



THE BOSTON CONSULTING GROUP

USING AGILE TO HELP FIX BIG DATA'S BIG PROBLEM

By Fabrice Roghé, Erik Lenhard, Burt LaFountain, and Giulia Airaghi

AGILE HAS TRANSFORMED SOFTWARE development and taken hold in other business functions. But it has not been incorporated to the same degree in big data projects. Although organizations of all kinds are modernizing by using big data to power important decisions, the way they develop big data projects remains decidedly old school.

The reasons for the lag are many, but in many cases, the biggest reason is simple. Data engineers and scientists often do not have the depth of business acumen needed to fully address the key questions that their projects are meant to answer. They may be great at mathematical modeling, computer science, and statistics but less adept at detecting the specific business meaning in the data or assessing the real-world business changes needed to capture the value of their analytics.

At the same time, business executives may not understand data science well enough to fully appreciate how their key business questions could be answered. They often

underestimate what is possible and what is practical when implementing big data solutions. As a result, the possible outcomes of a big data analytics project, and what is necessary to generate those outcomes, can easily get lost.

Traditionally, IT departments and business executives have used classic project management methods to overcome these shortcomings. Such techniques include establishing milestones and scheduling meetings to align business and technical staff. But those classic methods may not be enough to compensate for scientists' specialized knowledge and for business leaders' limited understanding of data science.

Organizations can overcome these challenges, however, by incorporating agile practices throughout big data analytics projects. If they do, they can better focus on both their internal and external customers. Using agile can also help employees feel empowered, which can lead to higher quality data and better project outcomes.

Big Data's Big Problem and the Agile Solution

At their best, big data analytics detect patterns that would require considerably more time and effort to uncover using traditional analytics tools. The widespread use of big data analytics has powered breakthroughs in areas as varied as medical diagnostics, people management, and the ways organizations respond to consumer behavior.

Data scientists start big data projects with a theory about a business problem, such as how to predict demand for a car model on the basis of new features and past sales or how to determine how many employees an organization should hire to staff a new venture given existing personnel and previous staffing levels for similar projects. They then build algorithms to test the theory using one or more forms of artificial intelligence, machine learning, optimization, or traditional statistics on a massive scale. If the theory is shown to be false, they may continue refining and testing it until they reach a valid conclusion. Or they may drop it and move to a different or more pressing business problem.

The results of big data analytics can be remarkable. In observing many companies, though, we have also seen significant failure rates, particularly when organizations attempt to roll out those analytics on a wide scale. When problems do occur, deliverables such as any expected insights or process improvements may not materialize. More often than not, the fault lies not with the data but with the methods used to verify, process, and act on it. (See “How to Avoid the Big Bad Data Trap,” BCG article, June 2015).

The heart of the problem is the manner in which big data analytics are developed. Most are built sequentially, applying the waterfall method of project management traditionally used in software development. In the waterfall method, data scientists acquire, verify, and integrate data; develop a model or algorithm to test it; run the test; and then either act on the results or continue refining the model. Work on one task waits until the preceding task is finished. But that process is inefficient. In many cas-

es, people spend more time sitting in meetings and managing handoffs than they do on actual data-related activities. The work itself may be misdirected and wasteful. The final product is often late and difficult for a lay executive to understand, and the impact is less than anticipated.

Frustrations with the waterfall method eventually led software developers to improve the process by adopting agile ways of working. Agile calls for working in a way that is iterative, empirical, cross-functional, focused, and continually improving. (See “Five Secrets to Scaling Up Agile,” BCG article, February 2016). Common agile methods include assembling cross-functional teams, which improve communications and reduce handoffs, especially when team members work in the same location. They also include developing minimum viable products (MVPs), rapid updates, and frequent feedback to ensure that the finished product delivers on expectations and goals. (See Exhibit 1.)


Using Agile in Big Data Projects

Agile has shown great promise in many fields besides software, including financial services, marketing, and consumer goods. (See “Agile to the Rescue in Retail,” BCG article, October 2018, and “Taking Agile Transformations Beyond the Tipping Point,” BCG article, August 2018.) In client engagements, we have seen agile empower teams to do their best work while assuring that they are aligned with an organization's strategic goals. Given those successes, we believe that agile could bring several specific benefits to big data projects.


Rapid Experimentation. Historically, testing occurs near the end of big data projects, which means that business executives might not see results until then. For example, a team building a predictive analytics model to help sales people convert leads might wait until late-stage testing to show executives the results. However, offering results so late in the process could lead to unclear expectations for the work to be done, methods to be used, and possible outcomes. With agile,

EXHIBIT 1 | Agile Speeds Up Traditional Big Data Projects and Improves Outcomes

Traditional



Agile



Source: BCG analysis.
Note: MVP = minimum viable product.

projects are broken down into manageable chunks that can be built and tested quickly. Teams develop and test MVPs continuously. If the data analytics do not yield the expected results, business executives find out right away and can correct their course. They could ask the team to analyze different data, make other modifications, or even, in some cases, abandon the project—all moves that save time and money compared with other methods.

A specialty retailer took this concept of rapid experimentation to heart when it convened an agile team that adopted the motto, “Get 1% better each week.” The agile team—which was composed of personnel from the data engineering, data science, marketing, and creative functions, among others—was tasked with developing innovative big data marketing and sales solutions. To do that, they ran agile sprints that produced incrementally new omnichannel marketing programs every seven days. The new programs led to direct increases in revenue in a short period of time.

Rapid experimentation is appropriate not only for big data analytics algorithms and the data that a project is based on but also to ensure that an organization can understand and act on the results. For this reason, in addition to the algorithm and output, a

big data project MVP includes the business and behavioral changes needed to attain real results. For example, developing an algorithm to improve scheduling and dispatch for service technicians could also include developing a different process for notifying customers of upcoming appointments.

Early Customer Feedback. The overriding goal of big data projects is not to build brilliant mathematical models but to solve practical business challenges or discover insights leading to actions that could benefit customers. That makes it important to include customers in the process. If a project is for an external client, a representative of the client could be embedded with the team. If a team is working on a big data project for an internal client, a member of that department might be on the team. When a European oil refinery created a big data application that its engineers could use to optimize the maintenance cycle of their key equipment, for instance, some process and maintenance engineers were assigned to the teams that developed the app.

Prioritizing Value. Accomplishments that add value without increasing cost take precedence over completing tasks in a predefined order. If the team determines that including a particular feature will take significantly longer than expected without

providing a lot of additional value, the product owner—the team member responsible for representing the customer—can take this into account. The product owner can drop it and move on to lower-cost or higher-value items on the project backlog, a prioritized list of work items.

Cross-Functionality. Traditional big data projects fail most often for reasons that are largely unrelated to data analysis. In our experience working with clients, 70% of a cross-functional team’s efforts reach beyond strict analytics into the business processes, operational behaviors, and types of decision making that the analytics suggest. To accommodate that scope, big data project teams typically include personnel with a variety of backgrounds. (See Exhibit 2.) The team can make decisions without members needing approval from their individual bosses. Team members work out tradeoffs, conflicts, and compromises in real time, which explains why it is so critical for them to work in the same location—preferably in the same room. At the same time that they are working on algorithms and data, teams may also be making changes to operating models and business processes.

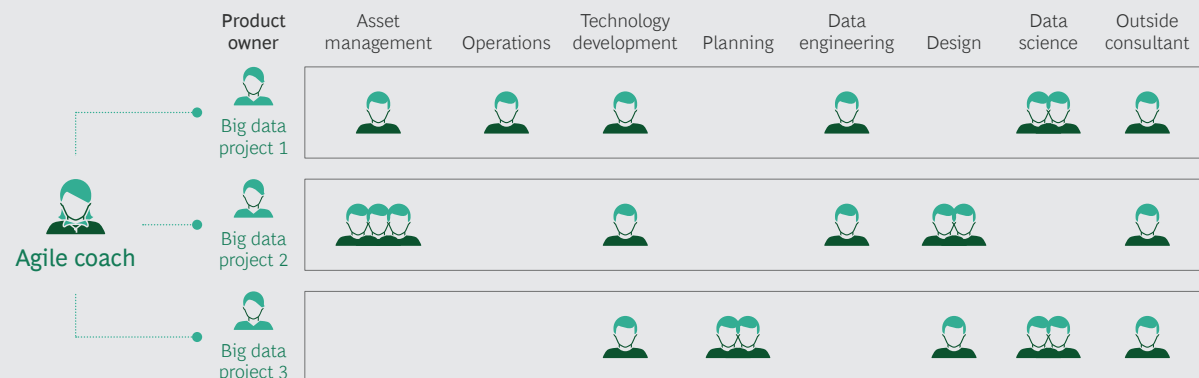
The European oil refinery mentioned above took such an approach when it incorporated agile ways of working into three big data projects launched as part of a larger digital transformation initiative. Each project had a dedicated scrum team with a scrum master, product owner, and person-

nel from such key functions as operations, technology development, asset maintenance, and IT. Data scientists from an outside advanced-analytics consulting firm were included on the teams. The teams and agile approach helped the refinery produce multiple MVPs within a four-month span and industrialize a final product considerably faster than it could have done in the company’s typical product development cycle. Multidisciplinary scrum teams also contributed to a more collaborative corporate culture and drastically increased employee engagement.

People Empowerment. In a traditional big data project, a project manager decides which priorities are most important and how they will be met—even though he or she may not understand the development process. When that power is delegated to a team, people become more engaged in their work and are more invested in the outcome. Unlike data scientists who might work on multiple engagements, for example, agile team members are not staffed on several projects simultaneously. Rather, they devote all their time to the team, thus becoming more invested in the work. This singular focus also builds accountability.

It’s no wonder that companies that have adopted agile are more successful than others in attracting digital talent and younger workers—two groups of people who prioritize work that gives them a sense of purpose. (See *How to Gain and Develop Digital Talent and Skills*, BCG Focus, July 2017.)

EXHIBIT 2 | Agile Teams for Big Data Projects Have a Cross-Section of Expertise



Source: BCG analysis.

The Secrets to Using Agile in Big Data

Agile and big data may sound like a perfect pairing, but getting them to work together is not as easy as it may seem. To be successful, keep several critical factors in mind.

The algorithm is not the finished product.

An algorithm can be a thing of beauty or a waste if it cannot deliver results-oriented output in a way that a business department or organization can understand. Agile teams should develop dashboards, infographics, or other visuals that readily communicate the results of their findings. They must also include as part of their work the end-to-end operational and behavioral changes that are necessary to get real results.

MVPs are distinct from prototypes. Prototypes come first and generally are built with historical data in order to verify that an algorithm can do what it is supposed to do. If a prototype works, it is used as the frame for a more all-encompassing MVP that could potentially be released to end users. An MVP also includes up-to-date data, a user-friendly interface, core features, and operating instructions. And since it has to work in a business context, it includes relevant changes to processes and operating models as well.

Stakeholders must accept imperfection.

The iterative nature of agile development means that works in progress might not look great or perform as well as possible. Nevertheless, and despite their flaws, they

may indicate signs of progress toward a satisfactory solution. Accepting such imperfection may require stakeholders—who, in the past, saw only near-final versions—to shift how they think about big data analytics projects. By encouraging trial and error, stakeholders improve the odds that a project will move successfully from MVP to full-scale production.

Include more than just data science

personnel. Agile teams should include a mix of talent assembled on the basis of need rather than on a standard structure or past experience. As a general rule, it makes sense to staff a project team with data engineers who can prepare the data, data scientists who can conduct the analysis, designers who know how to present the data, and a variety of business personnel who are familiar with the project's business objectives and implications for existing processes. The goal is to blend people's talents into a whole that is greater than the sum of its parts.

WITH SO MANY positives to be gained from incorporating agile into big data projects, companies might be tempted to dive in right away. But getting the best outcomes takes substantial planning, including a thorough examination of what is needed and how it would affect existing personnel and processes. Conducting a pilot is a good way to start. If it succeeds, agile can be added to more big data programs, a step that requires thinking about how to set up teams and get customer feedback.

About the Authors

Fabrice Roghé is a senior partner and managing director in the Düsseldorf office of The Boston Consulting Group. You may contact him by email at roghe.fabrice@bcg.com.

Erik Lenhard is an associate director in the firm's Munich office. You may contact him by email at lenhard.erik@bcg.com.

Burt LaFountain is a BCG Gamma principal in the firm's Boston office. You may contact him by email at lafountain.burt@bcg.com.

Giulia Airaghi is a consultant in the firm's Milan office. You may contact her by email at airaghi.giulia@bcg.com.

The Boston Consulting Group (BCG) is a global management consulting firm and the world's leading advisor on business strategy. We partner with clients from the private, public, and not-for-profit sectors in all regions to identify their highest-value opportunities, address their most critical challenges, and transform their enterprises. Our customized approach combines deep insight into the dynamics of companies and markets with close collaboration at all levels of the client organization. This ensures that our clients achieve sustainable competitive advantage, build more capable organizations, and secure lasting results. Founded in 1963, BCG is a private company with offices in more than 90 cities in 50 countries. For more information, please visit bcg.com.

© The Boston Consulting Group, Inc. 2018. All rights reserved.

For information or permission to reprint, please contact BCG at permissions@bcg.com. To find the latest BCG content and register to receive e-alerts on this topic or others, please visit bcg.com. Follow The Boston Consulting Group on Facebook and Twitter.

10/18