into objective, homogeneous segments

Lack of availability of timely and relevant data at time of decision

Web-based/mobile-based tool to provide real-time pricing guidance

Source: McKinsey analysis

Another component manufacturer used data mining and analysis to assess the attractiveness of its products against those of competitors. The company found that customers regarded roughly half of the 5,000 products it examined as leading edge, 35 percent as generic, and 15 percent as pure commodities. It combined this information with pricing and share-of-wallet data to create an index of product strength and to deploy its sales force more effectively.

Advanced analytics can also improve account management by pinpointing a company's top customers, as opposed to more routine transactional ones or targets that turn out to be unattainable. By combining these customer analyses with product analyses, semiconductor companies can get a clearer view of how and where they need to compete.

How to make this happen

Companies can increase the likelihood that advanced analytics will gain traction by ensuring that all programs have four core elements: robust and actionable data, enterprise-wide support, well-trained analytics teams, and a suitable IT infrastructure.

Ensuring robust and actionable data

Our experience with clients has demonstrated that any solid advanced-analytics effort must start with the creation of a robust data set that encompasses all necessary inputs, can be stored effectively, and allows for easy retrieval. While this may sound simple, it can be extremely challenging. For example, in manufacturing, most leading-edge fabs have well-defined data models for individual tools. Each model can require sorting through 300 or more individual variables per tool or process step (repre-

Whenever possible, companies should compile their own information instead of relying on possibly incomplete data from outside sources.

senting such diverse inputs as electrical current, resistance, temperature, pressure, and robot location). Some of these variables may be collected every few seconds or milliseconds, resulting in large volumes of data. In addition, inputs for each model may differ, depending on what equipment or sensors are in use or whether data are being analyzed while a process is under way or after it is completed.

Many companies also have difficulty aligning data sets or making comparisons. Information is typically generated and stored in different formats, and much of the data are “noisy”—they cannot be understood and interpreted correctly by machines, because they include unstructured text or have other limitations. Overcoming these challenges requires companies to go back to the original data and ensure that they are aggregated in a consistent way. This may require companies to work with vendors, to invest in additional tools, and to build internal data-collection and management capabilities.

To ensure robust data and facilitate analytics, companies must carefully extract, validate, and visualize information before creating any models. They should also ensure that they examine complete information rather than relying on aggregate data sets that capture averages or on a sampling of inputs, since such methods can lead to false positives or missed patterns. Manufacturers, for instance, would need to compile all sensor, process, inspection, and metrology data.

Whenever possible, companies should compile their own information instead of relying on possibly incomplete data from outside sources. For instance, equipment vendors often gather and manage data and then provide fabs with a summary analysis and statistics, rather than supplying full data sets along with process and outcome data. If companies rely solely on this partial information when they conduct analytics, their results may not be accurate.

Companies must also ensure that their data are gathered and stored effectively, so they cannot take common shortcuts. When data are lost, for example, many companies just replace the missing information with averages or extrapolated data points. While that approach might save time, it reduces the quality of the analytics, so companies should identify and fill any gaps.

[Building support at all levels](#)

Across industries, companies have made large investments in advanced-analytics initiatives only to receive little in return. All too often, the problem is that the top leadership views advanced analytics as a tactical responsibility for midlevel managers rather than a strategic priority for the entire company. In the absence of a corporate directive, some business units embrace advanced analytics, while others treat it as an afterthought because they question its worth. Companies can avoid such situations if their boards have one or more members who actively advocate advanced analytics. Management tactics that may help build support include the following:

- newsletters that describe interesting advanced-analytics findings and show how they improved revenues and other metrics
- routine meetings where data scientists ask for input from all stakeholders, including frontline employees—information that is then used to develop customized approaches
- dashboards that show preliminary results from recent analyses; companies can then ask employees if the information would help them improve performance
- quick wins gained by applying advanced analytics to areas that do not require a large-scale transformation, such as pricing or customer acquisition

Training, staffing, and supporting analytics teams

Companies usually need new staff to handle advanced analytics, especially an experienced leader who can help define the strategy. Every analytics team should also include the following roles:

- data engineers responsible for managing data and preparing data for analysis
- people who create models to predict business outcomes, optimize processes, segment customers, and accomplish other important tasks
- business liaisons who serve as a bridge between the modelers and the decision makers on the business side; for instance, they may work with the leadership to identify important business problems that models should address

Although many companies now situate analytics teams within IT, this placement may be problematic because all advanced analyses should focus on solving

important business issues and changing employee behavior. Data engineers and modelers, especially those who are just out of school and have little workplace experience, may not be able to identify the most important business problems. Placing advanced-analytics teams within business units, where they can serve as internal consulting groups and interact closely with business liaisons who understand operational issues, may be preferable.

Creating the right IT infrastructure

Some companies already collect the internal information necessary for advanced analytics but do not organize it well or make it readily accessible. They can address these problems by creating data marts and related support teams. While companies may be tempted to include all easily accessible information in the data marts, it is better to restrict them to data that will help address pressing business needs. Such a data mart will need extraction, transformation, and loading (ETL) routings that continuously pull information from different sources and link it together using common identifiers. Companies must establish a rigorous process to ensure the data's quality, since even a single error in the ETL may decrease confidence in the data mart.

Analyzing external data may pose a greater challenge, since many businesses capture incomplete information on the activities of their competitors or the behavior of their customers. To close this gap, companies could consider establishing a market-research team specifically to scan public sources for competitive insights. If such teams feel that public sources are inadequate, they can work with third-party aggregators to gain more insights.

In certain cases, companies may need to seek assistance from data experts to manage information. The point when they do will vary greatly, depending on their technological capabilities and the type of

data under analysis. For instance, most businesses can manage structured (usually SQL-based) databases effectively, but they may have more difficulty working with unstructured data, such as text, images, and videos.



The payoff from advanced analytics is a more aggressive, alert, and competitive enterprise that makes better decisions and uses more accurate and reliable information to analyze risk. Once the infrastructure is in place, fabs can create better predictive models to enhance manufacturing. Meanwhile, sales and marketing teams can more efficiently sell what R&D teams and design engineers create. ■

¹ Hadoop is an open-source software framework for the distributed storage and distributed processing of very large data sets.

² See Aaron Aboagye, Dorian Pyle, and Alexander Silbey, “By the numbers: R&D productivity in the semiconductor industry,” *McKinsey on Semiconductors*, Autumn 2014, mckinsey.com.

The authors wish to thank Aaron Aboagye, Eric Egleston, Taras Gorishnyy, Mithun Kamat, Eoin Leydon, Mark Patel, Dorian Pyle, and Bill Wiseman for their contributions to this article.

Gaurav Batra is a principal in McKinsey’s Silicon Valley office, where **Nick Santhanam** is a director; **Zach Jacobson** is a consultant in the New York office.

Copyright © 2016 McKinsey & Company.
All rights reserved.