Harvard Business Review

Technology And Analytics

Analytics for Marketers

by Fabrizio Fantini and Das Narayandas

From the Magazine (May–June 2023)

Andrew Strauss

Summary. Advanced analytics can help companies solve a host of management problems, including those related to marketing, sales, and supply-chain operations, which can lead to a sustainable competitive advantage. But as more data becomes available and advanced... **more**

Advanced analytics can help companies solve a host of management problems, including those related to marketing, sales, and supply-chain operations, which can lead to a sustainable competitive advantage. For example, firms can integrate decisions and optimize the entire value chain by modeling individual customers' behaviors and preferences and offering tailored products priced as close as possible to shoppers' willingness-to-pay price points—all while reducing the cost of servicing individual transactions.

But as more data becomes available and advanced analytics are further refined, managers may struggle with when, where, and how much to incorporate machines into their business analytics, and to what extent they should bring their own judgment to bear when making data-driven decisions. The questions they need to answer are: When does it make sense to shift from traditional human-centered methods to greater automation of analytics and decision-making? And how can we strike an appropriate balance between the two?

One of us (Fabrizio) founded a practice that helps clients optimize performance using AI to automate pricing and supply-chain decisions; the other (Das) is an academic who has developed an MBA course that incorporates field cases focused on using AI to enhance marketing, sales, and support functions. Together we set out to understand how to maximize the potential of both humans and machines to arrive at the best business decisions.

In general, humans are more capable in the areas of intuition and ambiguity resolution; machines are far superior at deduction, granularity, and scalability. How can you find the right balance? There are three common approaches to analytics: *descriptive,* where decisions are made mainly by humans; *predictive,* where machines determine likely outcomes but humans choose which course to follow; and *prescriptive,* which usually means autonomous management by machines. This article describes when and how to use each approach and examines the trade-offs and limitations. (Although the focus here is on marketing and sales, the principles may be applied more broadly.)

Three Approaches to Analytics

The role of machines differs significantly in these approaches—from a tool to help managers understand a business situation, to an aid that supports managers' decisions, to a decision-maker that relieves managers of that duty. Let's explore each.

Descriptive: Aggregated Observations

In descriptive analytics—commonly termed "business intelligence"—managers use machines to make sense of patterns in historical data. They essentially ask, "Help me understand what happened." That help commonly takes the form of dashboards that highlight the input and output performance variables, enabling managers to decide "which dial to turn" and "by how much" on the basis of historically observed facts.

Descriptive analytics is about making sense of the past to inform the future. Past data is specific, clear, and certain, and this approach is rooted in verifiable and objective facts. We expect that descriptive analytics will remain part of business managers' daily experience. But because humans can't process enormous amounts of granular data, they must rely on highly aggregated information. Decisions based on that data tend to be coarse in nature, and they require the nontrivial step of extrapolating past trends and projecting them into the future.

Furthermore, descriptive analytics tends to be overly reliant on internal transaction data, which is the lowest-cost, most readily available data. External data, such as customer-related data (a Net Promoter Score, for example) and market survey data, are more expensive and time-consuming to source; they are also difficult to analyze and synthesize in real time. Consequently, the most common types of data used in descriptive analytics are internal and industry-performance variables, which are historically observed facts. Somewhat instinctively, managers complement backward-looking data with their own experience or received wisdom, especially when using this approach for diagnostics. Therefore, a descriptive analytics approach is heavily dependent on the intuition of specific decision-makers and on their ability to overcome their biases, such as by not cherry-picking data that validates preexisting views.

Three Approaches to Analytics

Different management problems are best solved by different analytics approaches. As decisions require less intuition and ambiguity resolution, and more deduction, granularity, and scalability, data and algorithms play a bigger role.

APPROACH	**Descriptive: Business intelligence** *What happened?*	**Predictive: Prediction engines** *What will happen?*	**Prescriptive: Decision automation** *What should I do now?*
ROLE OF MACHINE	Helps me understand	Supports my decision	Tells me what to do
SIZE OF VALUE-CREATION OPPORTUNITY			
EXAMPLES	Strategic planning — Initial product pricing — Scenario planning — Investor reporting	Demand planning — Discount/ promotion management — CRM segmentation — Maintenance	Inventory optimization — Price optimization — Markdown optimization — Risk optimization
RATIONALE	Typically little data available compared with the problem — High levels of uncertainty — Simplified manual approach	Quick-win opportunities — Relatively frequent decisions and observations — Semi-automation	Larger size of improvement prize — High frequency of decisions — Full automation

HBR

See more HBR charts in Data & Visuals >

In short, the descriptive analytics approach tends to lack external perspective and to be limited to high levels of aggregation. Managers provided with business-intelligence tools rely on past experience and high-level pattern recognition to project the past into the future, often relying on their gut. That can lead to repeating time-trusted approaches to solving problems rather than finding innovative new paths. Despite the subjectivity issues associated with this approach, it is still widely used because it's relatively simple and inexpensive to develop and implement. And it relies on humans for sensemaking, which puts it squarely in the comfort zone of most managers reared in the analog world.

Predictive: Limited View of the Future

With predictive analytics, machines determine the likely outcome or outcomes of a particular situation for different combinations of input variables, giving managers insight to choose the course of action whose expected result best meets their objective. Predictive analytics can be used to forecast wins and losses, calculate price elasticities, predict the impact of marketing actions on specific customers, and dynamically cluster customers in market segments. These predictions allow managers to drill down and make decisions at the transactional and tactical levels as opposed to the typically high level of descriptive analytics.

The predictive analytics approach is structurally limited. It's nearly impossible to predict future demand (let alone the future itself) with much certainty. Furthermore, even predicting individual input variables can be highly complicated: Weather, competition, and supplier performance, for example, may require their own prediction models. Such models can be not only difficult to build but also problematic because the inputs and outputs often depend on one another, forcing managers to predict input and output variables concurrently.

There are also limits to the number of input variables that can be modeled and the level of granularity that can be achieved. Although multiple factors typically influence purchase decisions, common predictive techniques such as regression, clustering, and time-series forecasting usually accommodate only a small subset of variables. That is because for a model to be valid, its variables must be independent of one another—but adding more input variables creates complex interdependencies that render the model statistically unfit. In addition, to make more-granular predictions, firms must collect more-granular data. For example, to predict sales of a specific product, they must collect data at the SKU level rather than the category level.

Well-designed prescriptive models can deliver greater financial rewards and better business performance. But they can be very expensive and complex to set up.

Another issue in predictive analytics is the burgeoning gap between data scientists and business scientists in terms of objectives. Data scientists are focused on improving statistical rigor, while business scientists are focused on optimizing the analytics to enhance business outcomes. For data scientists, the goal of predictive analytics might be to increase the accuracy of their model, whereas for business scientists the goal is business impact. Business scientists focus on maximizing the benefits of predictive analytics by accounting for the economic impact of a false positive (when the prediction is positive but the outcome turns out to be negative) or a false negative (when the prediction is a negative outcome and the firm decides against taking any action but would have achieved a positive outcome had it pursued the opportunity). For example, in a win/loss prediction-analytics exercise, a false positive typically results in wasted sales and marketing effort, while a false negative typically results in a wasted opportunity or lost business. Focusing only on increasing accuracy might result in a model that reduces false positives (a good outcome) but also has a high degree of false negatives, which would lead to wasted opportunities and suboptimal overall performance.

In short, predictive analytics can be problematic. Relying only on machines may lead to suboptimal business decisions and a loss of profit potential. Managers can, of course, perform manual diagnostics and predictive analyses on top of descriptive data to enhance the quality of decision-making. But that sort of ad hoc effort is subject to the same kind of biases as those observed with descriptive analytics.

Prescriptive: Granular Guidance

With prescriptive analytics, machines make decisions that are based on managers' defined objectives, by employing large amounts of data to rapidly analyze market conditions and learn by designing and running large numbers of low-cost experiments and what-if scenarios. Although many of their experiments might initially be suboptimal or even downright wrong, the machines can learn rapidly, getting closer to the optimal outcome targets quickly and inexpensively. They then tell the manager what needs to be done, shifting focus from inputs (such as ensuring the accuracy of decision variables) to outputs (such as optimizing the business impact of decisions), while explicitly modeling risk and economic costs.

The optimal prescriptive decision typically depends on market prediction, which drives the expected revenues, and on uncertainty, which drives the expected costs. In predictive analytics the focus would be on forecasting the number of units expected to be sold while ignoring the level of error in demand uncertainty. The prescriptive approach takes this uncertainty into account to make profit-optimizing decisions and continually adjusts as new information becomes available. For example, a retailer with low inventory on the shelves and relatively low logistics costs might respond to the possibility of a demand uptick with an aggressive inventory-replenishment strategy. However, the same retailer, in the face of high logistics costs and market uncertainty, might find a more conservative replenishment strategy to be optimal and profit-maximizing.

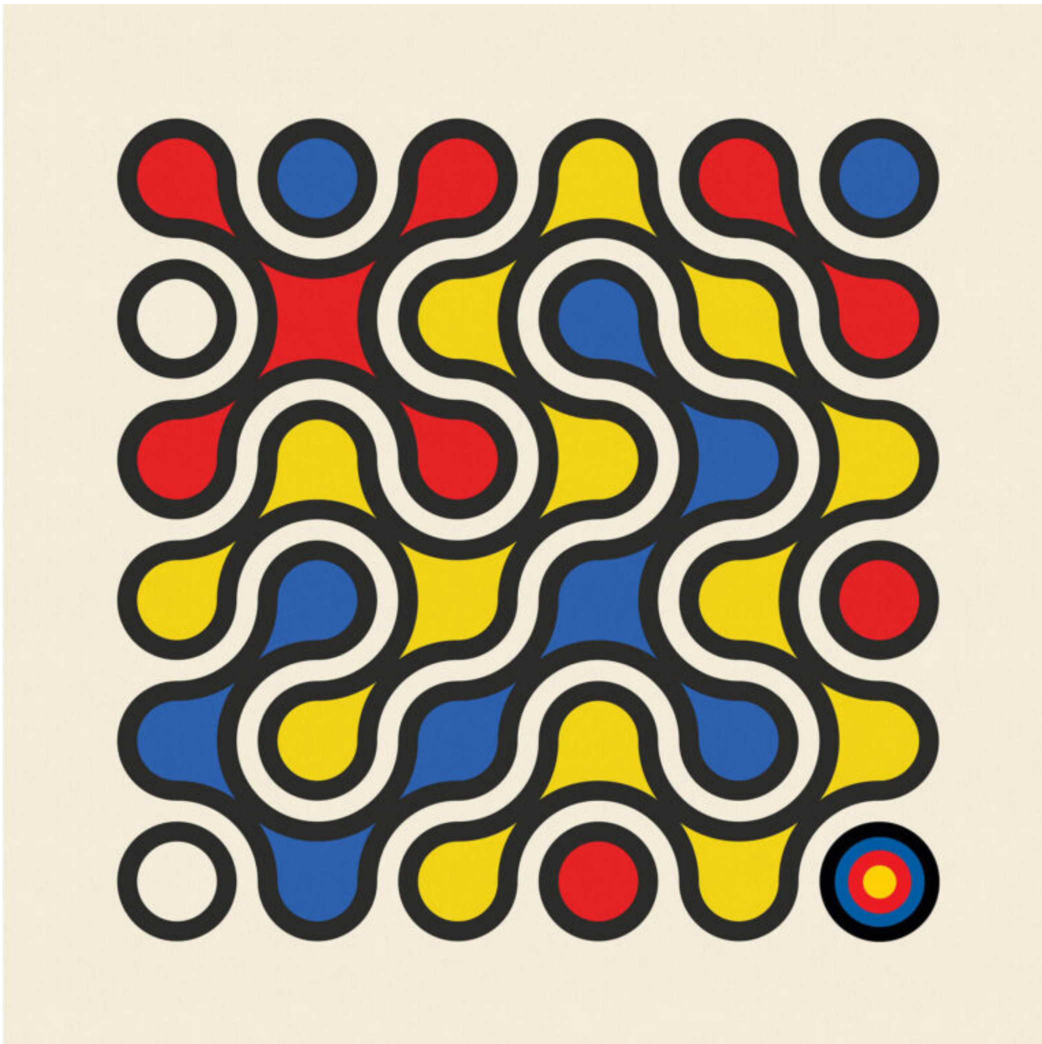

The visual artist Andrew Strauss works with algorithms and code to reimagine abstraction, repetition, and patterns through generative automation.

Well-designed prescriptive models can deliver greater financial rewards and better business performance than descriptive or predictive models can. However, they can be very expensive and complex to set up: They require dedicated software and hardware solutions and specialized human expertise to translate management strategies into mathematical, machine-friendly optimization objectives and business rules.

The human role in all this—defining the business rules and objectives—is tremendously important. Predictive analytics depends on the ability to translate business objectives, rules, and constraints into unambiguous directions to the prescriptive machine. That, in turn, enables the prescriptive model to dynamically calibrate its own recommendations toward the direction that management has specified while guaranteeing optimal outcomes and the systematic fulfillment of all rules and constraints.

When to Use Which Approach

Moving beyond descriptive analytics to more-advanced and costly approaches requires a cost/benefit assessment. Whereas costs are related to the infrastructure, expertise, and leadership required to collect and analyze data, the benefits depend on the opportunity for extra profits that can be captured through more-granular and relevant decisions.

Therefore, which approach to use in a given situation depends on two factors: the relevance of the available data and the strength of the business case. A successful balance between human and machine maximizes the contribution of each.

Data: When available data is limited and high levels of uncertainty exist, descriptive analytics is the most viable option for providing directional guidance to managers. As the frequency of decision-making increases, more granular data becomes available, and the relevance of the data to the problem increases, more-autonomous prescriptive analytics approaches tend to perform best. In intermediate cases, where only limited relevant data is available, a predictive analytics approach is preferred.

Business case: The profit-improvement potential derives from the amount of inefficiency that data-driven insights can be expected to address. But inefficiency isn't a characteristic of every business problem. And when it is an issue it may be addressable only with data that is not readily available. Therefore, not all problems are amenable to advanced approaches.

> **When choosing an analytics approach, we must rethink the role of the manager: from the person who has all the answers to the one who asks the right questions.**

For example, machines may struggle with problems related to setting long-term strategy and innovation, for which the initial definition of the question is actually more important than the formulation of accurate answers. But when it comes to the optimization of prices, inventories, or marketing investments, analytics offers companies substantial opportunities because accurate answers will better serve their customers' needs. For business problems with long time horizons, like planning, or high levels of intrinsic noise at the granular level, like CRM segmentation, or low marginal benefit from extreme optimization, like operations maintenance, a predictive approach tends to work best.

In a cost/benefit analysis, descriptive analytics is a "low pain/low gain" approach. It is most relevant in cases where limited data is available and a high level of uncertainty surrounds the outcome. While the absolute economic impact of each decision may be very high, the resulting improvement in performance does not justify the investments needed to incorporate machine input to enhance the quality of the predictions and decisions. At the other end of the spectrum, when a lot of data is available and there is an opportunity to enhance the economic impact in each single prediction with a high level of certainty, then prescriptive analytics makes the most sense, justifying its relatively higher degree of complexity and cost with its high return on investment. Often in these situations the absolute economic impact of individual decisions is not high, but the number of decisions being made, the upside potential in each of the decisions, and the higher levels of certainty of the outcomes over time combine to make the investment in prescriptive analytics worthwhile. Predictive analytics is the best fit in the intermediate region.

In Practice: The Evolution of Price Markdowns at Event Network

Excess inventory is a common problem. It must be sold, and usually at a discount, making price markdowns a pervasive and necessary part of inventory management. The root cause is the structural impossibility, even with a theoretically perfect forecasting model, of precisely predicting sales. Given the uncertainty of factors such as weather, competitors' actions, and macroeconomic shocks, managers tend to maintain high levels of inventory to avoid losing sales and customers.

Let's look at how Event Network (EN), which operates gift and memorabilia stores throughout the United States and Canada, handled the challenge. (Disclosure: EN is a client of Fabrizio's company, Evo Pricing.) Customer traffic at its stores, which are located in museums, zoos, aquariums, and other cultural attractions, is highly seasonal and relatively unpredictable. Each EN location carries unique inventory, often customized to the location (San Francisco or New York, for instance), the theme of the attraction (plants at a botanical garden), and the time of year (sweaters in winter). The chain's high number of SKUs—more than 100,000—posed a formidable challenge to price-markdown management.

Over time EN has used all three analytics approaches. Here's how each one worked.

Approach #1: Descriptive Analytics

EN managers started by using a simple method: They offered deeper discounts on products with higher inventories that resulted from disappointing sales. To decide which products to mark down and by how much, EN managers considered measures such as historical sales per week, inventory levels, and coverage ratio (the number of days that the inventory will last given the current rate of sales).

To calculate the markdown for a product with a $10 unit cost and 10,000 units on hand, they multiplied the proposed markdown (30%) by the number of units on hand (30% × 10 × 10,000). They started with the SKU with the highest coverage ratio and worked down the list of SKUs until the total available markdown budget was spent.

This approach was ultimately unsatisfactory because it relied entirely on historical internal inventory-performance data. It did not consider customer- or context-related factors that have a significant impact on consumer demand.

Approach #2: Predictive Analytics

Next, the managers used regression-based techniques to discount products with the highest price elasticity (the percentage change in sales volume expected from a given percentage change in price). They calculated price elasticity by running the regression of historical sales volumes on historical prices by category by store by week. For example, a price reduction of 10% for an SKU with a price elasticity of –2 yields a volume of sales increase of 20% (a product of –10% × –2). So going from a baseline of 100 units at $10 each earning $1,000 in revenue to selling 120 units at $9 each would lead to $1,080 in revenue, representing a gain of 8% in revenue. Similar calculations can be made for metrics such as margin and inventory level. By simulating scenarios, the managers could pick their preferred strategic objective and determine the optimal markdown mix according to its expected impact. Doing so could take into account not just internal inventory data but also the customer-demand expectation and therefore the market impact of their decisions.

The optimal markdown varied according to the managers' objective rather than sales or inventory level. Although the results of their regression models were statistically significant, the EN managers found the explanatory power of the models to be relatively low (price explained just 10% to 20% of the variance in the sales of a product). That's because many other factors than price influence sales, including weather, foot traffic, and the range of products available. Adding such variables to the model would have incurred the cost of collecting the additional data in a timely manner. Moreover, more data would increase the complexity of the calculations by introducing more noise and causing unwanted interdependencies among the variables.

The EN managers went ahead with the simple one-dimension regression of volume versus price, however crude, since it yielded results superior to those obtained using the descriptive analytics approach. The resulting improved performance also built up the EN management's appetite for the use of more advanced approaches to analytics. They became open to using a different approach altogether to overcome the structural limitations of the predictive analytics approach.

Approach #3: Prescriptive Analytics

The prescriptive analytics approach that the EN managers eventually used improved on the prior two approaches by accounting for the broadest range of factors affecting consumer behavior. Using multiple data sources and advanced techniques such as machine learning and automated optimization, EN could identify which products to discount at any particular time and by how much.

The managers recognized that it was virtually impossible to rely on intuition at this level of granularity and nonlinearity. Furthermore, their journey across the different analytic approaches led them to appreciate the benefits of using automation and machine learning to make sense of complexity and to build self-learning systems that improved profitability significantly over time.

. . .

When it comes to choosing an analytics approach, it is imperative to rethink the role of the manager: from the person who has all the answers to the one who asks the right questions. The framing of problems, which can then be given to machines to solve, remains squarely a human ability. But managers can wisely cede some control to machines. The primary considerations when choosing the best approach are known and clear: the relevance and availability of data, and the potential for improvement in business impact expected from investing in more-sophisticated analytics.

Humans and machines excel at different tasks: humans at dealing with limited data and applying intuition in unfamiliar contexts, and machines at making decisions, however granular and sparse, that are repeated in time or space or both, and in environments flooded with rich data. Provided with too little data, in highly ambiguous situations, or in the presence of conflicting objectives that limit what can be inferred from data, machines struggle to produce relevant outcomes. But for complex problems that have abundant relevant data and whose solutions could significantly improve business performance, managers should buy or build the right machines and set the right goals for them to do what they can do so well.

A version of this article appeared in the May–June 2023 issue of *Harvard Business Review*.

Fabrizio Fantini is the founder and CEO of Evo Pricing, a firm offering AI solutions to automate and optimize pricing and supply-chain decisions.

DN

Das Narayandas is the Edsel Bryant Ford Professor of Business Administration at Harvard Business School.