

1

The Broken Promise of Customer Data and Technology

Of all the different types of data within organizations, customer data is perhaps the most important. It informs a company about what customers have bought, what kinds of people they are, and what they might buy in the future. It is the key to knowing whether customers are satisfied or dissatisfied. It reveals what products and services they might find appealing in the future. At the most basic level, customer data includes where to ship the product, how it's being paid for, and whether the customer is likely to renege on payment or return the product.

We typically think of only companies as having customers and customer data, but government agencies and nonprofits have their equivalent. Governments can apply the science of customer analysis. Knowing how constituents feel about their leaders and the services they receive is extremely important to their success. Nonprofits have both contributors and recipients of services; they're different types of customers, but keeping track of both is critical to success.

We will argue in this book that customer data and technology have often not delivered on their promises or potential. Perhaps that's a good thing for

us, or you would have little motivation to buy and read these words and images. For more than 30 years, we've heard about one-to-one marketing, a 360-degree view of customers, personalization and hyper-personalization, multi- and omni-channel relationships, etc., *ad infinitum*. Despite decades of progress in the science and technology of analytics and customer modeling, the practical application of these methods has remained elusive.

If you're not skeptical of these concepts, you probably aren't paying a lot of attention to the gap between how marketing is described and how we observe it in practice. Of course, the appeal and importance of these concepts and capabilities is unquestionable. It's certainly possible to imagine a world in which customer data and technologies were used very effectively to improve the course of customer relationships. Analytics and artificial intelligence have already had a positive impact on these relationships, and generative AI and agentic technology will go much further.

If these technologies are used effectively, what will the future look like? If customer data were captured, stored, and used effectively, for example, we would know about virtually every interaction a customer has with an organization—purchases, returns, product or service usage for online offerings, web and mobile views and clicks, related social media comments, summaries of calls to the customer service center, and even subjects addressed in informal chats with salespeople or other company representatives. All of that information could be combined into customer lifetime value predictions, propensity-to-purchase models for particular products and services, churn predictions, accurate personalization of virtually every ad or offer, and detailed attribution of a customer's channel activity and path-to-purchase. Many of those analyses would be powered by traditional analytical AI based on machine learning and structured data.

Generative AI has further potential to enhance customer relationships. By harnessing the capabilities of machine learning and data integration, it could automate many textual and visual offerings to customers and provide extensive personalization of each interaction. It could analyze a customer service email or the content of a phone call or text chat, pass along relevant feedback to the individuals involved in addressing the customer need, and alert the customer about what action has been taken. The idea of standard product copy or images would become obsolete. "Agentic" AI, or a somewhat autonomous AI system that can make decisions and take actions rather than just responding to prompts, could automatically determine a

customer need, make an immediate recommendation of how to address it, and complete the purchase transaction with customer approval.

Much of this glittering image of data- and technology-enhanced customer relationships is technically possible today. However, as we'll discuss later in this chapter, at various times in recent history we've been promised other revolutions in customer interaction, and we're still waiting to see them. It's a rare event for either of us, for example, to receive a marketing offer from a company for a product we really want but haven't already searched for. In the remainder of this chapter, we'll discuss the history of customer data and technology and explain why a scientific approach to customer relationships hasn't been achieved.

The Nature of Customer Data

As in any scientific endeavor, effective customer relationship management requires rigorous collection and analysis of data to generate meaningful insights. Customer data—for better or worse—is multifaceted. It consists of such elements as these:

- **Transactions:** Sales, shipments, payments, website and mobile clicks, customer support tickets, and the like.
- **Attributes:** Names and addresses, demographic information, credit card numbers.
- **Structured opinions:** Customer satisfaction ratings, net promoter scores, whether purchases are returned; web and mobile site views.
- **Unstructured opinions:** Website comments, social media comments, customer service calls, comments to salespeople.
- **Calculations:** Segmentation, engagement, acquisition cost, recency, frequency, monetary (RFM) analysis, customer lifetime value.
- **Predictions:** Recommended products or services, likelihood of attrition (churn), price the customer is willing to pay.
- **Experimental results:** A/B tests, controlled trials, and other scientific methods of testing hypotheses about customer behavior and preferences.

Customer data, then, can be a rich repository of content about one of the most important resources in business. It's hard to imagine how an organization can be very successful without managing customer data well.

Amazingly, however, most don't do it well at all. They usually manage to stay in business, but they don't prosper as much as they might if they had and used good customer data.

Why Customer Data Is Difficult to Manage

One reason customer data often isn't managed well is simply because it's complex. Vendors often speak of their "Customer 360" capabilities, suggesting that they can handle and make sense of all forms of customer data. But as you can imagine, capturing and analyzing all the different types of data we've just described can be difficult. And doing it successfully goes well beyond software alone.

Just to keep that set of customer data elements we described earlier in play, imagine the difficulty of combining all of them in a 360-degree view of customers. Customers are always changing (addresses, credit cards, opinions), so it's hard to keep up-to-date with them. Mixing structured and unstructured data is both a technical challenge and—more importantly—an analysis and sensemaking challenge. What if the customer continues to buy from us but trashes us on social media? What if they buy a lot of our product online but return much of it after purchase? What if structured data predicts churn, but the sales rep says the customer is happy? Such potential contradictions can be difficult or impossible to resolve.

There is also the difficulty of even learning who the customer is. You may have wondered, for example, why manufacturers of durable goods often package in little cards to encourage the customer to register for warranty service. It's usually not technically necessary to register to get your warranty service, but the manufacturer wants your name and address data to try to sell you additional products. Incentives and offers included in the product packaging often have the same motivation.

This is a particular challenge for consumer goods that are sold through retail channels. In such cases, the retailer knows what you have bought, but that information is seldom passed along to the manufacturer of the product because retailers want to own the relationship with the customer. This issue has bedeviled consumer packaged goods (CPG) manufacturers for decades. As we'll discuss, some CPG firms have made considerable progress in selling to—and getting data from—consumers

through ecommerce channels. Not only does that usually increase their profit margins, but it also permits an ongoing marketing, sales, and service relationship with the customer.

That previous sentence suggests another difficulty with customer data: it's often captured and owned by different parts of an organization. Marketing owns personalized customer offers and perhaps leads from all channels. But to convert, those leads must be handed off to sales (inside or outside), which then owns them and customer information encountered through the sales process. Customer service is responsible for the problems that customers bring to an organization and the solutions to them. If there is a "customer success" function, it has custody of the information about how the customer is using the product or service and how they feel about it. Analytics or AI organizations may own the data used to train models and the predictions they make. If that's not enough owners, the research and development (R&D) group may own customer comments in focus groups or suggestions for new products and services.

Unless your company has an unusually friendly and collaborative culture, these different owners may clash when someone suggests pulling all the information about customers together into one database or customer data platform. And the ownership of that database or platform may be contentious as well. Those political and cultural issues can't be handled at all by software, even if it's highly flexible and functional. All of this has led many managers to give up on the 360-degree-view idea and to just try to manage the data they own or control about customers.

Perhaps the most basic problem of all in managing customer data effectively is defining the customer with precision and employing that definition around the company. Table 1.1 describes a set of alternative terms that could all be used to describe a "customer," for example. Tom has a vivid memory of this issue from a consulting engagement he performed with a computer company a few years ago. The company had two major distribution channels: direct sales to the end customer (generally businesses), or sales through distributors. The different channel organizations couldn't agree on whether a distributor counted as a customer. Even more problematically for financial reporting, they couldn't decide on what constituted a sale—if a product was sent to a distributor, did it count as a certified sale?

Table 1.1 Multiple terms that could mean the same individual

Companies	Government
User	Citizen
Visitor	Voter
Shopper	Resident
Buyer	Inhabitant
Purchaser	Constituent
Consumer	Taxpayer
Customer	Defendant
Client	Petitioner
Patron	Subject

The direct sales unit was the larger and more powerful of the two, so their definition was more likely to be employed in the organization, but the indirect sales organization didn’t change their definitions—sometimes leading to multiple versions of the truth. Even with the direct sales organization, there was debate about whether a prospect should be viewed as a customer and included in the customer database. Some people included prospects; some didn’t. Both of these arguments led to long delays, for example, when a new chief executive officer (CEO) asked for a list of the top 100 customers for a letter he planned to send to them upon assuming the role. As a result of these and other interactions with clients, Tom formulated a general law of common information in companies (see “Davenport’s Law of Common Information”).

Davenport’s Law of Common Information

A couple of decades ago, Tom was researching how companies managed information. He visited several of them in a short period. When he visited American Airlines, the head of information management told him, “We have 11 different meanings of the term *airport*” and proceeded to describe some: maintenance people think that anyplace you can fix a plane is an airport, passenger-focused employees think that anyplace you can pick up or drop off passengers is an airport, the employees who work with the International Air Transport Authority (IATA) think that their list of airports is definitive, etc. Then Tom heard

something similar at Union Pacific Railroad the following week: “We can’t agree on what a *train* is.” Some believe it’s an abstract scheduling entity, some believe it has to have a locomotive, some think it’s any collection of cars on a track. Then the next week someone at the US Department of Justice commented, “We have great arguments about what constitutes a *trial*.” Finally, Tom got the picture. He formulated Davenport’s Law of Common Information, which goes this way:

“The more an organization knows or cares about a specific entity within its business, the less likely it is to agree on a common term and meaning for it.”

Tom has seen many examples of this law in practice since he formulated it, and *customers*—being an entity that many organizations deeply care about—are one of the best examples. However, the fact that Davenport’s Law of Common Information exists does not absolve organizations from trying to do something about it! Even if it’s difficult, for example, to agree on a common definition of *customer*, it’s important to do so, and managers should argue around a table until they agree on what it means for the entire organization. Perhaps needless to say, for IT people to create complex customer data models—some rival microprocessor circuit diagrams in visual complexity—does not generally solve the problem.

There is no right answer to these questions, but not having agreement across an organization on what constitutes a customer (and a sale to a customer) is problematic. It may be perfectly reasonable to treat prospects as customers for some purposes (marketing, for example), but in general, it’s better to clearly define *customer* in a particular way (e.g. someone or some organization that has bought something from us in the past two years) and treat others as “prospects” or “lapsed customers” as their official status. Then if there is an ad or an offer that applies to more than one type of customer, the different groups can be combined temporarily for that purpose.

Organizational structures can also lead to other types of customer data problems. Take, for example, the merger or acquisition. Even if the acquired company isn’t in the same industry, there may well be overlapping customer records—perhaps with slightly different data attributes. Companies have

worked for decades to create integrated, high-quality “golden records” or important data types like customers, but it has been a challenging process that is rarely finished successfully.

There are AI technologies now that can help with data integration processes.¹ They typically work by using traditional machine learning to predict whether two different records are actually for the same customer. Such a system might predict with 95% confidence, for example, that “John Smith, 100 Main St., Lenexa, KS” is the same customer as “J Smith, 100 Main, Lenexa, KS.” That level of confidence would generally be high enough to automatically merge the two records. These types of systems can substantially lower the time needed to integrate data, although they usually still require occasional human intervention.

Finally, there are the pernicious data quality problems that plague virtually every organization. These often originate at the front line, when a sales rep enters the customer’s name, address, or other information incorrectly at first contact. Then it becomes difficult to fix the mistake. We won’t go into detail on data quality issues—and there are also some partial technological solutions to the problem, like AI that identifies unlikely names or addresses that don’t exist—but they continue to be found almost everywhere. Fixing them is possible but typically a labor-intensive and expensive proposition.

Where Did We Go Astray?²

So, customer data is hard to manage well. But as we argued at the beginning of this chapter, it’s more important than just about any other type of data, so it’s worth some special effort. Perhaps the real problem is that we have been overly optimistic about our ability to use technology to master customer data. Like the customer 360-degree meme mentioned earlier, there has long been a set of messages that we are just about to turn the corner on customer data management and the resulting ability to understand, act on, personalize, and profit from customer activity.

¹ Thomas H. Davenport and Tom Redman, “How AI Is Improving Data Management,” *MIT Sloan Management Review*, 20 December 2022, <https://sloanreview.mit.edu/article/how-ai-is-improving-data-management>.

² This section draws in part on Jim Sterne and Thomas H. Davenport, “A Brief History and the Future of Customer Data,” submitted to *Applied Marketing Analytics*, 2024.

Table 1.2 Technologies for customer transaction capture and execution

Mature or obsolete	Well-established	Emerging
Customer databases	Ecommerce platforms	GenAI service
CRM systems	Customer data platforms	CX platforms
Self-service kiosks	Mobile commerce platforms	Composable CDPs
Interactive voice response	Service management/ticketing	Augmented reality
Retargeting	Data management platforms (ads)	Voice search
	Search term capture	GenAI prompt capture

Information systems have long been proposed as the answer to deeper, better, more rewarding customer relationships, and companies have been focused on this issue for decades. Various new technologies (databases, the Internet, analytics, customer relationship management [CRM] systems, etc.) have been advanced over time as the way to execute and capture customer transactions. Other technologies (descriptive customer analytics, predictive customer analytics, machine learning, and now generative AI) have been held up as the answer to all our customer insight and prediction challenges. We review the current state of these technologies in Table 1.2 (and Table 1.3 later in this chapter). They could ostensibly support greater customer awareness, understanding, affinity, engagement, conversion, and loyalty, all with the click of a mouse.

It didn't turn out that way, at least not yet. Customer transactions are no doubt easier with the Internet and mobile apps, and we have massive amounts more data about our customers than in the past. However, marketing effectively to them remains problematic. We still struggle with capturing, understanding, and acting on the behavior, sentiment, and voice of the customer. We have difficulty making high-quality offers or purchase recommendations. Post-sale customer service is hardly a pleasant experience in most cases. There is little or no evidence that our technology-enabled marketing to customers has increased customer satisfaction or lifetime value—the primary outcomes for which marketers strive. We'll next review why these technologies have not yet yielded the promised benefits to either customers or marketers.

Getting Customer Data into a Database

In the beginning, the first transactional technology for customer data was database marketing. It was to be the driver of dramatic change in the relationships among customers and their suppliers. The idea was that technology could capture—all in one place—important attributes about customers and their relationship with a company. This trend emerged in the 1980s, when companies began to compile data on how their customers transacted with them—what they bought, how much they spent, and as much demographic information as they could compile. The goal was to begin to personalize marketing offers based on historical relationships with customers and future predictions of what would appeal to them.

One early example of database marketing was at AT&T, which in 1984 created the “Opportunity Calling” marketing campaign based on a customer database. It was intended to provide personalized offers to the company’s 85 million customers as it divested itself of the Bell Operating Companies.³ The field of direct marketing was still relatively new, and the *Journal of Direct Marketing* (in which the AT&T initiative was described) began to be published in 1987.

In 1991 two marketing academics, Robert Blattberg and John Deighton, published an article discussing the possibilities for using data to personalize marketing. Like many other articles since, it could still be relevant today. It said, for example:

To appreciate the power of a customer database, one must see it not merely as a mailing list, but as the memory of the customer relationship: a record of every message and response between the firm and each address. Add artificial intelligence and the system can design new messages, and even product offerings, at the individual level to reflect everything learned from past interactions.⁴

Although these feats sound like something requiring generative AI, they meant rule-based expert systems at the time, and neither the authors nor

³ Gary F. Beck, Sherry M. Karas, and Ann E. Skudlark, “The AT&T opportunity calling program: Responsiveness and implications,” *Journal of Direct Marketing* 4:1 (1990), pp. 24–29.

⁴ Robert C. Blattberg and John Deighton, “Interactive marketing: Exploiting the age of addressability,” *MIT Sloan Management Review*, Fall 1991, <https://sloanreview.mit.edu/article/interactive-marketing-exploiting-the-age-of-addressability>.

anyone else was aware of what might be done with newer forms of AI in the future. Putting all customer information in one place was certainly a good idea, but there were both organizational and technical barriers to integrating the information and using it effectively.

Our One-to-One Future, and Why It Has Been Slow to Arrive

Shortly thereafter in 1993, two consultants, Don Peppers and Martha Rogers, coined a term that has remained popular for 30 years, although it has seldom been fully realized. *One-to-one marketing* promised highly personalized offers, ads, and communications to customers based not just on their attributes and purchase histories but also on their individual interactions with the marketer.⁵ They suggested that the end of mass marketing was upon us, as companies would soon be worrying less about how to sell an individual product or service to as many customers as possible, and more about how to sell an individual customer as many products or services as possible.

Since that time, marketers have struggled to achieve the one-to-one future for a variety of reasons. In addition to the problems we described earlier, customer data—particularly from third-party providers—is often inaccurate.⁶ Companies don't usually have enough different offerings or messages to fully personalize them to each customer. And the individually targeted matching of customers to marketing content has been beyond the capability of the largest computers and strongest algorithms. Creating truly one-to-one content would be very labor intensive—at least until the recent emergence of generative AI.

To understand what happened with the one-to-one concept in practice, we talked with two people who were and are intimately familiar with it. Don Peppers, the co-author (with Dr. Martha Rogers) of *The One to One Future: Building Relationships One Customer at a Time* (and the author of the Foreword to this book), is justifiably impressed with his and Rogers' book's predictions but dissatisfied with actual outcomes. The book anticipated

⁵ Don Peppers and Martha Rogers, *The One to One Future: Building Relationships One Customer at a Time* (Crown Business 1993).

⁶ Alyssa Boyle, "CTV ad targeting is getting more advanced, but data quality is not," Adexchanger, 20 October 2023, <https://www.adexchanger.com/tv/ctv-ad-targeting-is-getting-more-advanced-but-data-quality-is-not>.

widespread personalization, a focus on customer relationship management, greater use of analytics in marketing, and even the rise of social media, remote work, and direct-to-consumer ecommerce—despite that the Internet was a largely academic resource in 1993, when the book was originally published. The first commercially useful web browser, Netscape Navigator, did not even become available until 1994.

In a conversation with Tom, Peppers laid the blame for the one-to-one future not being more broadly realized on organizational structure. Most companies, he said, are still organized by product and are effective practitioners of product management. Rarely, however, is anyone responsible for making customers happy and increasing their lifetime values, one customer at a time.

Peppers commented further:

One of the first things we said would happen in the book is that while mass production generates economies of scale, making it possible to produce each product for less and less unit cost, one-to-one marketing generates “economies of scope,” which reduce the cost of selling to any individual customer, based on the scope of a marketer’s relationship with that individual customer. But one-to-one marketing requires more than interactive technology (like the internet); it also requires computer-assisted production of individually tailored products. We called it “customerizing” the product, describing a process that soon came to be called “mass customization.”

Ultimately, however, Peppers says the success of any comprehensive one-to-one marketing effort depends on whether and how a company assigns responsibility within its organization for selling more products and services to individual customers (as opposed to selling their product or service to more customers, which is what product managers do). Peppers said that software-as-a-service companies are among the few that are beginning to achieve this one-to-one ideal, because not only do they treat different customers differently, customizing their offerings to specific customers, but they have also established customer success managers who are individually *responsible* for maintaining and increasing the profitability of each customer. And he predicts now that as more and more products and

services are connected to the Internet, this “customer success” organization model will become more pervasive, even among consumer products.

However, Peppers’ and Rogers’ 1993 book had a chapter entitled “Make Money Protecting Privacy, Not Threatening It,” and in other writing he has lamented the behavior of the big tech companies in general. In a 2020 blog post called “The One to One Future—Are We There Yet?” he answered the question:

So we’re *definitely* not there yet, folks. No way. Because not only has this prediction not yet become reality, but most of the “Big Tech” businesses driving so much economic growth today make their money by undermining privacy, not by protecting it.⁷

One of Peppers’ and Rogers’ colleagues in their consulting firm Peppers & Rogers Group was Hamit Hamutcu, who now works at the Institute for Experiential AI at Northeastern University and also runs a company focused on data science skill standards development and assessment. He was a managing partner at Peppers & Rogers Group, launching several offices globally, and at one point became the global analytics lead for the company.

When we asked him about his perspectives on the difficulties companies had with one-to-one marketing, he said there was often a gap between executive sponsorship and the reality on the ground. Client executives—typically the CEO or chief operating officer (COO)—wanted the work to be transformative and dramatic. They wanted a culture of customer centricity, with channels and marketing initiatives built around it, daily business processes designed around it, and technology to support it.

However, the easiest and most obvious changes and investments involved technology—data warehouses, CRM systems, and analytical tools. These were necessary investments, Hamutcu said, but the technology wasn’t quite up to the job and was certainly not sufficient to drive the desired impact. There were no real-time analytics or machine learning models; any analytics usually had to be done after the customer transaction. There were many different repositories of data, and integrating them was always

⁷ Don Peppers, “The One-to-One Future—Are We There Yet?” LinkedIn post, 22 July 2020, <https://www.linkedin.com/pulse/one-future-we-yet-don-peppers>.

challenging. Companies would define customer segments, but they would change faster than the organization could keep up and define new ones.

There were also organizational challenges, Hamutcu noted. Organizations were siloed by business unit and didn't want to collaborate on marketing or share data. The budgets and power, as Peppers told us, were often with the product organization. The consultants would devise customer-oriented actions, but deploying them consistently across a large and complex organization with a variety of goals and incentives was often hard. And the sustainability of the one-to-one focus was often a problem. A CEO makes a strategic mandate, but if the CEO changes, the next incumbent might have a different vision.

Three lessons are clear from Peppers' and Hamutcu's perspectives. One is that technology—albeit technology that is much more powerful and capable than it was in the 1990s and early 2000s—is not enough to create a customer-centered organization. Two is that organizational and cultural issues are bound to come into play within specific organizations. Three is that the one-to-one marketing organization is not a passing fancy; it must be articulated clearly, be built over time, and be a persistent goal for multiple leaders and the employees who follow their lead.

From Database to Digital

As we've noted, Peppers' and Rogers' one-to-one vision largely predated the Internet and e-commerce. But the advent of these technologies brought a resurgence of hope to the prospect of one-to-one marketing due to the enormous amount of behavioral data collected in the serving of web pages. The ability to collect information about an individual visitor's clicks, time spent on a web page, operating system, browser type, screen resolution, (approximate), geo-location of Internet Protocol (IP) addresses, and even the X-Y coordinates of their mouse movements seemed like a data dream come true.

However, data difficulties immediately surfaced that have yet to be rectified, including the following:

- Cookie deletion
- Cookie blocking
- Multiple machine browsing
- Multiple browser browsing

- Multiple people on the same cookie
- Nonhuman traffic
- Dynamic IP addressing
- Page caching
- JavaScript loading
- Pixel placement

Across a large organization, the data was further mangled by the following:

- Different tools using...
- Different date stamping routines and...
- Different methods to capture...
- Different types of data to store in...
- Different kinds of databases with a...
- Different method of data cleansing followed by...
- Different segmentation analysis to produce...
- Different kinds of reports that ended up in a...
- Different pipeline for integration into...
- Different data warehouses.

Marketers attempted to derive indicators of individual interest by reading the tea leaves of search, navigation, recency, frequency, and more. They successfully identified website friction as a signal of customer experience distress ripe for improvement, but dynamically serving content for each individual proved complex and painful, and even the best efforts failed to produce a decent return on the effort.⁸

As a result, even in a highly configurable medium like digital marketing, it is still rare for consumers to receive well-personalized marketing content from companies. One need only view the prevalence of retargeting in digital channels—assuming that if customers were once interested in a product or service, they will be interested again—to realize that we are far from the one-to-one goal.

⁸ René Arnold, Scott J. Marcus, Georgios Petropoulos, Anna Schneider, “Is data the new oil? Diminishing returns to scale,” 29th European Regional Conference of the International Telecommunications Society (ITS): “Towards a Digital Future: Turning Technology into Markets?” Trento, Italy, 1–4 August 2018, <https://hdl.handle.net/10419/184927>.

When social media took off, marketers saw a new data stream to attach to individuals: their own words posted in public. The sheer amount of data led to a boom in sentiment analysis companies peddling “insight” into individuals’ frames of mind. But responding to a social post by hand was too labor-intensive, and doing it programmatically was seen as creepy. Then came the Cambridge Analytica debacle, and the legislation started rolling out. Not only did trying to use individuals’ social posts as input for altering their opinions violate Facebook’s terms of service, it violated society’s sense of decency.

Today, organizations are beginning to collect and manage all types of customer feedback, including surveys, questionnaires, and reviews. This information is increasingly compiled and accessed through *customer experience platforms* (CXPs). But this is still an emerging category of customer technology, with variations involving survey-oriented feedback, service feedback, sales account feedback, customer journey mapping, and customer success tools.

From Call Centers to Chatbots

Technologies for customer service have also evolved over the years but remain frustrating to many customers. Many companies are heavily focused on the efficiency with which customer problems or questions are addressed in call or contact centers. They view customer service largely as a cost to be minimized. Therefore, they have embraced technologies that allow fewer human agents to be employed in these roles. Many organizations also outsourced customer contact to offshore providers to lower costs.

Call centers date to the 1960s, but the earliest automation tools for them—known as *interactive voice response* (IVR)—only became popular in the late 1970s. We are all painfully familiar with that technology; it involved pressing touch-tone buttons to reach different types of answers. The most popular button, of course, was 0 to reach a human agent. Later IVR technologies included some speech recognition capabilities, and customers were able to (occasionally) meet their service need by talking to the IVR system. Again, however, the most popular phrase uttered by customers was, undoubtedly, “speak to an agent.” And whether the caller got to speak with an agent or not, there was usually no attempt to capture the nature of the

customer's issue for later attempts at resolution. There are exceptions to this, of course—in the past typically involving pre-generative forms of AI.⁹

A further attempt to reduce human customer service assistance came with the introduction of first live human text chat and then chatbots. Chat with humans often reduced the resolution time compared to calls, and—perhaps most appealing to companies—it allowed agents to handle chats with multiple customers at once. This was often frustrating to customers who had to wait for a response.

Fully automated chatbots, also known as *virtual assistants*, made their corporate customer service debut in the 2010s. One such chatbot, Ask Jenn at Alaska Airlines, answered more than 28 million customer questions—probably with greater and lesser degrees of success. Like IVR, chatbots can be frustrating to customers, but they are slowly becoming more capable. Smart companies allow frustrated customers to reach a human agent if a chatbot can't solve their problem. Some chatbots collect customer query content to understand where the chatbot fails and to improve its performance over time.

The next phase of chatbots is clearly represented by generative AI. We are still in the early days of applying generative AI to customer service, and we'll describe that further in Chapter 6, "Better Customer Voice Analysis and Action." But customers are justifiably skeptical. We've been promised better post-sale service from chatbots, AI-based knowledge management, and previous technologies as far back as the 1990s.¹⁰ AI-enabled or not, however, customers remain dissatisfied with machines as customer service agents. A recent survey found, for example, that only 14% of customers would prefer to speak to an AI-based service system than to wait (between 1 and 11 minutes) for a human service agent.¹¹ We hope that generative AI will change that preference, but it's too early to tell whether it will.

⁹ See, for example, Thomas H. Davenport, "Extracting Customer Support Knowledge with SupportLogic," *Forbes*, 29 October 2021, <https://www.forbes.com/sites/tomdavenport/2021/10/29/extracting-customer-knowledge-with-supportlogic>.

¹⁰ Thomas H. Davenport and Philip Klahr, "Managing Customer Support Knowledge," *California Management Review*, Spring 1998, <https://journals.sagepub.com/doi/10.2307/41165950>.

¹¹ Christopher Zara, "AI Agents Are Taking Over Customer Service, but Most of Us Would Rather Wait for a Human," *Fast Company*, 30 April 2024, <https://www.fastcompany.com/91115626/ai-chatbots-vs-humans-customer-service-tasks-survey-ratings>.

From Index Cards to Customer Relationship Management Systems

A key technology that—once again—promised to revolutionize customer focus was the customer relationship management system (CRM). Initially developed in the 1980s and early 1990s, there were two primary types of customer transactions that CRM was designed to support. One was sales processes; that technology was initially called *sales force automation* (SFA) and was eventually wrapped into CRM systems. The idea of SFA was to enable salespeople to automate many time-consuming activities—entering orders, checking available inventory, monitoring their progress against goals—to free their time up to build customer relationships. It also served a valuable function for the organizations that implemented it, in that it shifted much customer information from index or Rolodex cards to a system of record. If the salesperson left, the customer information stayed with the company. That was, of course, one factor that inhibited many salespeople from adopting it.

The other primary function of CRM systems was automating customer service transactions. They enabled service departments to create “tickets” that kept track of service requests. It was relatively rare for the information in such systems to be analyzed. Instead, the focus was on automating service transactions. By the mid-1990s, vendors like the market-leading Siebel Systems (now owned by Oracle) had combined SFA and service capabilities and began to call it CRM. Some marketing functions were included as well in CRM, but again the focus was on transactional capabilities—campaign management, outbound offer delivery, and so forth. Few CRM systems had, for example, any ability to predict what a customer might do.

However, capturing the various forms of data involved in customer relationships was a big step forward for organizations. Having all or most transaction data in one place—e-commerce transactions were often not included—made it much easier to monitor, analyze, and act on customer behavior. Since the early days of CRM, cloud-based providers like Salesforce.com have added increasing numbers of functions, including data analysis and prediction, e-commerce, mobile access, and even some AI capabilities. Enterprise resource planning (ERP) providers like SAP and Oracle also added CRM offerings to their software suites, making it possible to connect financial, supply chain, and HR information to sales and service information.

From Multiple Databases to a Customer Data Platform

Places to store customer data have proliferated over time. As we've mentioned, CRM systems ended up being one repository, but they often didn't include content such as ecommerce transactions, social media comments, and customer emails. Many companies often had separate and overlapping customer databases for sales, marketing, service, and other functions. The "360-degree view of customers" concept became an objective again when vendors began to offer "customer data platforms." There are many versions of the concept, but the overall idea is, as Forrester Research describes, "A CDP centralizes customer data from multiple sources and makes it available to systems of insight and engagement."¹² This means that customer data is extracted from the transactional systems that gather it, collected in a CDP, and analyzed there for decisions and actions. The customer identifiers may differ slightly across these different sources, so some "identity resolution" is often necessary.

CDPs have, according to surveys of marketers, achieved high levels of adoption in organizations. A 2023 Gartner survey, for example, found that 67% of marketing technology respondents in companies, technology survey said they had installed a CDP. However, only 17% of marketers reported high utilization of the CDP.¹³

There are several reasons why utilization may have lagged. The effective use of CDPs requires a high level of discipline within an organization to aggregate customer data in one place, and even more to refine it over time to maintain and improve customer data quality. They also require a high level of collaboration across customer-related functions. Marketing departments, as the most analytical customer-oriented function, tend to be the primary beneficiary of CDPs, but other functions that deal with customers have to supply data in the right formats. In addition, a lot of investment is required before companies see returns on the investment in terms of increased customer purchases or loyalty.

Although we believe CDPs are a good idea, we haven't seen them revolutionize customer relationships yet. Perhaps the low level of utilization

¹² See, for example, "Explore the CDP Vendor Landscape in Asia Pacific," Forrester Research website, <https://www.forrester.com/blogs/explore-the-cdp-vendor-landscape-in-asia-pacific>.

¹³ Hana Yoo, "CDPs Are in the Gartner Hype Cycle 'Trough of Disillusionment'" AdExchanger website, 22 February 2024, <https://www.adexchanger.com/data-exchanges/cdps-are-in-the-gartner-hype-cycles-trough-of-disillusionment>.

found in the Gartner survey is a factor in that not happening. CDPs are best suited for business-to-consumer (B2C) companies that have substantial amounts of first-party customer data. Business-to-business (B2B) companies typically have substantially fewer customers, and some are more important than others—perhaps weakening the case for onboarding all customers into a CDP. Companies that distribute through retailers or other intermediaries to consumers (B2B2C) typically have challenges in getting detailed data on their end consumers.

Nevertheless, some B2B2C companies are making progress toward customer data integration on CDPs. The consumer products giant Nestle, for example, reported to investors in 2022 that it planned to increase the number of consumer data profiles in its CDP from 248 million halfway through 2022 to more than 400 million by 2025.¹⁴ This increase in customer data is consistent with the company's goals for direct consumer sales through ecommerce, which it hopes will reach 25% of revenues by 2025 (from about 16% in 2022). Nestle said it was already seeing substantial benefits (measured in individual country performance) from better consumer data, including a much lower cost-per-click in digital advertising, improvement in return-on-ad-spend, and better ad recall through content personalization.

If broad utilization of CDPs for customer data management spread through more organizations, it is likely that the organizations that can afford them will see considerable improvement in customer relationships and the ability to market, sell, and service customers effectively. If Nestle—a company that historically sold through retailers and did not possess much consumer data at all—can transform itself in this fashion, many other companies can, too.

Some vendors are beginning to argue for “composable” CDPs, which don't move data into a new platform but make it accessible from existing sources. Given the challenges of moving customer data into a single location, this may be an attractive option for some companies, although it is still in the early stages. Some organizations use the broader concept of “data fabric,” a virtualized data layer that allows users to access data without moving it from its original location, such as a data lake, data warehouse,

¹⁴ “Accelerating Our Data-Driven Digital Transformation,” Nestle investor seminar 2022 presentation, <https://www.nestle.com/sites/default/files/2022-11/investor-seminar-2022-digital-transformation.pdf>.

relational database, or ERP system. Data fabric can also mix and integrate data from on-premises, cloud, and hybrid cloud platforms. A composable CDP is effectively a data fabric for customer data.

From Transactions to Analytics and AI

As customer data has gotten better over time, companies and other types of organizations have begun to use analytics and AI to understand, predict, and classify customer behaviors. Now, with generative AI, they can (with some cautionary steps) communicate directly with customers using AI. This is a big step forward and a primary purpose for which much customer data was gathered in the first place.

So what’s the state of customer data analytics and AI? The technology itself for this purpose is well-established. Every company has access to the cloud processing and machine learning algorithms to analyze customer data and put it to work. They’re even cheaper than they were in the past—largely because of open-source algorithms and cloud processing that doesn’t require a big computer purchase. Many analytical tools are now mature (Table 1.3).

Table 1.3 Technologies for customer analytics and insights

Mature or obsolete	Well-established	Emerging
Predictive purchase models	Multi-channel attribution	Automated machine learning
Descriptive customer analytics	Next best offer	Deep learning models
Spreadsheet models	Personalized content	Generative AI-based models
Rule-based segmentation	Digital attribution	Generative AI chatbots
	Dynamic pricing	Gen AI sentiment analysis
	Churn modeling	Machine learning operations
	Sentiment analysis	
	Statistical segmentation	




Figure 1.1 Companies vs. consumers perception gap.

Some other factors remain problematic, however. High-quality customer data, as we've discussed, is still hard to create. Data scientists—or even marketing, sales, or service professionals with a strong quantitative orientation—are still expensive and somewhat rare, at least outside of Silicon Valley. Business leaders in those domains may not be aware of what can be done in data-driven customer relationships.

The overall situation is that companies generally believe they are making considerable progress in analytics- and AI-based customer relationships, but that progress is not often noticed by customers (see Figure 1.1). For example, a 2024 survey conducted by CDP vendor Twilio found that while 81% of medium to large global companies say they have a deep understanding of their customers, less than half (46%) of global consumers agree.¹⁵ Similarly, 84% of companies feel that they provide good or excellent customer engagement, but only 54% of consumers agree. We tend to side with the customers on this one and are only surprised that more than half are happy with the engagement level.

On the plus side, 64% of companies say they use AI to build a unified profile of every customer, and 70% say they're already using AI to personalize content and marketing offers. Among the companies that have adopted AI,

¹⁵ 2024 State of Customer Engagement Report, Twilio, <https://www.twilio.com/en-us/state-of-customer-engagement/report>.

45% have seen improved customer satisfaction scores, 41% report better data-driven decision-making, and 41% experienced improved market segmentation and targeting.

Again, however, consumers aren't quite as bullish on AI. While 85% of the company executives surveyed "strongly/moderately" agree that AI will enhance customer engagement, only 22% of consumers share that degree of optimism. This suggests that companies attempting to measure their progress in this domain should at least occasionally ask their customers what they think about it.

Among the challenges to customer science—using data, analytics, and AI to better understand and target offers to customers—company executives mentioned "responding to changing customer needs/preferences" (44%); "connecting customer data across channels and sources" (42%); "navigating compliance and data privacy standards" (41%); and "turning data into customer insights" (40%). In other words, despite some progress, many organizations still struggle with effective customer science.

A New Era of Customer Science

Although we've described an environment for customer science that has promised great advances but delivered only spotty progress, we're optimistic about the future for several reasons. As we've discussed, organizations have been putting the foundations in place for data and its effective management. There is still more progress needed in customer data, but let's say we are halfway there. Even with the data that's already in place, companies like Nestle are making considerable progress toward more effective marketing, online sales, and customer-oriented decisions.

Another factor that gives us hope for a brighter future is artificial intelligence. It has been present in customer contexts for several decades now, but there is increased maturity and capability in traditional "analytical" AI, and now of course the highly innovative and rapidly maturing generative AI. We believe that both types of AI have a bright future in enhancing customer engagement and relationships, and we'll describe their impact throughout this book.

We're also optimistic because companies are becoming more data-driven in their cultures and decision-making. In 2024, an annual survey of data and analytics leaders in which Tom participates revealed a major change from

the previous year. The percentage of company executives (primarily chief data/analytics/AI officers in large companies) reporting their organizations had “established a data and analytics culture” increased from 21 to 43% over the previous year. The organizations that had “created a data-driven organization” doubled from 24% in 2023’s survey to 48% in 2024’s. This was far more improvement than at any time over the 12-year history of the survey; the percentages had languished at the same level or even decreased over the previous few years. Tom and his survey colleague, Randy Bean, argued that the change was primarily driven by generative AI.¹⁶

However, in the thirteenth version of the survey (labeled 2025), the numbers fell back to earth a bit. In that survey, 37% indicated that they work in a data- and AI-driven organization, and 33% said that organization has a data- and AI-driven culture.¹⁷ It appears that generative AI alone was not quite enough to create lasting cultural and organizational change. In fact, another question in the same survey asks what the principal barrier to becoming data-driven is in their companies. Ninety-two percent say that people/culture/change management issues are at fault; only 8% blame technology.

Of course, the key to success in science-driven customer relationships is to take a holistic approach across the organization. This transformation needs to encompass multiple business functions, channels, technologies, and customer-related activities if it is going to be successful. It needs to address organizational and cultural issues as much as technological ones. In the next chapter, we’ll describe two companies that have addressed all of these issues and made effective use of technology for customer relationships. One did it for a while but has fallen back a bit; the other is still going strong.

¹⁶ Thomas H. Davenport and Randy Bean, “Survey: Gen AI Is Making Companies More Data Oriented,” *Harvard Business Review*, 15 January 2024, <https://hbr.org/2024/01/survey-genai-is-making-companies-more-data-oriented>.

¹⁷ Data and AI Leadership Exchange, “2025 AI & Data Leadership Executive Benchmark Survey,” January 2025, <https://static1.squarespace.com/static/62adf3ca029a6808a6c5be30/t/67642c0d40b42a7d7e684f49/1734618125933/2025+AI+%26+Data+Leadership+Executive+Benchmark+Survey+120624.pdf>.