CHAPTER 1

Reimagining the Enterprise

The Stained Glass Bistro in Evanston, Illinois, is a bustling wine bar with impressive wine flights and a one-of-a-kind cheese and charcuterie plate. Nonetheless, your coauthors weren't paying much attention to the ambience during our dinner there together on November 12, 2013. By then, we were immersed in an animated conversation—begun hours before in Mohan's office, just a short walk away at the Kellogg School of Management—about the startling evolution now underway for analytics in large enterprises.

Introduced by our mutual friend Mary Gros, we had come together as two veteran technology professionals with decades of experience under our collective belts, trying to chart our industry from two very different perspectives: Mohan as an academic researcher, consultant, and technology company board member; Oliver as a longtime analytics practitioner and executive at major companies like eBay, Sears, and now Teradata Corporation.

As the ideas and the wine continued to flow around the table, we came to realize how our different perspectives were like complementary puzzle pieces that when fit together suddenly painted a clear picture of the journey that a data-driven enterprise needs to undertake. We began to see that rapid developments in technology and the explosion of data are now transforming the very nature of large enterprises—and that maturing analytics capabilities are the key to future survival.

Our insights that cold Tuesday in November were as enduring as they were sudden. Even the calendar date, 11/12/13, seemed auspicious as we began our own step-by-step journey in developing a capability maturity model for large data-driven companies, and putting it together in the book you're now reading. As you will see in the pages that follow, we're on an evolution toward an end state—a journey every big company should take, but only a few brave ones have started—that we call the "Sentient Enterprise."

We've since filled in the details of the five stages—the Agile Data Platform, the Behavioral Data Platform, the Collaborative Ideation Platform, the Analytical Application Platform, and the Autonomous Decisioning Platform—but everything is based on the insights sparked during that first meeting at the Stained Glass Bistro. Even the initial "Sentient Enterprise" term that we came up with over our dinner has endured: As we talked through the need to make decisions in real time

and at the speed of data, with the enterprise ingesting information and using algorithms to make the bulk of decisions on its own, Mohan observed how such an enterprise was almost like an organism, a sentient organism . . . a "Sentient Enterprise."

The name stuck because it is persuasive not just as the title of a maturity model or a book, but because it summarizes the end state our whole analytics journey is leading us toward: the Sentient Enterprise is the North Star that every large business should aspire to as it struggles to make decisions at the speed of data.

DISRUPTION AND DECISION MAKING

Henry Ford is famous for reportedly (though not definitively) having said, "If I had asked people what they wanted, they would have said faster horses." By choosing instead to create the Model T automobile in 1908 and introduce the assembly line approach to production, he disrupted and redefined an entire mode of transportation. A century later, Steve Jobs fostered much the same attitude and outcome in developing the Apple iPhone, an instant and total game changer in how we view phones and mobile capabilities overall.

These are prime examples of disruptive innovation, a term coined only in 1995 but already a fixture in modern business theory and practice. Today, disregarding or deconstructing the status quo is embraced in countless business plans. Entrepreneurs are mounting wholesale reworkings of entire industries and product lines. Many investors put their money on the disruptive playbook, favoring the revolutionary over the incremental.

Serial disrupter Elon Musk is legendary for upending e-commerce with the 1998 advent of PayPal. Five years later, he echoed Henry Ford's makeover of the auto industry by founding Tesla Motors. From design and manufacturing to service and operation, Tesla's electric, software-intensive vehicles have redefined what the automobile can be today. (Full disclosure: both of us are proud Tesla owners and big fans of the company and its approach.) In a further instance of disruption, Tesla even released all of its patent holdings in 2014 in the belief that open-source innovation can accomplish more than any single company can achieve with its own proprietary ideas.

By definition, disruptive innovation displaces companies and sectors that remain vested in the status quo. There will always be losers. Borders Books and RadioShack, for example, are two retailers that failed to straddle the online/brick-and-mortar divide with seamless, multichannel customer engagement models. They learned the hard way how disruption and bankruptcy share the same linguistic root.

Modern analytics, however, raises the stakes and brings disruption to another scale altogether. As data-driven becomes the norm across all industries, we're no longer just facing obsolescence of particular products, sectors, or services; we're now seeing the extinction of fundamental business models that most major companies have been founded on.

Indeed, big data and the new analytic capabilities that go along with it are changing everything from how large enterprises structure and finance operations to how they pursue opportunities and engage their workforce. And analytics can revolutionize an organization's ability to listen to data sources, understand what the data is saying, and use it to make informed decisions in near real time.

SELF-DISRUPTION AT CISCO: ON PURPOSE AND AT SCALE

Foundational writings on disruption are required reading for anyone in business today. So embedded is the concept, in fact, that the question is no longer how and why disruption happens, but who the winners and losers are. That has given rise to a second wave of insight around the mantra to "disrupt or be disrupted" (a very large wave indeed, judging by the 14 million search results when we recently Googled the term).

Consider the case of Cisco Systems, a hugely successful global networking company that includes more than 70,000 employees and 240,000 industry partners. Some 80 percent of the world's networking traffic crosses Cisco infrastructure at some point in time; the company consistently ranks number one or two in every market where it competes. Still, Cisco is pursuing an aggressive and company-wide self-disruption effort as if its survival depends on it. That's because it does!

We recently caught up with Kevin Bandy, senior vice president and chief digital officer for Cisco Systems. As he shared Cisco's

self-imposed transformation from a hardware-intensive model to a software- and consumption-based model focused on recurring revenue, Kevin explained that the company is not racing so much against competitors, but against the future needs of its own customers.

"Business models are changing every 18 to 24 months with Moore's law," he told us. "Our trigger to change was the voice of our customers and the forward visibility they expect us to have when it comes to innovation and how they'll be consuming it in the future."

Marathon runners know to hydrate before they get thirsty. The same can be said of companies needing to self-disrupt before they get desperate. "Rather than let someone else disrupt us, we chose to disrupt ourselves," Kevin explained. "That's especially critical with the operational level we're at; 80 percent of global networking traffic is too important to let fail."

SELF-DISRUPT IN SUSTAINABLE WAYS

When you're a small start-up, disruption is like intellectual Red Bull that powers you through a few market cycles. You're agile because you're small. And you can risk a huge crash because, given how 90 percent of start-ups fail within a few years, long-term survival is mostly an abstraction.

Large enterprises, with an ecosystem of customers relying on them, can't afford to think this way—but neither can they afford to sit still. To thread the needle, big companies like Cisco are fostering disruption and entrepreneurialism within the context of sustainable and scalable models. You're building a digital operating model of people, processes, behaviors, and competencies in the spirit of what the Wharton School's Eric Clemons calls an "all-pervasive" approach to disruption across the "structure and strategy of the entire business."

Unless you're constantly anticipating tomorrow, even today's biggest successes will always be on borrowed time. This is especially true in analytics, where clients may be buying not just a product but an entire digital environment that their whole business relies on.

"Think about the logic of Moore's law, and the reality that corporate timetables for standing up innovation at scale can be 18 to 24 months," explained Kevin. "If it's only then that we realize we stood

up the wrong solution, we can spend another two years of unraveling and rebuilding. That whole cycle counts for eons on the technology clock—plenty of time to put yourself and your customers out of business if you make the wrong call."

Cisco's story shows how, especially for large enterprises that serve in a trusted adviser role, getting ahead of disruption is a make-or-break proposition. Whatever customers think of you now, they'll abandon you—or go out of business along with you—if you're not there with the right innovations needed for tomorrow.

"The further along you go on this Sentient Enterprise maturity model, you encounter the challenge of people relying on you," echoes Brett Vermette, director of big data infrastructure and platform engineering at General Motors. "Delight becomes demand. Experiment becomes expectation."

We interviewed Brett about GM's own proactive transformation to consolidate what turned out to be hundreds of disconnected data marts into a more unified and agile environment. "We had an intensive, six-week period in early 2013, launching our enterprise data warehouse program," he told us. "This was a major transformation program, including installation of 60 crates of infrastructure in our data center, building the foundation for a global data warehouse, and consolidation of more than 200 siloed repositories and data marts over time.

"As part of our IT transformation, GM hired thousands of new college graduates and experienced IT professionals to handle work previously done by third-party suppliers," Brett said. "We were fortunate to have top leadership support; the challenge was more how to mobilize company-wide in ways that remained agile and innovative. It's like fighting the agility war on all fronts." In Chapter 3, we'll take a closer look at how GM achieved success in this effort.

ANALYTIC PAIN POINTS AND A SELF-SERVICE REVOLUTION

GM is a manufacturing company with roots going back more than a century. The fact that such a legacy institution would embrace analytics shows just how fully data has penetrated all markets and industries. Indeed, financial implications for this global tsunami of digitized information are so important that the World Economic Forum has

now designated big data as a new kind of economic asset, just like currency or gold.

A study by the MIT Center for Digital Business, meanwhile, is representative of many in showing how data-driven businesses do indeed have the edge. That survey of 330 leading U.S. businesses showed companies that focused strongly on data-driven decision making had an average of four percentage points higher productivity and six percentage points higher profits overall.

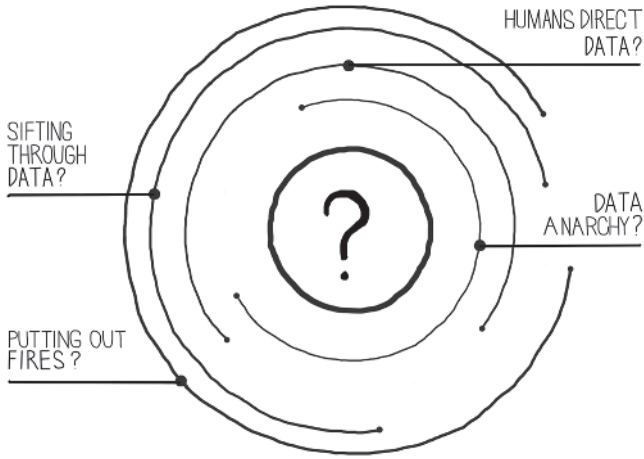

Many businesses nonetheless struggle through a combination of huge data volumes and organizational hurdles that creates analytic pain points familiar to any data scientist who works at scale. For one thing, we spend the majority of our time just sifting through data instead of making decisions. We're constantly on our heels in reaction mode, putting out fires instead of charting the future. We can't seem to make decisions fast enough, given that our brains don't scale the way data can.

Consensus among computational neuroscientists puts the human brain's storage capacity somewhere between 10 and 100 terabytes. Compare that to a worldwide data explosion that's more than a trillion gigabytes annually (1,000 gigabytes = 1 terabyte) and on track to hit 30 exabytes (or 30 billion gigabytes) every month by 2020—and you begin to see the root of nearly all our enterprise challenges and opportunities.

Success involves recognizing the changing nature of both data and the human expectations and interactions around that data as technology transforms the entire world into an on-demand, self-service community. A recent study by Forrester Research showed 73 percent of online customers consider time as their most important customer service priority, while another survey showed three-quarters of consumers prefer to solve customer service issues on their own.

Self-service has even come to dominate tasks involving personal or sensitive information, like accessing insurance records or checking in at a doctor's office. Among the other data points, a *Harvard Business Review* report found self-service on the rise in restaurants; and back in 2013, a large Cisco survey of consumers in 11 different countries found 61 percent were willing to shop in a completely automated store. Now the advent of Amazon Go has made that a viable and widespread option.

Self-service options have revolutionized the average person's ability not just to receive products and information, but to create them, too. A prime example is the do-it-yourself mobile app creation trend, with the advent of the software development kit (SDK) or devkit putting advanced software development skills at virtually anyone's fingertips.

In 2011, for example, a 14-year-old named Robert Nay used a free version of Corona SDK to create Bubble Ball, a mobile game for iPhone, iPad, and Android that was downloaded 400,000 times the first day of its release and went on to unseat Angry Birds from its No. 1 spot in the Apple Store. Robert programmed the game all by himself in just over a month's time. If this is how far we've come in the realm of shopping and games, you begin to see why nearly every employee working today within data-driven companies is clamoring for the same freedom around business intelligence and analytics.

ACCESS AND CONTROL

Imagine for a moment that you're a midlevel data scientist at a large corporation. Back home, your teenage kid can follow Robert Nay's example and use a free online devkit to manipulate everything from

audio, graphics, cryptography, and networking to device information like accelerometer, GPS, and user input data.

Unfortunately, while he runs along to school to show off his new smash-hit game to classmates, you run to the office and smash your fist against your IT department's door, demanding why it is going to take 18 months to get that data research project done for your finance department to make decisions about asset allocation!

You're a resourceful employee, so when you don't get a good answer or a better time line, you find a better way. You quickly build a data mart—a collection of copies from the company's broader data infrastructure—that you customize for your own specific business needs: in this case, finance. In short order, this data mart is churning out insights that feel useful and relevant. This is your ticket to agility in a demanding environment where time to market is everything.

The problem is that data that may seem useful turns out to be unruly; information that seems relevant could be just plain wrong. It seems your own little slice of the company's data is copying information and introducing data drift—the unpredictable and unending mutations of data caused by operations, maintenance, updating, and replatforming of databases. You've just created chaos in the guise of agility. This is one of many nightmare scenarios that are anything but hypothetical; they're the unfortunate reality for most large enterprises today.

These problems of control, access, and value around data are nothing new. Before our first meeting in Evanston, we had each given dozens of keynotes and lectures on various parts of the picture. Oliver routinely shared with business leaders the trials, tribulations, and opportunities that came with his having built production-scale analytic capabilities from the ground up. Mohan, meanwhile, had been making similar rounds with executive audiences from his vantage point as a busy researcher, consultant, and board member of a public company with more than 1,500 analytics professionals.

Presentations tend to be one-way conversations, however. So it was a revelation for the two of us to finally meet, compare notes, and hit upon a vision for the future. We realized how Oliver had been in the trenches and Mohan had been in the clouds. Long before the

term *big data* became popular, Oliver was toiling at ground zero, living through a larger analytic journey that he was too busy to recognize or name. Mohan, meanwhile, was looking at the same dynamics from 30,000 feet, charting broad industry trends without the "how to" insights that made up most of Oliver's daily life. At last, our combined perspectives connected the dots between things like vision, strategy, finance, and execution.

This expanded dashboard—our combined perspectives on our data-driven economy—brought new clarity to both the pitfalls and the payoffs involving huge growths in data volumes, complexity, and velocity. Today, well over 1.7 megabytes are generated per minute for every person on Earth. Much of this growth involves so-called behavioral data, which are all those events and data points in between or across transactions. These are the bread crumbs and fingerprints, derived from both human- and machine-generated data, which let us chart behavioral patterns and understand interactions. We no longer just count or aggregate data points; we are now connecting them in complex ways.

One of the side effects of dealing with behavioral data, however, is an explosion in data volumes that quickly outpaces our ability to extract insights. Whether you're talking about metadata management or related techniques to assign meaning and value, the simple fact is that our human brains don't scale the way data does. For every product manufactured there can be hundreds or thousands of data points leading up to the assembly. And for every purchase transaction there can be more than a hundred interactions, especially with e-commerce.

Think of the humble visit to a brick-and-mortar store in the era before big data, or even today. By our own estimates, that process of walking the aisles and buying something at checkout may create about 100 bytes of data involving the item, the price, credit card information, and probably not much more information than that.

By comparison, today's average visit to an online retail site can generate hundreds of kilobytes of data—everything from the items and customer reviews you viewed (and for how long) to what products you placed in your online shopping cart and whether they stayed there through checkout. Any one of these points is subject to analysis and further generation of data.

It has been said that what you can learn from an online customer today is akin to tracking every single step during a physical visit to a store; it's like having a GoPro video camera mounted on the head of every customer who walks through the door, and then collecting and analyzing the entire video stream, frame by frame. But if you don't have structures to capture and responsibly manage that data—including provisions for privacy, data protection, and information security that we'll talk about toward the end of this book—then you won't really be able to leverage that information as fully and responsibly as possible.

A NECESSARY EVOLUTION

These and other issues reverberated in our heads long after our first meeting in Evanston. Over the next days and weeks—in airplanes and taxis, in gaps between meetings and other spare moments—we formulated a nascent framework to align our thoughts and energy. We kept in touch throughout, and every call and e-mail helped flesh out our understanding of what has become the Sentient Enterprise journey.

The more we aligned, the more we were able to understand that this journey is an evolutionary one that is disrupting entire industries and changing the very nature—the very identity—of large companies across nearly all sectors of our economy. Many talk about what the enterprise can do with data and analytics; but it's crucially important to also understand just what data and analytics are doing to big business.

As Cisco's and GM's own stories earlier in this chapter make clear, no industry is safe from disruptions in how we leverage data and data products, and no large company can afford to sit on the fence. Big data is a disruptive and unstoppable wave, and the leaders who drive change in their organizations are the ones who recognize they have no choice but to become disrupters themselves, before some other company does it first.

These pressing realities position the Sentient Enterprise journey as a necessary evolution for survival—a mandate rather than an option. It's what we call "change management on steroids," and we need to heed the clarion call for change in ways that don't make the problem worse.

We'll explore in these pages, for instance, how important it is to democratize access to data throughout your organization for faster decision making. And we'll see how the key to scaling insights and keeping your business user community aligned is to tailor the environment for maximum collaboration—a "LinkedIn for Analytics" approach that borrows lessons from social media, gaming, and other areas where users naturally are driven to engage.

We can't achieve this agility, however, through shortcuts like throwing away documentation and governance. That kind of Wild West approach will do more harm than good via data drift, duplication, and destructive information anarchy. There are many such challenges in the overall bid to adapt our technology to the needs of our employees, customers, and the entire business ecosystem.

PUTTING IT ALL TOGETHER

A couple of months after first connecting in Evanston, the two of us met again, this time in San Diego, to flesh out more fully what this evolution toward the Sentient Enterprise looked like. We were certain by then that our collaboration around this journey was an important one that was just waiting to be rendered. In fact, our mutual friend and Oliver's Teradata colleague, Mary Gros, arranged for our discussion to be literally rendered by a graphic artist who joined us in a conference room overlooking the nearby peak of Mount Woodson. As the artist scribbled away, we talked through a skein of insights until we aligned on five foundational qualities that define the Sentient Enterprise.

The Sentient Enterprise is:

- **Proactive** and able to sense micro-trends signaling the next crisis or the next opportunity.
- **Frictionless** in that it can act as one organism, human and machine, without any impedance or bottlenecks created by silos.
- **Autonomous** in its ability to listen to data and make decisions in real time without too much human intervention.
- **Scalable** to virtually any size company, with the ability to leverage unlimited amounts of data for making decisions.
- **Evolving** through intelligence that is native and emergent.

In short order, we continued to refine this journey into its systematic, five-stage evolution and hit the road with a joint slide presentation to executives at the Walt Disney Company, Boeing, BMW, and other major corporations. As the nods of agreement added up in boardrooms and C-suites around the world, we knew we had to take the next step and write the book you're now reading.

As mentioned in the Introduction, our long-term vision for the Sentient Enterprise is not unlike the popular Six Sigma model for manufacturing in that we've established a specific framework for excellence that can be broadly applied. Just as Six Sigma has been modeled by countless business strategists since its inception by Motorola in 1986, our hope is that the five-stage Sentient Enterprise model will gather momentum across the many industries that rely on data for decision making and value.

Like Six Sigma, the Sentient Enterprise is a business phenomenon that is already underway, yet always aspirational. We can and will always do better. But we have seen enough of the Sentient Enterprise puzzle in place to know it's a viable model for orders-of-magnitude leaps in business performance and value that aren't really possible any other way. It's the threaded needle, the secret sauce, a recipe for agility and productivity at any scale.