Hyper-Personalization for Customer Engagement with Artificial Intelligence

Thomas H. Davenport
Babson College and Oxford Business School

Personalization based on customer attributes and behavior is a familiar concept among marketers, and artificial intelligence is making it increasingly effective. AI-based hyper-personalization employs both sophisticated methods and far more data than previous methods and is far more precise as a result. Thomas H. Davenport discusses the role of AI in personalization as well as the growing backlash against personalization fueled by data privacy concerns.

Personalization is one of the most important ways marketers can use data, analytics, and artificial intelligence (AI) to increase customer engagement. It suggests to potential customers that the offered product or service is particularly suited to their specific needs and desires. It promises to value customers' attention and time, bringing to their attention only offerings consistent with their interests.

Personalization—also known as one-to-one (or 1:1) marketing—is not new. The concept has its roots in the idea of segmentation, in which marketers treat different groups differently. Personalization is ostensibly about appealing to specific individuals, but has traditionally relied upon placing people in groups such as gender, sociodemographic status, geographical residence, age, etc. Generally speaking, the more attributes marketers use to segment customers, the closer segmentation comes to personalization. However, until recently, few firms had the data or modeling ability to actually create a different offer for each consumer. Full personalization was more a concept than a reality in the first decade or two of its existence; in practice it was not distinguished from segmentation.

Both segmentation and one-to-one marketing were primarily focused on increasing sales rather than improving customer engagement with a product, brand, or company. Customer engagement generally refers to developing an ongoing relationship and a broader customer experience, rather than a particular sales transaction.[1] One popular measure of engagement proposed by consultants Bain & Co., the Net Promoter Score (NPS), has been tied to company-level shareholder return growth through repeat and referred sales revenues.[2] However, many academic studies have found little or no relationship between NPS and consumer-level sales.[3] The idea that there is a relationship between personalization, engagement, and sales is thus rooted only in logic, rather than deep empirical results.

Over the last decade, with the rise of e-commerce and online marketing, personalization has been used to target digital advertising, offers, and other customer-oriented content. The vast amount of data involved in digital customer relationships, and advances in AI methods, make it increasingly possible to approach the ideal of a unique offer for each customer.

However, most approaches to personalization have not been very sophisticated or effective. They have usually fallen well short of the 1:1 ideal because they lacked accurate and detailed data for personalization, or because they used relatively primitive methodological approaches and algorithms. Today, when sufficient data is available, some methods can help marketers to realize the promise of true personalization, here termed *hyper-personalization*.

Many managers lack a good means of characterizing the technology and sophistication of alternative personalization approaches. Vendors tout their own approaches without mentioning limitations. Each approach has its own data requirements and technical underpinnings.

There are also privacy tradeoffs involved in personalization. The potential value of personalization is highly appealing to marketers and product or service designers, but we must address important issues concerning customer privacy and perceptions of invasiveness. Personalization based on data should therefore generally be conducted with transparency and customer permission.

What to Personalize?

Personalization is applicable to a wide variety of marketing approaches. Perhaps the most common is personalized advertising, either for digital or analog (direct mail, television, etc.) ads.

> The potential value of personalization is highly appealing to marketers and product or service designers, but we must address important issues concerning customer privacy and perceptions of invasiveness.

Offers and discounts on particular products and services can also be personalized. In online commerce, search rankings and other aspects of web and mobile sites can be personalized. These personalization approaches are typically intended to increase conversions, from initial click to final purchase, rather than overall engagement.

Marketers today can easily personalize digital content, whether in marketing messages, web pages, or other forms of expression. Many assume that it builds customer engagement more than transaction-oriented digital advertising. And personalized content has moved beyond business. Political campaigns are increasingly personalized, as are news feeds, which has given rise to concerns about the filter bubble effect, in which people only see content from sources with which they agree.

We can also personalize the attributes of products and services. Financial services firms personalize investment portfolios with AI advice bots. Clothing companies offer shoes and other goods which can be personalized at manufacture. Even medical companies increasingly offer services and treatments that are targeted to a patient's genome or other biological makeup, termed *precision medicine* or *personalized medicine*.

Many companies now offer personalization in behavioral recommendations. Startups employ *precision nudges* to encourage weight loss, good economic habits, and overall health and nutrition.[4] Because these personalized services benefit customers beyond encouraging them to purchase an individual product or service, they do seem to advance customer engagement.

Personalization has often been viewed as relating to offers—ads, promotions, new product announcements, and the like. A personalized *next best offer* is intended to motivate only the customer's next purchase. Firms aiming to increase customer engagement are more partial to the term *next best action*. AI systems designed for this broader concept may advise customers to purchase an additional product or service, but they may also recommend actions or content that advance the customer relationship without making a new sale.

Recommended next best actions may include content about how the customer can improve their experience with the product or service, how the product or service can be used most effectively, or simply how to live a better life. Whenever possible, managers should ensure that this advice be personalized not only with regard to the product or service purchased, but also according to other customer attributes that increase the likelihood of their relevance.

Morgan Stanley's wealth management team, for example, has created a next best action system that provides machine learning-based personalized investment recommendations to its customers (mediated through investment advisors, who send the recommendations to customers). Investment advisors who use the system significantly outperform those who do not in growing the assets under management, and have higher productivity and frequency of engagement with customers.[5]

Benefits of Personalization

Personalization has many potential benefits. Survey research attests to some of the reputed benefits, while others are suggested by online consumer behavior. Most involve increased sales or conversion rates rather than broader metrics of customer engagement. A McKinsey 2021 survey, for example, found that 71 percent of consumers say they expect businesses to recognize them and to personalize product or service offers to their interests. Seventy-six percent are frustrated by an absence of personalization. In addition, companies surveyed by McKinsey that say they personalize report higher levels of revenue growth than those who do not.[6] Other less recent surveys also show a strong consumer preference for personalization.

However, the landscape surrounding the tradeoff between personalization and privacy is changing. Two 2021 KPMG surveys of businesses and consumers depict the challenge companies face when collecting data on consumers, which is necessary for effective personalization. Among 250 business executives, 70 percent reported that they expanded their collection of personal consumer data during the previous year. However, 86 percent of consumers said they have growing concerns about data privacy, and 30 percent said there are no circumstances under which they would share data with businesses. Only 12 percent said they would share data to personalize ads, and 17 percent would do so to help companies improve their products and services.[7]

As customers become more sensitive to data privacy, they may be less interested in personalization, or at least may require a higher degree of accuracy in personalizing to justify the tradeoff. It remains difficult to establish the exact frontier at which the desirability of personalization is outweighed by the desire for personal privacy, in part because the two traits are hazy in the minds of many consumers.

Contemporary AI-Based Approaches to Personalization

Firms have employed AI, in the form of rule-based systems, for personalization for decades, but until recently it has been a relatively blunt instrument. The most sophisticated, precise, and difficult form of personalization—in other words, hyper-personalization—requires machine learning models. Unlike rule-based personalization, machine learning can employ multiple different customer, product, or contextual attributes with few complications, while tending to be much more precise than rule-based approaches. It allows a company to produce millions of unique offers, drawn from many different groupings or segments. Machine learning-based personalization can even approach the elusive 1:1 segmentation to which marketers have long aspired. This approach has been theoretically possible for several decades, but lacked the necessary data. For example, the UK-based supermarket chain Tesco pioneered hyper-personalization in the late 1990s and early 2000s using its Clubcard loyalty program data to generate 12 million unique offers for grocery promotions. More recently, in the inflationary economy of 2022, Tesco focused on personalizing discounts for Clubcard holders.[9]

Hyper-personalization approaches that use machine learning can be centered on the item, the customer, or a combination of the two. Item-centric approaches like collaborative filtering rely on the "people who bought this item also bought this other item" approach, generating coefficients of item relatedness based on purchase data. All a company needs to know to use this approach is what items customers have expressed interest in or bought.

Some companies go deeper into the item-centric approach by classifying the attributes of their products. Netflix's well-known recommendation engine, for example, classifies each movie and TV show by multiple product attributes, such as subject, stars, directors, and the like. It can then recommend content to customers which has the same attributes as content they've purchased.

In addition to information about titles, Netflix's personalized recommendations are based on self-reported customer entertainment preferences, viewing history, and ratings of watched titles, as well as other members with similar tastes and preferences (it categorizes more than 2000 different taste communities), and situational factors such as time of day, device type, and length of a typical viewing session. The engine uses the information to present a personalized set of titles to each viewer.[10] Netflix's personalization has clearly contributed to a better customer experience, and led to high subscriber growth for many years (though it has declined somewhat recently due to economic and post-pandemic contraction) and to higher combined usage than cable and satellite viewing combined.[11]

Customer-centric approaches to personalization might employ not only past purchases, but also customer demographic data, recent life events, estimated income levels, communications channel preferences, and responses to previous offers. They combine these variables to develop a predictive model of how a customer will respond to the personalized offer. They then use the model to score each potential customer in terms of the likelihood of purchasing a product or category. By deploying many different models, a company could approach making a different offer to each customer. The grocery store chain Kroger, for example, has 60 million customers in its loyalty programs and delivered 1.9 billion personalized offers to them in 2021, using a large-scale machine learning model developed by 84.51°, its data and analysis subsidiary.[12]

These offers are primarily focused on encouraging sales transactions rather than on customer engagement. Yet 84.51° and Kroger are also beginning to emphasize personalized nutritional information and recommendations that may boost engagement over time. The retailer's OptUP program uses Kroger loyalty card data to calculate a nutrition score from a customer's recent purchases. Shoppers can also browse an app while shopping to see nutrition scores of individual products and receive "Better for You" recommendations of similar but healthier products.

While companies like Kroger have voluminous data on shopper behavior, what usually makes customer-centric models difficult is obtaining the necessary data. To employ them, a company needs extensive data on customer attributes and labeled outcome data, such as whether customers purchased the product or category, or responded to an offer. Loyalty programs allow companies to track many aspects of customer behavior over time. In placing digital ads, companies have traditionally used cookies as an excellent source of data about websites that customers have visited, which can be used to predict customer interest in ads. Consumers are beginning to be wary of cookies, though, and some companies, like Google, are beginning to phase them out of their web browsers.

Third party data aggregators and brokers are increasingly drawing data from multiple sources to provide personalization attributes.[13] They might combine, for example, a consumer's web browsing history with credit card purchases, social media activity, email domain name, type of device used to access the Internet, and other characteristics. The widespread availability of these types of data make it easier to create sophisticated machine learning-based personalization models. Unfortunately, they will probably also hasten the backlash against personalization.

In addition to the attributes of customers and products or services, successful personalization models often include contextual factors about the transaction or offer—including the season, time of day, or specific customer location—that can influence the nature of the offer. Companies may also decide to make offers only on products that are in stock at a local store or through e-commerce purchases. Only traditional supervised or deep learning models allow for very granular segments using multiple contextual and offer-oriented features to make predictions, so many factors, indeed, that no human could keep them all straight.

In addition to the lack of adequate data, another traditional constraint on marketers seeking to personalize for greater customer engagement is now being eased. AI companies have begun to automate the creation of machine learning models (called automated machine learning, or AutoML) and ongoing maintenance (machine learning operations, or MLOps), making the benefits of machine learning more accessible to non-data scientists. Business or marketing analysts who have some quantitative orientation and who understand customers and markets are, in many cases, now able to create personalization and other types of models using machine learning. They can also ensure that the models do not drift, but continue to effectively predict customer behavior over time using MLOps systems. At Kroger and 84.51°, for example, insights specialists work alongside professional data scientists to create machine learning models, and their greater business insights sometimes make their models more useful than those of data scientists.[14]

Personalization Model Types

Companies can use many machine learning models for personalization, see Figure 1. The most sophisticated AI companies typically combine multiple types for different circumstances. The most common type is supervised learning models, in which systems are trained on labeled data and then used to make predictions.

One example of a labeled outcome is whether or not the customer bought a particular type of product. The model uses that to make a prediction of how likely the customer is to buy it in the future and what factors are statistically associated with that buying behavior. Companies have used traditional supervised ML models, without deep layers, for many years in personalization. Among the companies using supervised learning, among other methods, for personalization are Disney Parks and Resorts (using Magic Band data for personalized itineraries), Netflix, Nike, and Instagram.

Some companies have begun to employ multi-layer or deep learning neural network models for personalization.[15] Among the early adopters of this approach was Dynamic Yield, which was acquired by McDonald's to support its personalization efforts (and those of other companies as well) in 2019, and then sold to Mastercard in 2021. These models require more data and are less readily interpreted than traditional ML models, but can often supply more accurate predictions. In addition to McDonald's, companies using the Dynamic Yield personalization technology include Lands End, PacSun, Sephora, and Forever 21. Netflix has developed some of its own approaches and models using deep learning for personalization.

Unsupervised machine learning models can also be used to identify segments or clusters of like customers on several features or variables, and can thus also support personalization.[16] Segments, however, are not a highly precise form of personalization. These models also require extensive data and don't necessarily provide higher value than traditional analytical approaches to data-based segmentation, such as factor analysis or K-means clustering. There are many discussions of unsupervised learning for customer segmentation in the academic literature, but these methods are not widely used in business personalization.

In order to employ multiple models and choose between them, some leading companies are using reinforcement learning, sometimes in combination with deep learning, to find the models best at optimizing longer-term rewards such as a set of clicks or conversions over time.[17] AI systems can automatically evaluate and compare different reinforcement learning models. Netflix, for example, uses reinforcement learning to optimize customers' long-term satisfaction with its entertainment content.

Finally, many companies combine experimental results with other models. Using A/B and multivariate testing approaches helps them to understand customer preferences, particularly for personalizing online content. The outcomes of such experiments are causal rather than correlational, which often provides marketers with a higher degree of certainty about outcomes. Stitch Fix, for example, has built a centralized experimentation platform which uses a variety of experimental

Figure 1
Types of AI Models Used in Personalization

Types of Model	Definition	Strengths	Weaknesses
Supervised learning	Training models with labeled outcomes	Prediction	Labeled data requirement
Unsupervised learning	Requires little or no model training	Clustering like customers	Segmentation only; not better than other methods
Deep learning	Supervised multi-layer neural network models	Adds precision to some predictions	High volume of data required; low interpretability
Reinforcement learning	Maximizes longer-term objective with rewards	Prediction over time	Requires data on sequence of actions
A/B and multivariate testing	Compares two or more online alternatives	Causal outcomes	Becomes complex with multiple arms and when combined w/ML

designs to optimize the effectiveness of clothing recommendations.[18]

An Example of Successful Hyper-Personalization

One highly successful hyper-personalization is Starbucks' AI platform, Deep Brew. Starbucks had historically relied on baristas in physical stores to build customer engagement, but this became increasingly difficult as customers began to order in advance through the company's smartphone app, and was impossible when stores were closed during the COVID pandemic except for drive-through orders or drink pickup. The company launched Deep Brew, which was initially focused on English-speaking markets in the US, Canada, and UK, in 2019. It included not only personalization functions for the Starbucks app, but also assisted workers with administrative activities in the stores, such as reorder points for supplies and labor scheduling.

Deep Brew was Starbucks' first major foray into machine learning, and the company formed a small dedicated team of data scientists to prototype the recommendation engine and other models. Starbucks had massive amounts of data on past customer purchase patterns, which it applied to emailed promotions, in-app featured products and discounts, and games and prizes.

The team concluded that deep reinforcement learning would create the most powerful models for recommendations. This method tests alternative models against each other using an A/B testing approach. In this case, it optimized the criteria of total revenue from a sale and the likelihood of the customer buying additional items beyond their normal purchase. Their goal was to personalize across all touchpoints and channels. Starbucks now uses these models to create more than 10 billion hyper-personalized recommendations a year. The models learn rapidly and continuously from new data.

Now data and customer behavior-based recommendations include likely customer preferences like vegetarian food, price sensitivity, tea vs. coffee, baked good preferences, and more. Recommendations also personalize the drive-thru experience, using not the customer identity, but instead contextual factors such as location, time of day, and weather. During the pandemic, when drive-thru lanes were the only way customers could get Starbucks products from stores, data scientists added a feature assessing the length of the drive-thru line. When lines were uncomfortably long, the drive-thru screen recommended drinks that were easy and quick to prepare.

Within Starbucks, Deep Brew is viewed as successful not only for personalization but for its other administrative functions, and as a driver of engagement with customers who increasingly order through the mobile app; about a quarter of all orders now go through that channel. Kevin Johnson, Starbucks' CEO in 2021, credited Deep Brew with boosting Starbucks' same-store sales, mobile app sales, and drive-through sales during the COVID pandemic period.[19]

The Deep Brew AI project wasn't easy to accomplish, though. Some managers, used to an intuitive and consensus-based decision culture, were not initially receptive to personalization based on opaque machine learning models. Knowing this, the data science team used agile development and reviewed prototypes frequently, carefully tracking the costs and benefits of adopting each model. They maintained a dashboard of key performance indicators during the development of the new model and process behind Deep Brew. They also focus closely on preventing a security breach or hack. Senior executives calmed Starbucks front-line associates by assuring them that Deep Brew would not replace human workers, but would instead free them up to develop closer personal ties to customers.[20]

These important management tasks of persuading stakeholders, ensuring security, and demonstrating the method's value were perhaps more difficult than the actual modeling. AI managers at Starbucks also emphasize that the model development process was only a small part of the technology development for Deep Brew. The systems or processes necessary to surround Deep Brew's ML models included configuration, data collection, feature extraction, data verification, machine resource management, analysis tools, process management tools, serving infrastructure, and performance monitoring. Together the company had to put much more time, effort, and investment into these infrastructure activities than into the modeling itself.

> While most consumers still seem to appreciate some degree of personalization, they may react negatively to offers that seem to be too personal, based on attributes or activities they view as private.

How Much Personalization is Too Much?

The rapid growth in digital media and tools for personalizing ads and offers has prompted a growing consumer concern about privacy, although it has not yet had a substantial effect on personalization or its related customer engagement in the United States. Some have criticized common approaches used for personalization, including capturing and analyzing customer online behaviors and purchases, as amounting to surveillance capitalism.[21] While most consumers still seem to appreciate some degree of personalization, they may react negatively to offers that seem to be too personal, based on attributes or activities they view as private. Or they may require a higher level of personalization, more accurately attuned to their needs and desires, in compensation for the amount of privacy they believe they are giving up.

We don't really yet know the limits of personalization. For better or worse, surveillance capitalism is in its early stages, and has until recently been hindered by poor quality data, insufficient data, challenges in establishing a persistent customer identity, and lack of methodological sophistication among marketers. Most consumers probably do not view advertising and marketing offers as attacks on their privacy if they are closely targeted with the products and services they really desire.

In addition, individual consumers will have a different sense of the appropriate balance between privacy and personalization. Some will perceive any loss of privacy as a fair value exchange for more efficient and enjoyable shopping. Others will be placated by companies being transparent about consumer data, and even explanations of why they may have seen a particular personalized ad or offer. Facebook, for example, has offered a Why Am I Seeing This Ad? feature since 2014, and in 2019 added further details about personalization approaches, such as the attributes the advertiser was attempting to appeal to, and the source of the data used for personalization.[22] Still other consumers may want no data-based personalization at all.

Personalization of ads and news on mobile devices has been challenged over the last year by changes in vendor policies. Specifically, Apple announced that beginning in 2022 its device Identifier for Advertisers (IDFA) would require users to opt in. IDFA previously allowed apps to personalize ads, offers, and content based on third party app iPhone activity. Prominent app providers like Facebook and YouTube estimated that this app tracking transparency would cost them billions of dollars in advertising revenue.[23]

The ability to personalize marketing is also governed in part by regulatory constraints. The European General Data Protection Regulation (GDPR) does restrict companies' ability to use data for personalization, although its utility in this regard is limited both because individual consumers don't understand the full implications of their consent for personalization and because enforcement mechanisms are limited.[24] The California Consumer Privacy Act (CCPA), the state-level data privacy legislation that took effect in 2020, allows consumers to opt out of having their data sold to another user. However, CCPA has also had little impact on personalization of advertising thus far, in part because few consumers opt out.[25]

Personalization is continuously evolving in response to increasing levels of digitization and data, advancing methods for artificial intelligence, and the changing perceptions of customers and regulators about the tradeoffs between privacy and personalization. Astute marketers can make increasingly effective use of sophisticated personalization approaches, but they should be aware that the value of and reaction to any particular approach is likely to be limited in time. Because of the difficulties of defining and measuring customer engagement accurately, marketers should also address how personalization affects long-term customer behavior and measure multiple aspects of it. It seems likely that well-executed hyper-personalization, driven by artificial intelligence, can help marketers to preserve customers' time and attention amidst an overwhelming flow of information, increasing customer engagement over time. But this expectation is based on logic and certain limited and narrow empirical findings. We will only know more when both research and practice have been pursued over a longer time and with greater breadth than they have been thus far. ■

Author Bio

Tom H. Davenport is a Distinguished Professor of IT and Management at Babson College, a visiting professor at Oxford Business School, a fellow of MIT's Digital Economy Initiative, and a senior AI advisor to Deloitte. He has published twenty-three books and over 300 articles, consulting for leading companies on data, analytics, and AI strategies, organizational structures, and deployment. He is among the world's top twenty-five consultants, top three business/technology analysts, and 100 most influential people in IT.

Endnotes

1. Weng Marc Lim et al, "Past, Present, and Future of Customer Engagement," *Journal of Business Research* 140, Feb. 2022, https://doi.org/10.1016/j.jbusres.2021.11.014
2. The Net Promoter Score is described in Frederick F. Reichheld, "The One Number You Need to Grow," *Harvard Business Review,* December 2003; the newer research on "earned growth," is described in Frederick F. Reichheld, Darci Darnell, and Maureen Burns, "Net Promoter 3.0," *Harvard Business Review,* Nov.-Dec. 2021.
3. Sven Baehre et al, "The Use of Net Promoter Score (NPS) to Predict Sales Growth: Insights from an Empirical Presentation," *Journal of the Academy of Marketing Science* 50, 67–84 (2022). https://doi.org/10.1007/s11747-021-00790-2
4. Thomas H. Davenport, James Guszcza, and Greg Szwartz, "Using Behavioral Nudges to Treat Diabetes," *Harvard Business Review,* Oct. 10, 2018, https://hbr.org/2018/10/using-behavioral-nudges-to-treat-diabetes
5. Thomas H. Davenport, "The Future of Work Now: Morgan Stanley's Financial Advisors and the Next Best Action System," *Forbes,* May 16, 2020, https://www.forbes.com/sites/tomdavenport/2020/05/16/the-future-of-work-now-morgan-stanleys-financial-advisors-and-the-next-best-offer-system/?sh=7fd496fb7027
6. Nidhi Arora et al, "The Value of Getting Personalization Right—or Wrong—Is Multiplying," McKinsey & Co., 2021, https://www.mckinsey.com/capabilities/growth-marketing-and-sales/our-insights/the-value-of-getting-personalization-right-or-wrong-is-multiplying
7. "Corporate Data Responsibility," KPMG, August 2021
8. Thomas H. Davenport and Jeanne G. Harris, *Competing on Analytics* (Boston: Harvard Business Review Press, second edition, 2017), pp. 138-139.
9. James Davey, "Loyalty Pays Off for Tesco as Britons Feel the Pinch," Reuters, Feb. 23, 2022, https://www.reuters.com/business/retail-consumer/loyalty-pays-off-tesco-britons-feel-pinch-2022-02-23/
10. Netflix website, "How Netflix's Recommendation System Works," https://help.netflix.com/en/node/100639, accessed August 4, 2022.
11. "The Netflix Effect: How Netflix's Commitment to Extreme Personalization Is Impacting Customer Expectations," WBR Insights, no date, https://nextgencx.wbresearch.com/blog/netflix-effect-commitment-extreme-personalization-customer-expectations, accessed August 4, 2022.
12. The company is named after the longitude of Cincinnati, where it is based. See Thomas H. Davenport, "84.51 ° Builds a Machine Learning Machine for Kroger," *Forbes*, April 2, 2018, https://www.forbes.com/sites/tomdavenport/2018/04/02/84-51-builds-a-machine-learning-machine-for-kroger/?sh=3bd47a0c64e1
13. For an excellent description of the types of data available from third party providers and how it is used, see Stuart A. Thompson, "These Ads Think They Know You," *The New York Times*, April 30, 2019, https://www.nytimes.com/interactive/2019/04/30/opinion/privacy-targeted-advertising.html
14. Thomas H. Davenport, "The Future of Work Now: AutoML at 84.51° and Kroger," *Forbes*, Oct. 21, 2020, https://www.forbes.com/sites/tomdavenport/2020/10/21/the-future-of-work-now-automl-at-8451and-kroger/?sh=6839782f4b0e
15. Deep learning models are neural networks (themselves a form of machine learning) that use nodes based on features or variables to convert data inputs to outputs such as predictions; they are "deep" with several or many layers of intermediate nodes.
16. Unsupervised learning attempts to find clusters, patterns, or anomalies in a collection of data with unlabeled outcomes.
17. Reinforcement learning is a form of unsupervised learning in which a program attempts to maximize a reward in a particular sequence of situations over time. It can employ a variety of algorithm types. It is most commonly used in games, including autonomous chess, Go, and Atari video games.
18. Aaron Bradley, "Building Our Centralized Experimentation Platform," Stitch Fix website, July 30, 2019, https://multithreaded.stitchfix.com/blog/2019/07/30/building-centralized-experimental-platform/
19. Emma Liem Beckett, "Starbucks' Digital Success Partially Driven by AI Engine, CEO Kevin Johnson Says," RestaurantDive website, April 28, 2021, https://www.restaurantdive.com/news/starbucks-digital-success-partially-driven-by-ai-engine-ceo-kevin-johnson/599182/
20. Jennifer Warnick, "AI for Humanity: How Starbucks Plans to Use Technology to Nurture the Human Spirit," Starbucks website, Jan. 10, 2020, https://stories.starbucks.com/stories/2020/how-starbucks-plans-to-use-technology-to-nurture-the-human-spirit/
21. Shoshana Zuboff, *The Age of Surveillance Capitalism* (Public Affairs, 2020).
22. Iris Hearn, "Facebook Expands "Why Am I Seeing This Ad" Feature to Provide More Detailed Insights to Users," *Impact*, July 15, 2019, https://www.impactplus.com/blog/facebook-expands-why-am-i-seeing-this-ad-feature-to-provide-more-detailed-insights-to-users
23. Matthew Fox, "$315 Billion in Market Value Erased After Apple IDFA Privacy Changes," Business Insider Markets, Feb. 3, 2022, https://markets.businessinsider.com/news/stocks/facebook-meta-stock-apple-idfa-ios-privacy-change-social-media-2022-2
24. Michèle Finck, "The Limits of the GDPR in the Personalisation Context" in U. Kohl, J. Eisler (eds), *Data-Driven Personalisation in Markets, Politics and Law* (Cambridge University Press, 2021) https://ssrn.com/abstract=3830304
25. Kate Kaye, "California's privacy law has had 'no impact' on ad revenues or inventory, but indirect effects could hurt." Digiday website, Feb. 25, 2021, https://digiday.com/media/ccpa-early-impact/