



# Seven rules for spinning analytics straw into golden results

Richard Kelly, Subu Narayanan, and Mark Patel

While IoT-enabled advanced analytics could be worth trillions to manufacturers, turning insights into outcomes requires more than just the right technology.

**Richard Kelly** is a partner in McKinsey's Stamford office, **Subu Narayanan** is an associate partner in the Chicago office, and **Mark Patel** is a partner in the San Francisco office.

*The authors wish to thank Markus Hammer, Malte Hippe, Olivier Noterdaeme, and Joris Van Niel for their contributions to this article.*

Manufacturing may generate a wealth of data, but companies' efforts to use those data to drive performance improvement have only scratched the surface thus far. But now, lower-cost sensing, better connectivity, and ever-increasing computing capabilities are combining to push analytics and intelligence far beyond what was possible in the past.

The challenge is knowing how to start—and how to achieve measurable, sustained impact. Our work with manufacturers around the world suggests that by following seven golden rules, companies can start capturing the benefits of IoT-enabled advanced analytics more quickly, and build a solid foundation to keep improving.

## Rule #1: Start simple, with existing data

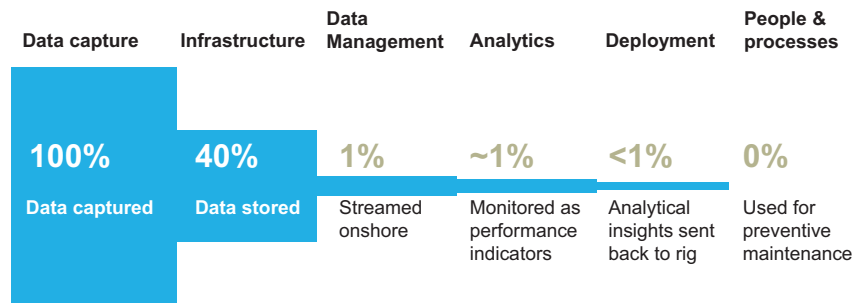
With the increased buzz around the Internet of Things (IoT) in manufacturing,

many companies are excited about deploying thousands of low-cost sensors within their operations. While we do think this idea shows value, our experience shows that most of the data currently being generated is unused (exhibit). Simple analytics, done right, and with the existing treasure trove of data can yield tremendous value for manufacturers in the near term. Those early victories help win the hearts and minds of frontline employees while strengthening a data-driven decision culture—and the business case for further advanced-analytics investment.

For example, digitizing performance management—such as through real-time data visualization of human and machine performance—requires minimal resources, as it relies on simple, rapidly deployable solutions. Yet its easily quantified results can serve as the gateway to rapid improvement and management buy-in.

## Exhibit. Most real-world data go unused.

EXAMPLE: OFFSHORE OIL RIG



One manufacturer seemed to be performing well: it was already mature in its implementation of lean management and had a robust problem-solving culture in place. Nevertheless, sensing an opportunity to improve still further, it deployed an analytical solution that applied sophisticated real-time analytics to existing (but previously unused) data, producing user-friendly visualizations. Up and running in just weeks at a capacity-constrained plant, the system delivered previously unavailable details to daily area huddles and operator-driven problem-solving sessions, revealing several previously unknown causes of slowdowns and minor stops. At the most important bottleneck, the ensuing changes increased overall equipment efficiency by 50 percent.

### Rule #2: Capture the right data, not just more data

Having the right data is more important than having lots of data. One basic-materials company invested several million dollars installing a “smart” manufacturing system that tracked more than a million variables. When the company analyzed 500 data tags from the system pertaining to a specific analytical use case, however, half of them were shown to hold limited or duplicated information. Another 25 percent of the data was discarded by a panel of process experts and data scientists as not being helpful for analytics. Further into the exploratory-analysis stage, the company found 20 critical variables—including

a key dependent variable—that were not being measured, making precise predictive analytics impossible. This formed the case for deploying new sensors in a targeted fashion within the plant, while the company used analytics to provide critical decision-support tools for the process engineers as a first step in a quest to increase yield by 1 percent.

### Rule #3: Don’t let the long-term perfect be the enemy of the short-term good

Missing data can threaten to stall analytics projects while they wait for a multiyear data architecture transformation. We acknowledge that capturing the full value of IoT-driven advanced analytics will require an investment in the technology stack. But companies don’t have to be bogged down by long IT projects. Minor investments can deliver much value.

One no-regret move is to develop a “data lake”—a flexible way to integrate data across an enterprise and overcome silo-based data management without full centralization. Although data lakes need strong governance and accountability for data definition and quality, they can democratize data access. Typically, data lakes provide data to different user groups either by permitting access to raw data or through data distillation, which affords access to pre-defined data structures.

The development approach required to implement analytics adds to the case for an alternative IT

architecture. Analytical experimentation and exploration require agile software-development methods with daily or weekly release cycles. This short cadence is often a challenge for established IT processes and data infrastructure. The solution is a parallel “fast-speed” IT and data infrastructure, often a cloud-based system offering a range of deployment environments and tailored databases.

Data lakes and cloud solutions get companies’ analytics efforts off to a faster start, allowing them to develop, test, and implement new use cases quickly. That helps in the creation of the necessary proof of concept before the wider rollout of new solutions. It is also a valuable way to build the organization’s analytical muscles as people become accustomed to new ways of working and decision making using analytics.

#### **Rule #4: Focus on outcomes, not technology**

Investment in digital products and solutions without knowing how they will deliver meaningful impact will lead to frustrating discussions with business leaders. An approach based on use cases can help (see sidebar, “Successful analytics uses cases”). When defining a use case, be sure to answer four fundamental questions together with their follow-ups:

- **What is the desired business outcome?** Is it a new business opportunity, a cost reduction opportunity, an increase in innovation capacity?
- **What are the value levers?** Should the focus be on energy savings, more-efficient maintenance, higher asset utilization, lower inventory, higher throughput?
- **What technical requirements must the proposed approach meet for it to scale across the organization?** Are new data sources needed? How will the solution integrate with legacy IT systems? How will

we handle the volume of data securely? What analytical techniques will be used? What new dashboards are required?

- **How will the approach fit into our existing processes?** Who will use the new system? What behaviors and decision-making processes must change to take turn analytical insights into business outcomes?


#### **Rule #5: Look for value across activities as well as within them**

While advanced-analytics methods have been applied very successfully to many specific activities that take place within the four walls of a manufacturing plant, much of the value of digitization lies in the whitespaces between organizational siloes—by bridging the gap between design and manufacturing, manufacturing and the supply network, and finally connecting with the end user. A manufacturer of highly specialized equipment recently conducted a “digital thread diagnostic” that identified more than \$300 million of actionable productivity improvements that could be realized with using better data flow between design and manufacturing, real-time performance management, and other levers.

#### **Rule #6: Break out of the pilot trap**

A pilot project is a powerful, and important, way to demonstrate the value of advanced analytics, build momentum, and encourage buy-in. Capturing that value, however, means scaling the approach across the entire company. That’s hard, and failure to scale can turn supporters into critics very quickly. Leaders must therefore think through the full end-to-end journey needed to turn attractive use cases into widespread impact. Some common pitfalls:

- **Focusing on the technology or approach, rather than the real source of value.** When defining the use case, it is important to start with the true source of value, which is often the user or customer needs. A software tool is almost never a panacea; moreover, the selection



## SUCCESSFUL ANALYTICS USE CASES

The following four use cases, each applied within the four walls of a plant, have delivered significant value for companies in multiple sectors.

**Real-time performance tracking:** At the most basic level, manufacturers are seeing improvement just by tracking critical metrics in real time, applying simple analytics and visualization, and embedding the results into their regular root-cause problem-solving processes. As the solutions get implemented, a culture of data-driven decision making takes hold at all levels within a plant.

**Multivariate process optimization:** Several companies have adopted an integrated approach that optimizes a process's yield, energy usage, and throughput to maximize profit per hour, and has proven both simpler and more effective than alternatives. One European chemicals manufacturer, for example, found that it was seeing almost no impact from an expert system it had installed to stabilize its process and reduce energy consumption. The reason turned out to be that the system was too complex to use. The company then switched to a plant-level profitability-optimization approach, which used advanced-analytics tools and neural-network models to simulate its processes, and profit per hour as the primary performance indicator. The resulting analysis showed that optimizing a few important parameters in real time could reduce energy costs by 4 to 5 percent. Better still, the changes involved only minor process modifications that required little time or capital investment to implement.

**Predictive maintenance:** The concept of proactive and preventative maintenance is not new, but the ability to manage massive amounts of data through sophisticated analytical techniques can dramatically

improve failure predictions and asset uptime. A resources company, for example, discovered that despite a preventive-maintenance strategy for its fleet of heavy trucks, many critical components—such as transmissions—routinely failed before the planned maintenance interval. As a result, three-quarters of the company's maintenance effort was corrective. Applying advanced-analytics techniques revealed that a single, easily measured operating characteristic was a good indicator of developing transmission problems. By monitoring changes in that characteristic in real time, the company could alert operators days in advance, and with accuracy higher than 80 percent. Similar results have been seen in chemicals, semiconductors, transportation, oil and gas, and other industries.

**Labor productivity:** Advanced analytics applies equally well to the factors that drive human productivity, as one company found with its engineering staff. The approach pulled data from a wide range of sources, such as project records (for day-to-day project plans and timesheets), product-life-cycle-management systems (bills of materials and document version-control entries), knowledge-management databases (wikis and discussion forums), supplier data (vendor data, requests for proposals, and requests for quotations), and metadata (such as emails, calendar entries, and employee demographics). Putting this information through a machine-learning tool revealed many counterintuitive insights and quantified other "known truths." For example, project starts and stops resulted in an 8 percent loss in productivity. A 7 percent drop in productivity was observed for every 10 percent utilization increase over 70 percent—and another 7 percent drop for every engineer added to a project beyond seven. The analysis identified total productivity-improvement potential of between 15 and 25 percent.

of the right technology depends on the universe of use cases a company wants to deploy.

- **Solving for one use case at a time:**

Focusing too closely on a single use case can lead to choices that limit scalability later on. Important technical requirements to achieve scale include advanced operational and analytical data architecture, such as data lakes and data-search layers, together with IoT platforms, tools for digitization and analytics, and a repository of modeling tools and techniques.

In a factory setting, the right IoT platform can help analyze many functions regardless of the specific application, and thereby scale a variety of use cases at once. The underlying technology needs are essentially the same whether the organization is trying to optimize yield or to predict failure of critical equipment. An IoT platform can provide common capabilities for computing power or storage or security, while reducing the cost of developing and maintaining applications.

In assessing IoT platform needs, companies should bear five factors in mind: the application environment and the proposed platform's connectivity to existing IT infrastructure; the platform's ability to ingest high-velocity and -variety data streams while providing context to the data; its compatibility with a broader enterprise-cloud strategy; data sovereignty and security questions; and its capacity for edge processing and control, meaning it allows for processing and data storage close to the source, rather than only centrally.

- **Prematurely celebrating success:**

Companies should think through the entire end-to-end journey, beyond the technical elements needed to achieve scale beyond a single proof of concept. Data-

governance issues such as domains, critical data elements, accountability models, and role definitions can pose tricky organizational and personnel questions, especially given the new analytical and technical positions that may be required. And analytics-generated insights must be integrated into existing workflows, often with attendant changes to business processes.

- **Nailing the technical solution, but forgetting the people:**

Technology is exciting—but it's people who capture the impact. While analytics can point to the right answer, people must act differently to capture the impact. Capturing the digital opportunity is a team sport, requiring close, cross-functional collaboration. A team of people with deep process knowledge, analytical acumen, and IT experience must work together to frame the problem, translate the business problem into an analytical problem, and define the right system and technical requirements from an IT perspective. Translating the analytical output into a form that can be used at the front line, and changing frontline behavior to make use of that new information, requires knowledge of human factors, persuasive design and change-management experience. Some companies find it useful to create a new role—digital translator—at the intersection of process knowledge, data science, and IT, to bring the required cross-functional teams together and steer the analytics effort from concept to bottom-line impact.

To avoid these pitfalls, companies need a structured approach to manage their analytics efforts, identifying and managing a pipeline of use cases, for example, and building the right technology stack. Once a use case is selected, companies need to systematically plan, pilot, scale, and embed analytics into their everyday processes through large-scale change management and capability building.

## Rule #7: Build your capabilities

The application of analytics at scale will require organizational changes, too. For example, a company needs to define its talent strategy as new roles and new career paths emerge. There will be a need for data scientists, agile IT teams, and user experience (UX) designers, who play a crucial role in supporting real-world use of analytics. A persuasive design created with frontline involvement, is often the secret to high adoption levels for any analytical solution. Accordingly, UX professionals should be involved from the moment a use case is designed, not asked to apply a visual interface after a solution has largely been built.

In addition, a company needs “translators”—multi-skilled individuals who can shepherd the process from end to end. Translators need deep business knowledge and the ability to get into the workflow of operations and maintenance teams. They must be comfortable with analytics and able to challenge data scientists. They must understand IT systems and design thinking. And they must be able to communicate impact to the

leadership team. That’s a very tough combination of skills to find.

In addition to these internal roles, a clear partnership strategy is important. There is an explosion of both big companies and start-ups with unique IoT capabilities. The successful companies will very quickly home in on their unique value proposition and partner in areas that help accelerate their capabilities.



The potential impact from IoT-driven advanced analytics is game changing. While it is easy for companies to get started and get some quick wins on the board, it is much harder to scale across the company and deliver consistent bottom-line impact. The most successful organizations will be those that think through all of the implications, invest in both technology and people, forge smart partnerships, and maintain sufficient leadership appetite to persist.

Copyright © 2017 McKinsey & Company. All rights reserved.