IT

Why IT Fumbles Analytics

by Donald A. Marchand and Joe Peppard

In their quest to extract insights from the massive amounts of data now available from internal and external sources, many companies are spending heavily on IT tools and hiring data scientists. Yet most are struggling to achieve a worthwhile return. That's because they treat their big data and analytics projects the same way they treat all IT projects, not realizing that the two are completely different animals.

The conventional approach to an IT project, such as the installation of an ERP or a CRM system, focuses on building and deploying the technology on time, to plan, and within budget. The information requirements and technology specifications are established up front, at the design stage, when processes are being reengineered. Despite the horror stories we've all heard, this approach works fine if the goal is to improve business processes and if companies manage the resulting organizational change effectively.

But we have seen time and again that even when such projects improve efficiency, lower costs, and increase productivity, executives are still dissatisfied. The reason: Once the system goes live, no one pays any attention to figuring out how to use the information it generates to make better decisions or gain deeper—and perhaps unanticipated—insights into key aspects of the business.

For example, a system that an insurance company installs to automate its claims-handling process might greatly improve efficiency, but it will also yield information for purposes nobody articulated or anticipated. Using the new data, the company can build models to estimate the likelihood that a claim is fraudulent. And it can use data on drivers' speed, cornering, braking, and acceleration—gathered in real time from sensors installed in cars—to distinguish between responsible and less responsible drivers, assess the likelihood of accidents, and adjust premiums accordingly. Yet simply putting the system in place won't automatically help the company gain this knowledge.

Our research, which has involved studying more than 50 international organizations in a variety of industries, has identified an alternative approach to big data and analytics projects that allows companies to continually exploit data in new ways. Instead of the deployment of technology, it focuses on the exploration of information. And rather than viewing information as a resource that resides in databases—which works well for designing and implementing conventional IT systems—it sees information as something that people themselves make valuable.

Accordingly, it's crucial to understand how people create and use information. This means that project teams need members well versed in the cognitive and behavioral sciences, not just in engineering, computer science, and math. It also means that projects cannot be mapped out in a neat fashion. Deploying analytical IT tools is relatively easy. Understanding how they might be used is much less clear. At the outset, no one knows the decisions the tools will be asked to support and the questions they will be expected to help answer.

Therefore, a big data or analytics project can't be treated like a conventional, large IT project, with its defined outcomes, required tasks, and detailed plans for carrying them out. The former is likely to be a much smaller, shorter initiative. Commissioned to address a problem or opportunity that someone has sensed, such a project frames questions to which the data might provide answers, develops hypotheses, and then iteratively experiments to gain knowledge and understanding. We have identified five guidelines for taking this voyage of discovery.

Deploying analytical IT tools is relatively easy. Understanding how they might be used is much less clear.

1. Place People at the Heart of the Initiative

The logic behind many investments in IT tools and big data initiatives is that giving managers more high-quality information more rapidly will improve their decisions and help them solve problems and gain valuable insights. That is a fallacy. It ignores the fact that managers might discard information no matter how good it is, that they have various biases, and that they might not have the cognitive ability to use information effectively.

The reality is that many people—including managers—are uncomfortable working with data. Any information-based initiative must acknowledge that. It must place users—the people who will create meaning from the information—at its heart. It should challenge how they do or do not use data in reaching conclusions and making decisions, urging them to rely on formal analysis instead of gut feel. And it should question their assumptions about customers, suppliers, markets, and products.

Achieving those shifts in mind-set was the objective for a large European-based manufacturer of chemical products. The company, which we'll call ChemCo, had grown rapidly through acquisitions and had a new CEO intent on developing a coherent picture of customers. He also wanted managers and employees at all levels to use data to enhance their understanding of the business and to make decisions more effectively.

He and the executive team promoted the view that being data-driven and creating usable information should be part of "business as usual." Building a large CRM system right away, they believed, would send the wrong message—namely, that a new system would change how managers used and shared customer information. They were also concerned that the

initiative would be seen as only an IT project. One senior manager remarked, "We need to make it clear that we want managers at all levels to work in a more evidence-based way—that is how they should be doing their jobs."

ChemCo's first step was to assemble existing data analysts from across the company to form information support teams. Each team was assigned to one or two business units and charged with understanding their decisions and information needs in depth and then helping employees improve the way they accessed and used data. Initially, the teams shadowed employees on visits to customers and suppliers in order to learn what information was involved in customer-facing work, how it was used, where it was not available, and where it helped or hindered the accomplishment of a task such as negotiating a sale. Each team then held workshops with customer-facing employees to present what it had learned, offer ideas for delivering improved information, and get feedback.

On the basis of the workshops, the teams developed prototypes of various information reports and worked with the business units to try them out. Given that the brain finds it easier to process information if it is presented visually, the teams incorporated graphics, charts, and screen layouts in the prototypes. These experiments showed whether employees assimilated information, the behaviors they exhibited, and ultimately whether they won business. It was only at this stage—once the company had developed deep insight into how employees used information—that a CRM system was rolled out across the organization.

ChemCo did not customize its system any more than other firms typically do. But compared with most companies, it had a much clearer understanding of what information would be collected and maintained and how it would be applied. And because salespeople were involved from the outset, they strongly accepted the need to work in an evidence-based way.

As sales and service employees began using the new information more effectively, managers considered how to modify customer databases to support them. Over time, the CEO encouraged greater standardization of sales and information-use practices across business units and prepared the way for developing shared views and understanding of customers. His

mantra: "Do we *think* this is true, or do we *know* this is true?" The units identified what they did not know about their customers and the sales practices that could lead to poor customer interactions and lost business. As the company improved its interactions with customers, revenue grew; this, in turn, increased the appetite of salespeople for higher-quality customer and sales information that they could use to improve their performance—a virtuous cycle.

2. Emphasize Information Use as the Way to Unlock Value from IT

Initiatives designed to extract information from existing systems or new sources of data must acknowledge how messy—and complex—that process is. People don't think in a vacuum; they make sense of situations on the basis of their own knowledge, mental models, and experiences. They also use information in different ways, depending on the context. An organization's culture, for instance, can frame how people make decisions, collaborate, and share knowledge. Moreover, people use information dynamically and iteratively. The steps of sensing a potential problem or opportunity, deciding what information is needed, and then gathering, organizing, and interpreting it occur in cycles.

The conventional IT-development approach ignores those realities. The design of most IT systems takes into account the data that have been identified as important and controllable. Abstracting from real-world complexity in this way and creating formal, logical rules for processing data simplify system design and provide clearly defined deliverables. That approach is fine for activities that are highly structured and whose tasks can be precisely described, such as processing customer orders. It is ideal for moving information from the human domain into the technology domain so that organizations can exploit the phenomenal processing capacity of computers and remove human involvement wherever possible.

The problem is that many organizations mistakenly apply this design philosophy to the task of getting data *out* of the technology domain and *into* the human domain so that it can be turned into usable information—and in those instances the approach usually fails. In the case of managers, that's because their roles are often complex and have little structure. Even when an organization tries to capture their information needs, it can take only a snapshot, which in

no way reflects the messiness of their jobs. At one moment a manager will need data to support a specific, bounded decision; at another he'll be looking for patterns that suggest new business opportunities or reveal problems. He must be able to build both kinds of knowledge.

pproach. The same

engine manufacturer cannot expect diagnostic software alone to determine the causes of problems using the massive amount of engine-performance data the firm generates. Rather, the engineer must have considerable expertise and knowledge to identify relationships in and ask questions about the data, often through the testing of hypotheses. And in interpreting the results of any analysis, he or she must draw on experience to weed out misleading or false explanations. One engineer told us that he had 30 years of experience in vibration analysis but was still learning how to sift through and interpret data.

> # IT projects don't usually encourage people to look for new ways to solve old problems.

Analytics projects succeed by challenging and improving the way information is used, questions are answered, and decisions are made. Here are some ways to do this:

Ask second-order questions.

Instead of setting out to create a system that can help sales professionals easily answer the question, "What stock should we place on shelves today?" an initiative might begin by asking, "Is there a better way to decide how we replenish stock?" By posing second-order questions—that is, questions about questions—the project assumes that decision makers could improve the way they operate.

Discover what data you do and do not have.

Avoid being bounded by easily accessible data and systems, which are based on particular assumptions and logic about how the business should be run. While they may have been correct in the past, those systems most likely have not kept up with a continually evolving

business and competitive environment. And chances are that the mountains of data trapped in departmental silos such as R&D, engineering, sales, and service operations are not being exploited. At many financial institutions, for instance, the diverse lines of business don't share data, which prevents companies from forming a coherent view of individual customers and from understanding portfolios of customers relative to market trends.

Give IT project teams the freedom to reframe business problems.
An openness to looking at problems through a new lens led central bankers in countries such as the United Kingdom and Israel to find a strong correlation between broad economic trends and consumers' Google searches for washing machines, aerobics classes, cars, and other luxury items. The idea to look for this relationship began as a hunch at Google's headquarters, where a staff economist started exploring whether particular keywords could foreshadow the findings of traditional economic reports. The resulting paper was circulated among the central bank economists, spurring their interest.

We have observed that IT projects don't usually encourage people to look for new ways to solve old problems. This lack of creativity is often driven by a myopic view of data and their value to the business. To combat this attitude, some organizations have adopted techniques such as brainstorming and assumption surfacing and testing. We are increasingly seeing online discovery forums, where employees throughout a company are invited to contribute ideas about markets to be served, new customer trends, and new ways to exploit this knowledge.

3. Equip IT Project Teams With Cognitive and Behavioral Scientists

Most IT professionals have engineering, computer science, and math backgrounds. Not surprisingly, they are generally very logical and are strong process thinkers, and they tend to focus less on the "I" and more on the "T" in IT. For tasks such as processing financial trades or retail transactions, these are ideal skills. If, however, the goal is to support the discovery of knowledge, they become a hindrance.

To address this problem, many companies have added people with deep knowledge of the business to IT project teams, exposed IT professionals to complex business issues, and hired more data scientists. But those moves will not be enough. When working with big data sets, you can probably find statistically meaningful relationships between any variables you choose. What pulls you back to reality is knowledge of the business. The dilemma is that this knowledge can also limit your sphere of thinking.

For that reason, big data and other analytics projects require people versed in the cognitive and behavioral sciences, who understand how people perceive problems, use information, and analyze data in developing solutions, ideas, and knowledge. This shift mirrors the shift in economics to behavioral economics, which applies knowledge from the fields of social psychology and the cognitive and behavioral sciences to develop a new understanding of how people think and behave in markets and economies.

In some organizations today, big data and analytics projects already include people with backgrounds in those fields. Her Majesty's Revenue and Customs (HMRC), the British tax agency, has recently employed organizational psychologists, who help analytics teams improve their interpretive abilities by, for example, making them aware of their confirmatory biases: their tendencies to search for or interpret information in a way that confirms preconceptions. One such bias was that certain debt-collection approaches worked for particular categories of taxpayers.

HMRC's leaders recognize that in addition to knowing how the business works—for example, what kind of case can go to court, what that process entails, and why certain cases fail—data scientists also need to understand the mind-sets of debt collectors and the behaviors of debtors (for example, why some people who owe taxes pay before a case gets to court and others don't). The organizational psychologists assist in this. They also spend time in the field with inspectors (who conduct tax investigations) and call-center staffers (who negotiate with taxpayers).

Organizations that want employees to be more data oriented in their thinking and decision making must train them to know when to draw on data and how to frame questions, build hypotheses, conduct experiments, and interpret results. Most business schools do not currently teach this. That should change.

4. Focus on Learning

Big data and other analytics projects are more akin to scientific research and clinical trials than to IT initiatives. They typically start with sensing problems or potential opportunities, which may initially just be somebody's hunch. They then often move on to develop theories about the existence of a particular outcome or effect, generate hypotheses, identify relevant data, and conduct experiments. In short, they are opportunities for discovery.

The cycle of sensing, analyzing, and discovering can be repeated many times. Consequently, projects can last from a few hours to more than six months, depending on the complexity of the business issues, the availability and quality of external and internal data, the nature of the experiments, and the analytical techniques and tools that are employed. But the evolutionary, cyclical structure and relatively short duration make the costs of these projects much easier to control than those of traditional IT projects.

Traditional IT Project	Analytics or Big Data Project
Typical Projects	
Install an ERP system Automate a claims-handling process Optimize supply chain performance	Develop a new, shared understanding of customers' needs and behaviors Predict future growth markets
Typical Overarching Goals	
Improve efficiency Lower costs Increase productivity	Change how employees think about and use data Challenge the assumptions and biases employees bring to decision making Use new insights to serve customers better, build new businesses, and predict outcomes
Project Structure	
TRADITIONAL PROJECT MANAGEMENT: Define desired outcomes Redesign work processes Specify technology needs Develop detailed plans to deploy IT, manage organizational change, and train users Implement plans	DISCOVERY-DRIVEN PROJECT MANAGEMENT: Develop theories Build hypotheses Identify relevant data Conduct experiments Refine hypotheses in response to findings Repeat the process
Competencies Required	
IT professionals with engineering, computer science, and math backgrounds People who know the business	In some cases, IT professionals with engineering, computer science, and math backgrounds People who know the business Data scientists Cognitive and behavioral scientists
What Does Success Look Like?	
Project comes in on time, to plan, and within budget Project achieves the desired process change	Employees base decisions on data and evidence Employees use data to generate new insights in new contexts

Organizations can do several things to make learning a central focus of big data and analytics projects:

Promote and facilitate a culture of information sharing.

Most learning in organizations takes place in teams and in interactions among colleagues. It is therefore crucial to foster a collaborative culture in which transparency, trust, and sharing motivate managers and data scientists to contribute their best ideas and knowledge. An environment where information is not freely shared and failures and errors are hidden has no place in initiatives to generate knowledge.

One financial services company we studied has established a "data lab" to bring together managers from different functions as well as data scientists to work on specific problems in a discovery and learning environment free from the normal pressures of daily work. This allows new interpretations of data and business ideas to emerge from candid conversations across disciplines.

Expose your assumptions, biases, and blind spots.
Be willing to reframe the why, what, and how of your accepted business practices. Develop and test hypotheses to explore the limits of what you know and don't know.

Strive to demonstrate cause and effect.
Analytics is all about discovering relationships and meaningful patterns in data, such as the factors that seem to cause or are related to certain outcomes. It is therefore important to move beyond symptoms and instead address questions such as, What is the problem we are trying to solve? What are its root causes? What factors seem to contribute to particular outcomes? What can we do differently?

For example, HMRC receives around 300,000 paper returns annually on bequeathed estates, approximately 200,000 of which claim to be below the threshold for paying inheritance tax. Because of the large number of returns, the agency had trouble identifying the cases where more tax was due than was declared. So it sought to uncover relationships among data that would help spot such returns. Working backward from returns previously flagged as inaccurate, HMRC employees built theories about factors that could indicate underdeclaring. From those theories, they constructed hypotheses and tested them using historical data. After much iteration, the agency found that a combination of data on property ownership and

transactions, company ownership, loans, bank accounts, employment history, and tax records drawn from diverse public and private sources was effective in identifying potentially fraudulent returns. From those data the agency built a model—which it continues to fine-tune —to predict which estates would have tax liabilities. Estates reporting no liabilities but with particular profiles are now the object of further scrutiny. The result of this effort has been a significant increase in tax revenues.

Identify the appropriate techniques and tools.
Data scientists and analysts have their own favorite techniques and data sources. Managers must understand the strengths and weaknesses of those as they decide how to handle the deluge of newly available data.

One example of an industry confronting that challenge is pharmaceuticals, which is in the early stages of figuring out how to use monitoring technologies to lower the cost and improve the quality of drug trials. Achieving the U.S. Food and Drug Administration's approval for a drug can cost nearly $1 billion and involves conducting trials with hundreds, if not thousands, of patients. In the past, doctors monitored trial participants mainly by seeing them periodically in their offices. Today technologies, such as sensors that can be placed on a patient's body, offer the potential to monitor participants around the clock and capture real-time data about their adherence to the treatment regimen and the positive and negative effects of the drug.

> Managers should expect to get their hands dirty during the iterative process of generating business insight.

The challenge for pharmaceutical companies is to come up with ways of analyzing all this information and figuring out what's truly useful and what's just noise. This will require them to develop models and simulations that generate reliable and scientifically valid results of drug effectiveness that the FDA will accept.

Analytical techniques and controlled experiments are tools for thinking. But it is people who do the actual thinking and learning, so managers should expect to get their hands dirty during the iterative process of generating business insight. While there may be some "aha" moments, when ideas and insights emerge quickly, there will be many more occasions when managers—and not just data specialists and analysts—must rethink the problem, challenge the data, and put aside their expectations.

5. Worry More About Solving Business Problems than About Deploying Technology

Conventional IT project management is risk-averse. It concentrates almost exclusively on neutralizing threats to the successful delivery of a new system. In contrast, projects concerned with information use and big data should focus less on managing the risks of deploying technology and more on solving business problems—or, to put it another way, these projects should seek to avoid the risk of not achieving successful business outcomes. This focus is reasonable because, as we've noted, analytics projects are not nearly as large or as expensive as the deployment of an ERP or a CRM system.

Consider a European electrical goods retailer we studied, which wanted to give iPads to sales assistants in all its stores. The prime objective was to provide product information that would be useful in the sales process. The tablets would also help store managers and salespeople with merchandising and layout and would update them on promotions and marketing activities. One of the key business problems the project tackled was that in-store sales pitches were not always successful.

To assess ideas for improving interactions between salespeople and customers, the retailer conducted controlled experiments in which salespeople used various information layouts about products and presentation styles to communicate with customers. Initially, this experimentation delayed the deployment of the iPads for use in the stores, raising the risk that the project would miss its deadline and exceed its budget—which might have resulted in a reduction in the project's scope if the retailer had been applying the conventional IT project-management approach. However, learning which information layouts worked helped reduce

the business risk of lost sales. In our research we have regarded the connections between people and events as the main drivers of information use. People use information to understand social interaction (for instance, how sales conversations work with various customer segments) and to understand relationships (for instance, how customer segments respond to particular forms of product layouts). It is possible to provide relevant information only if this interconnectedness is understood. While deploying tablets appropriately in stores is an important concern, the goal of improving how sales managers and staff use information to drive sales should be the project's primary focus.Organizations have long looked to IT to bring data under control by automating transactions, streamlining information flows, and storing data for later recall. Conventional approaches to deploying IT work well in achieving this. The paradox is that the technologies that were supposed to help manage data are now causing a massive deluge. As organizations seek to exploit internal and external data, they run the risk of applying conventional methods when instead they need a fundamentally different approach and mind-set.

Improving how businesses extract value from data requires more than analytical tools. It involves creating an environment where people can use the company's data and their own knowledge to improve the firm's operational and strategic performance. In this new paradigm, the manager's priority is to make discoveries that could benefit the organization and identify unknowns that could put it at risk.

Donald A. Marchand is a professor of strategy execution and information management at IMD in Lausanne, Switzerland.

Joe Peppard is a professor at the European School of Management and Technology in Berlin.

This article is about IT

Related Topics: ANALYTICS | DATA | DECISION MAKING
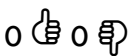
Comments

Leave a Comment

POST

1 COMMENTS

Chris Bergh   3 months ago

The article brings up a good point, that analytics is a continually flowing river of questions. The appropriate way to handle this is in a hypothesis driven approach. From a expense perspective, most resources should be spend on testing current hypothesis. But once a hypothesis is proved valuable, there is need to institutionalized and share it in metrics, models and monitoring reports.. So IT should deploy most of their resources on helping analysts answer the new hypothesis, while find quick ways to lower costs for those that hypothesis that are already answered. Only in this transition from high cost analytic activities in support analysts questions to low cost monitoring and reporting can enough resources be freed to maintain the overall process.

REPLY   0 👍 0 👎

⌄ **JOIN THE CONVERSATION**

POSTING GUIDELINES

We hope the conversations that take place on HBR.org will be energetic, constructive, and thought-provoking. To comment, readers must sign in or register. And to ensure the quality of the discussion, our moderating team will review all comments and may edit them for clarity, length, and relevance. Comments that are overly promotional, mean-spirited, or off-topic may be deleted per the moderators' judgment. All postings become the property of Harvard Business Publishing.