

Megadeals: How data and analytics can dramatically boost success

Megadeals are tricky, so rare that companies think there's not enough accumulated data to be useful. But the data isn't as scarce as some companies think.

Tobias Frese, Francesco Iannaccone, Mikhail Savin, and Michael Viertel



DECEMBER 2018 • MARKETING & SALES

Megadeals can be a crowning glory for both sales reps and businesses. A megadeal can propel a sales rep past her quota and, often, up the corporate ladder. For a company, it can make the difference between meeting the annual sales target and issuing an earnings warning. In many B2B industries, just 1 to 2 percent of deals can contribute more than 40 percent of annual revenues.

However, companies often fail to seize these golden opportunities. Their rarity and the unique characteristics of each deal make it hard either to apply lessons from previous successes or to create a standardized process for clinching subsequent deals. As a result, a new megadeal can trigger frantic and undisciplined efforts to win, deliver, and retain it that ultimately fall short.

Data can help, yet companies often hesitate to apply data and analytics to megadeals due to the perceived scarcity of information. Yet by drawing on multiple data sources—customer-relationship management (CRM), enterprise resource planning (ERP), sales reporting, and external data—it is possible to create small but rich datasets and derive insights that can significantly boost returns.

To understand more clearly how to convert megadeal opportunities into top-line results, companies need to invest to improve data quality, build capabilities in advanced analytics and machine learning, and know how to continuously generate and develop ideas on how to use data to make better decisions.

A powerful approach to boost megadeal success and maximize profitability is to understand how to apply data at each stage of the sales funnel: lead qualification, pre-RFP, negotiation, and delivery.

Lead qualification: Scarce data can still be rich data

When facing any business challenge, it is always tempting to mobilize the best people. When a

megadeal lead comes in, the first instinct can be to parachute top global sales talent into the local team. This can work, but it is not foolproof. Data can help find the most appropriate way of converting a lead.

A renewable-energy company urgently needed to close on a couple of megadeals to avoid an earnings warning. Sales managers immediately redirected a significant share of their sales support and expert teams to these opportunities. Instinctively, this made sense, and after three months, one deal was moving toward completion, thanks to some detailed conversations about the technical value proposition on the table. But the approach turned off decision makers at a deal that was being orchestrated by private-equity investors rather than by the usual industrial procurement team. The PE firms saw the technical conversations as a distraction from their focus on the deal structure.

The company decided to change tack. It combined CRM data on its customers from the past three years with external data (rich data on customer financials, industry classification, geography, etc.) to understand which key accounts actually benefited from investments in expert resources and which required other tactics, such as management visits, additional commercial terms, or a different go-to-market approach altogether.

By mapping the sales-expert deployment and investment to the outcomes, the company was able to allocate sales resources much more efficiently. For example, to management's surprise, the mapping revealed an extremely low win rate among one customer segment (farmer cooperatives in central Europe) that soaked up substantial investment in the expert

resources required to run product demos, Q&A sessions, and certifications. Meanwhile, the win rate for the PE investor segment bore no relation at all to the additional investment. Time to rethink. The company reallocated expensive resources away from both segments to opportunities where it knew they would have more impact, and senior management focused on building strong, centrally prepared business cases for the PE investors that did not require the sales experts' time.

Pre-RFP phase: Convince with customer data

Sales reps invest enormous effort in developing convincing arguments that help potential B2B customers understand why their company will deliver more value than their competitors will. These arguments are particularly important in megadeals, where a whole bench of competing firms may have to run through extensive pitches in front of the customer's procurement, technical, and executive teams. But what actually convinces a customer to go with sales team A over sales team B? The data is revealing.

A B2B technology-services company studied the success rate of its megadeal pitches and found that those based on real customer observations and real customer data had significantly better traction and win rates than those that relied on generic sales materials, especially when pitching to existing customers. Given the scale of individual opportunities on the market, the company decided to create a replicable approach that would bring the best information and customer-specific insights to each pitch.

Sales leaders took data from the software they used to manage operations with existing customers and combined it with multiple external data sources on the customer's market. Analysts then developed a sophisticated model that simulated the effect of changing various operational parameters and practices on production output.

The model generated a fact-based assessment of the improvement potential in customer operations that the company's service could bring. This could be translated into operational key performance indicators (KPIs) and—crucially—to customers' bottom line. The company quickly turned this tool into a digital application that each sales rep could take to the customer. This allowed the rep to shine in customer meetings, demonstrating a deep knowledge of the customer's equipment and even committing to the improvement potential by offering a lower price if the benefits did not materialize. The app immediately gave the company more traction with stakeholders and helped it shape the RFP together with the customer. The result was a one-to-two percentage-point increase in profitability for the relevant services.

Negotiation phase: Let machine learning weed out loopholes

Megadeal negotiations tend to be very complex, as both buyer and seller are discussing significant investments. This can lead to complex pricing structures and hundreds of pages of contractual terms. A single weak clause in this mass of documentation can result in major losses.

A large customer of an industrial-equipment company benefited from such a poorly worded clause. The contract failed to lock the customer into a commitment to buy a full range of equipment; instead it was able to cherry-pick only the lowest-priced items, which were loss making for the vendor. This inevitably meant the equipment company's margins were much lower than expected when the deal was signed.

It is extremely hard to spot such loopholes in a quotation approval or review process (they may even be deliberately hidden by the customer's procurement team). Going through the haystack to find the needle is time consuming and in some

cases not even feasible for a human, but it is easy for a machine equipped with the right analytics.

A telecom-equipment company built just such a tool for its megadeal pricing and contract-review process. Initially, it had a commercial war room staffed with 40 full-time pricing experts. The team logged the significant details of all its major sales, from megadeals down to midsize transactions. This generated more than a million data points from a sample of just a couple of thousand deals. For the complex megadeal bids, the team typically needed a full week to analyze and approve the terms—easily justified, given that the deals were worth more than €10 million. The team also delivered results, improving the profitability of megadeal bids by more than three percentage points.

To industrialize the approach, the company developed a machine-learning algorithm to analyze this data. Over several iterations, the algorithm was fine-tuned to achieve the same quality of recommendation as the war room on 75 percent of all deals, while passing the other 25 percent to the pricing experts for a final verdict. This not only allowed the experts to spend more time on new activities, it also reduced the time to review a megadeal from a week to less than a second, which let sales reps test different options and build strong negotiation strategies in real time. Building on its success, the company expanded the tool to identify overscoping risks within the offer and recommend adjustments to the sales rep, which helped make the quote leaner and improved win rates.

Delivery: Use algorithms to improve profitability

After the handshake, the sales team often moves on to the next big hunt. However, in megadeals, the final signature is only the start of the journey to realize the revenue and profitability and—equally important—deliver customer satisfaction. Success is far from guaranteed. Given their complexity,

megadeals often suffer from misjudgments at the tendering phase and lengthy delays. Large bridge or tunnel infrastructure projects, for example, overrun on average by 35 percent, while major IT implementation deals typically overrun by more than 50 percent.

To avoid such pitfalls, smart sales organizations must identify the risks early and act accordingly. One company often received enormous orders for multiple variations of its products as part of a megadeal. It struggled to deliver these orders on time due to production constraints, and could not keep items in stock due to the high cost and price depreciation of the equipment. Rather than continue to frustrate its customers and jeopardize its relationships, it decided to analyze its customers' demand patterns to find a solution. It found that a significant share of orders and exact equipment variants could be predicted by studying regulatory changes in customers' markets as well as the publicly available infrastructure-investment plans of all potential and existing customers.

The company created an extremely simple algorithm that was regularly fed by the market-intelligence department. This algorithm eliminated much of the uncertainty in forecasting the product-variation mix, which boosted the company's ability to meet large-customer purchase orders and streamlined production plans.

A chemicals company had similar challenges. It faced continuous operational issues with its production process, resulting in missed shipments to some of its biggest customers. To resolve this, it tended to be reactive, prioritizing customers that shouted the loudest. As a result it consistently failed to meet profitability targets as it fulfilled claims from other customers for delayed shipments. The company developed a sophisticated optimization model that accounted for contractual penalties,

production-network reconfiguration costs and options, spot-market purchase costs, transportation, and the sensitivity of customers to shipment times compared with alternative local options. The model helped the company fundamentally rethink its approach and to quantify the relevant trade-offs it faced. For example, in some cases it made sense to ship direct to smaller customers in competitive markets while using the spot market to fulfill selected megadeal contracts in order not to jeopardize the customer relationship.

How to get started

Unleashing analytics and machine learning on data can be tremendously lucrative for companies seeking, closing, and servicing megadeals. Given that each deal may be worth millions of dollars, it is usually worth making the investment in the technology and expertise required. However, many initiatives stall while debates roil on the collaboration model between commercial and IT,

on who owns analytics resources, or on technical analytics platform decisions.

It is rarely wise to go completely rogue and launch disconnected independent analytics efforts across the organization. However, building up a large central function can be slow and costly. The most successful companies have a federated analytics governance model that relies on a small center of excellence to set guidelines such as platform selection, to help build skills and to share best practices across functions. Each function then takes responsibility for specific use-case development and hosts small analytics teams that work closely with the business. The sales and commercial organization is often the perfect place to create such a center of excellence and prove the value of analytics to the organization. Once it has generated some wins, it can seed analytics teams into other functions. (See the sidebar, “One telecom’s analytics journey.”)

One telecom’s analytics journey

An interview with Vincenzo Pinto, head of commercial strategy and analytics at Nokia.

How did you embark on the analytics journey at Nokia?

When I joined commercial management, I quickly felt that analytics was a major opportunity for us. We work on megadeals that are of huge importance for our customers’ business, and therefore we are under strong pressure from customers to know them better and to use this information to improve how we serve them. Analytics and machine learning was a way to enable that. We started with very simple activities—

collecting the right data, ensuring its quality, and providing it to commercial stakeholders. Over time, we started to increase sophistication to descriptive analytics and later to machine learning.

What was your first analytics use case?

We were looking for activities that needed a lot of expertise and added a lot of value, but soaked up a lot of time. These are ideal situations, as machine learning is great at replicating expertise. The deal-pricing review process is conducted by

experts with five to ten years of sales experience and was perfect. We moved fast and set up a small group of data scientists and functional owners to model this supervised machine-learning case. We did not focus on the team's operating model, its mandate, or who would own the process in the future—this was simply something we wanted to try out. Indeed, for many colleagues this was more of a passion—a hobby they wanted to drive.

By moving quickly and getting a small team of data scientists and functional owners working on it, we were able to replicate the pricing recommendations of our experts with a 97 percent match. We were astonished—and of course, we cut the length of time the review took from several days to several seconds. This success case quickly propelled the reputation of our team and approach.

How did you create a portfolio of use cases?

This first success brought more success. If we had started from the top down, by seeking a mandate, creating a long-term road map and investing in technology, I don't think we would have gotten traction. It was critical to establish credibility with the pricing use case.

We then became evangelists, reaching out to different commercial and functional stakeholders. We would come with some ideas for their function, to stimulate discussion and show what was possible. The eventual use case was rarely our initial idea, but bringing something to start the discussion was crucial to making progress.

As we started to build credibility, we also started to take that longer-term view. I began building a road

map of where I wanted to see advanced analytics boosting Nokia's performance. I didn't use the road map to push any topics, but rather reviewed proposed use cases to see how they fit. It was like completing a jigsaw puzzle where you know what the end product looks like, but you can't control the order in which you pick up the pieces.

Q: How did you scale?

Our recent priorities have been to start standardizing ways of working. Once we had more than ten use cases being built simultaneously, we had to start structuring our operating model and making some choices. We set up a clear team structure consisting of squads that have clearly defined roles and profiles. Second, we defined a data strategy that includes our likely data needs over the next three to four years that we need to start collecting now. Finally, we developed a change-management process that brings the full organization up to speed on what is possible and how it can benefit them. This last point takes a lot of my time. It was very easy to get the first 10 to 20 percent of innovators in the organization onboard, but it gets harder to create buy-in and excitement as we start to cover the full organization.

Q: What is next?

The next step is to drive this in a big way beyond commercial use cases, while at the same time maintaining the business-driven approach where I act as an enabler for the business ambitions of my colleagues.



Megadeals are so valuable that it's tempting to throw the kitchen sink at winning them. But the reality is that—size apart—there is no single characteristic of a megadeal that dictates the right approach to finding, securing, and delivering it profitably. Although sales leaders have been reluctant to turn to analytics given the scarcity of megadeal data, the reality is that with some creative thinking and suitable rigor, machine learning can in fact have a substantial impact on the rarified world of the megadeal. ■

Tobias Frese is a consultant in McKinsey's Berlin office, **Francesco Iannaccone** is a consultant in the Milan office, **Mikhail Savin** is an associate partner in the Helsinki office, and **Michael Viertel** is a senior partner in the Munich office.

Copyright © 2018 McKinsey & Company.
All rights reserved.