Artificial Intelligence & Machine Learning: Emerging Legal and Self-Regulatory Considerations

A Report by the American Bar Association's Section of Antitrust

Part One

September 30, 2019

Foreword

The charge — assess the implications of data analytics, and new types of artificial intelligence, including machine learning and neural networks, on the law that drives our practices as consumer protection and competition lawyers. The call to participate was made in the fall of 2018 and almost immediately interested parties signed up. Not only did these individuals agree to commit substantial time over a ten-month period, many joined the Section of Antitrust just so they could participate on the Task Force project.

Every Task Force member contributed substantially to this Report, and the nature and quality of participation from every member was especially welcome. Having those individuals listed on the next page at the same table, working constructively to understand the evolving market for information and the applicable law, lent an unexpected energy to the entire undertaking. Clearly, all involved understood there was important work to complete here.

At the outset, the Task Force established its objectives:

- Identify industry structure and key participants in the creation, collection, sale, and use of big data in the U.S. economy, and define current and expected scope of practices;
- Frame the legal and ethical challenges related to artificial intelligence and machine learning through algorithmic decision-making, with emphasis on developments following the White House (2014) and FTC reports (2016); and
- Consider how consumer protection and competition law has been applied or likely will apply to artificial intelligence and machine learning, including private rights of action.

This report – *Artificial Intelligence & Machine Learning: Emerging Legal and Self-Regulatory Considerations* – is part one of a two-part project. The focus here is primarily on consumer protection implications. We look forward to part two in 2020, when the Task Force will consider more fully the potential impact of artificial intelligence on market competition.

On behalf of the American Bar Association Section of Antitrust, we sincerely thank John Villafranco (Chair); Vice Chairs Patrick Bernhardt, Aaron Burstein, Vandy Howell, and Janis Kestenbaum; and the Task Force members for an outstanding effort. Special thanks to Lead Editor Matt Sullivan and Privacy & Information Security Committee members who spent countless hours helping to disassemble this report and put it back together. Finally, we thank all of the companies, associations, and law firms that committed so many hours to this project.

Deb Garza
Chair, Section of Antitrust (2018-19)

Task Force Leadership

John E. Villafranco, Chair - *Kelley Drye & Warren LLP*
Patrick Bernhardt, Vice Chair - *Capital One*
Aaron Burstein, Vice Chair - *Wilkinson Barker Knauer LLP*
Vandy Howell, Vice Chair - *Cornerstone Research*
Janis Kestenbaum, Vice Chair - *Perkins Coie LLP*
Ben Rossen, Vice Chair - *The Federal Trade Commission*
Matthew Sullivan, Lead Editor - *Kelley Drye & Warren LLP*

Task Force Members

Imran Ahmad - *Blake, Cassels & Graydon LLP*
Vildan Altuglu - *Cornerstone Research*
Katherine Armstrong - *Drinker Biddle & Reath LLP*
Katherine Barbacki - *Blake, Cassels & Graydon LLP*
Deon Woods Bell *–Federal Trade Commission*
Ken Dai - *Dentons Shanghai*
Sean Flaim – *U.S. Department of Health and Human Services*
Christopher Fonzone - *Sidley & Austin LLP*
Aryeh Friedman - *Dun & Bradstreet*
Svetlana Gans - *NCTA*
David Golden - *Constantine Cannon LLP*
Kate Heinzelman - *Sidley & Austin LLP*
Avigail Kifer - *Cornerstone Research*
Arnd Klein - *Cornerstone Research*
Gabrielle Kohlmeier - *Verizon*
Rebecca Kuehn - *Hudson Cook, LLP*
Gita Khun Jush - *Cornerstone Research*
Richard Lawson - *Manatt, Phelps & Phillips, LLP*
Jay Levine - *Porter Wright Morris & Arthur LLP*
Sheng Li - *NERA Economic Consulting*
Lydia Lichlyter - *Wilmer Cutler Pickering Hale and Dorr LLP*
Michelle Machado - *Mattos Filho*
Rory Macmillan - *Macmillan Keck*
Maria Maher - *Cornerstone Research*
Elizabeth Mendoza - *Perkins Coie LLP*
Lauren Myers - *Kelley Drye & Warren LLP*
Amadeu Ribeiro - *Mattos Filho*
Thomas Ritter - *Thompson Burton PLLC*
Brian Scarpelli - *ACT*
Randi Singer - *Weil, Gotshal & Manges LLP*
Kim VanWinkle - *Texas Attorney General's Office*
Suzanne Wachsstock - *Walmart*
Shaundra Watson
Aubrey Wesser - *Verizon*
Claire Chunying Xie - *NERA Economic Consulting*

The contributions to this report were made by the individuals listed above and should not be attributed to the law firms, companies, agencies, and consulting firms that are listed as affiliations.

The Antitrust Law Section is aware that other sections of the ABA, including the Science and Technology Law Section (**https://www.americanbar.org/groups/science_technology/**) and Business Law Section (**https://www.americanbar.org/groups/business_law/**), have written materials and other projects related to the subject matter of this paper. Those materials may supplement, expand, or conflict with the contents of this report. We encourage you to review their work on this important subject.

TABLE OF CONTENTS

Page

PART I: BIG DATA – AN OVERVIEW

A. Introduction and Scope of the Project

Over the last generation, technological advances have dramatically increased the amount of information we produce and our ability to store and analyze that information. Disparate mechanical devices have become the connected, data-generating computers that constitute the Internet of Things. The cloud computing revolution has cratered the cost of digital information storage. And the inexorable progress of Moore's Law — the principle that the speed and capability of computers can be expected to double every two years, as a result of increases in the number of transistors a microchip can contain — has placed computing power that was once the sole province of nation-states into the hands of billions of smartphone users worldwide.

Put simply, the volume of data generated every day is massive. During *each minute* of 2018 there were 3,877,140 Google searches, Amazon shipped 1,111 packages, Venmo processed $68,493 in peer-to-peer transactions, the Weather Channel received 18,055,555 forecast requests, YouTube users viewed 4,333,560 videos, Uber users took 1,389 rides, and Americans used 3,138,420 Gigabytes of internet data.[1] The volume of available data and the computational power to analyze this data have led to technological advances in artificial intelligence including the development of machine learning, and deep learning such as neural networks. These advances have greatly expanded the analytic insights we can distill from this data. In short, in recent years, our sense of what data is, and what we can achieve through data analysis has developed significantly.

Colloquially, these intersecting trends constitute the concept of "Big Data." This catch phrase has no precise meaning or definition, and what it may evoke likely differs "depending on

[1] *Data Never Sleeps 6.0*, DOMO, https://www.domo.com/learn/data-never-sleeps-6#/.

whether you are a computer scientist, a financial analyst, or an entrepreneur pitching an idea to a venture capitalist."[2] Generally, the Four V's describe big data:

- Volume: Digital datasets can be quite large in terms of scale and size, which has implications for both the use of the data as well as the storage and analytical requirements to analyze it.
- Velocity: The rate or frequency at which data is being generated and collected is unprecedented and increasing.
- Variety: Numerous and increasingly different types and forms of collected information can be used together in new and novel ways to draw inferences.
- Value: As the amount of data generated and our ability to store and analyze it grows, the degree to which the data allows valuable insights or analysis also increases.[3]

Fueled by these Four V's, the combination of big data infrastructure and analytic capabilities is creating increasingly valuable opportunities for firms, individuals, and societies. Indeed, the possibilities appear almost endless, including the ability to (1) better target individuals and connect them to relevant services, products, or opportunities (*e.g.*, jobs or schools); (2) anticipate people's behavior to improve healthcare, health habits, or crime prevention; (3) improve efficiency in the delivery of goods, traffic flow, and connecting people; and (4) achieve previously unattainable scientific advances, product designs, and problem solving.

Despite these seemingly limitless possibilities, big data also raises a host of legal, ethical, social, and moral challenges, ranging from the relatively familiar to the cutting edge of philosophical insight. For example, despite the general social knowledge of the rise of data

[2] EXEC. OFFICE OF THE PRESIDENT, BIG DATA: SEIZING OPPORTUNITIES, PRESERVING VALUES (2014), *available at* https://obamawhitehouse.archives.gov/sites/default/files/docs/big_data_privacy_report_may_1_2014.pdf [hereinafter WHITE HOUSE BIG DATA REPORT].

[3] *See* NAT'L INST. OF STANDARDS AND TECH., *NIST Big Data Interoperability Framework: Volume 1, Definitions*, NIST SPECIAL PUBLICATION 1500-1R1 (2018), *available at* https://bigdatawg.nist.gov/_uploadfiles/NIST.SP.1500-1r1.pdf; *see also The Five V's of Big Data*, BBVA (May 8, 2017), https://www.bbva.com/en/five-vs-big-data/; *The Five V's of Big Data,* IBM (Sept. 17, 2016), https://www.ibm.com/blogs/watson-health/the-5-vs-of-big-data/; Jason Williamson, Getting a Big Data Job 29–31 (2015).

analytics, consumers nonetheless may be unaware of the full extent to which they are generating data, the scope of sources from which the data is collected, how the data is used, how decisions derived from the data may impact them, and other privacy-related considerations. An individual may willingly provide certain information (e.g., via shopping forms or inputs on social networks), yet may be wholly unaware of other types of data collection (e.g., license plate scans, email scans, or signal trackers). Further, these individuals may not have visibility into how entities use or combine the collected information with other data to infer information about them or others. At the other end of the spectrum, the likes of Elon Musk and Stephen Hawking have argued that we must develop general principles now governing the development of big-data-driven artificial intelligence, lest it cause massive economic dislocation or produce a loss of human control that could destroy human civilization.[4]

In the end, the value society extracts from big data will depend directly on the ability to address these legal, ethical, social, and moral challenges produced by the collection, storage, and analysis of massive amounts of data. This is likely to be a multi-generational challenge, and a full discussion of the issues big data raises is well beyond the scope of this report. Significant issues such as the "long-tail" concerns associated with maintaining human control over artificial intelligence, the national security implications of autonomous technologies, how increased predictive capabilities may affect law enforcement techniques, and others, are best left for future works.

Rather, this report endeavors to provide a basic introduction to the use of big data and the most relevant legal issues it currently raises, with a focus on its consumer applications and the

[4] *See* James Vincent, *Elon Musk and Scientists Agree: We Need to Make Sure AI Helps Humanity*, THE VERGE (Jan. 12, 2015, 10:33 AM), *available at* https://www.theverge.com/2015/1/12/7531269/scientists-sign-open-letter-focusing-on-AI-dangers-benefits.

legal regimes governing how the corporate use of data analytics affects consumers. To this end, this report – Part One of a two-part undertaking – begins with a brief discussion of the structure of the big data industry, focusing on its key players and technologies. The report then provides an assessment of the opportunities and challenges of using data analytics, with a focus on commercial applications. An overview of how some of the most relevant consumer and data protection legal regimes might apply to data analytics comes next, followed by a discussion of the ways industry is attempting to self-regulate its use of data analytics. The report then concludes with a short set of basic, practical measures that might begin to address the legal and ethical concerns raised by big data and data analytics.

B. Industry Structure and Players

A multitude of different types of information fuel the modern data economy. Electronic devices of all kinds; individuals from every walk of life and every corner of the earth; entities of every type, whether the largest corporations or the smallest non-profits; every government in the world—all are continually generating data. These entities and devices are doing so through activities such as the following:

- Interactions with websites, applications, service providers, appliances, cameras/satellites, video cameras, microphones, and social media accounts;
- The processing of historic and current written and recorded content, including email, news, social media posts, business documents, literature, and reviews;
- Tracking operations, including manufacturing, testing, inventory, and business activities; and
- Scientific research or studies, including traffic, weather, and health patterns.

This variety of activities produces a seemingly infinite amount of data types and categories. Typical categories of data include financial records, geographic locations, demographic information, personal information and preferences, consumer activities and purchases, documents, photographs and videos, patterns of individual behaviors and choices, personal networks and relationships, and voice recordings.

Given this diversity, it is unsurprising that there are numerous market participants in the collection, storage, combining, selling, and analysis of big data. Thus, to the uninitiated, understanding the market can be difficult. While a comprehensive description of all forms of data and data analytics is beyond the scope of this report, we lay out below some of the lenses through which market commentators and participants view the commercial data market, in the hope that they will provide at least a baseline of knowledge helpful to the legal and regulatory discussions that follow. These include how the data is collected and shared, what infrastructure is used to derive value from the data, what sort of data is at issue, and how data analytics are being used to extract useful inferences from the data.

C. Collection and Sharing

One of the most prominent taxonomies of the big data economy is to characterize information based on how close it is to its source or point of origination, such as data about an individual (e.g., customers or product users) collected directly from such individual. This taxonomy is helpful because the obligations of parties using or processing data, and the entities that support such efforts, can turn on the nature of their relationship with the data and/or the underlying data subject. Different taxonomies can be used to clarify which aspects of datasets can be sensitive. Further, while the gathering of some datasets may be innocuous by itself, combining different datasets could become more problematic.

1. First-Party Data Generation and Collection

Virtually every organization is now generating first-party data, whether intentionally or not. First-party data includes information that a company collects directly from its customers and users, its processes and finances, as well as any information digitized in the normal course of business. The collection of data about customers and users can occur in myriad ways, including the following:

- When customers or users themselves provide it, such as by actively submitting it by filling out a form or a survey;
- Generated through a customer's transactions with the company, including the type and frequency of product purchases;
- Gathered from customers' interactions with the company via in-store, app, or website monitoring;
- Observed from the customers' interactions with the company's products and services, such as Internet of Things ("IoT") sensors, fitness trackers, or vehicle dashboards; and
- Generated through customer preference testing and purchase habits, *e.g.*, identifying which types of customers are most "price sensitive" by repeatedly offering different prices to different groups of customers.

Over time, these forms of collection have evolved. For example, data about a user's interaction with a website has expanded beyond the pages the user visits, and how long the user spends on each page, and where the user goes next online, to now include "heatmaps," or graphical representations that indicate the parts of a webpage that users are most likely to engage with their mouse.[5] Data about consumers' in-store behavior is no longer limited to purchasing records; companies now use in-store surveillance cameras in combination with facial recognition technology[6] and smartphone signals or beacons to track users' behaviors, interactions, and locations within a store.[7]

[5] *See The Complete Guide to Website Heat Maps*, HOTJAR (July 11, 2019), https://www.hotjar.com/heatmaps.

[6] Molly St. Louis, *How Facial Recognition is Shaping the Future of Marketing Innovation*, INC. (Feb. 16, 2017), https://www.inc.com/molly-reynolds/how-facial-recognition-is-shaping-the-future-of-marketing-innovation.html.

[7] Stephanie Clifford & Quentin Hardy, *Attention, Shoppers: Store Is Tracking Your Cell*, N.Y. TIMES (July 14, 2013), https://www.nytimes.com/2013/07/15/business/attention-shopper-stores-are-tracking-your-cell.html?pagewanted=all&_r=0.

Companies sometimes also generate first-party data by conducting A/B testing on their user-base. For example, by slightly changing which customers are addressed by certain ads, companies can analyze what features increase their "click-through" and purchase rates, and how this varies by customer type; similarly, by slightly changing prices in a certain submarket, companies can better understand different customers' price sensitivity.

First-party data collection has also expanded as consumers adopt and interact with new technologies, particularly since organizations can monitor and track those interactions. For example, companies now offer home digital or online assistants that respond to voice commands.[8] These assistants provide those companies with access to consumer data that can provide rich insights into behavior within the home.[9] As smartphones and smart appliances evolve, and digital assistants become ubiquitous, the data these devices generate will increase in sophistication and potential value. For example, companies are collecting biometric information (such as fingerprints, voiceprints, and facial prints), along with biomarkers (such as typing patterns, physical movements or walking patterns, screen navigation patterns such as mouse movements and finger movements on touch-sensitive screens, and engagement patterns) to identify users and collect data.[10]

At the same time, new types of devices and applications are collecting new types of data about individuals. The growth and adoption of IoT technology is also generating more consumer

[8] *See, e.g.*, Darrell Etherington, *Amazon Echo is a $199 Connected Speaker Packing an Always-On Siri-Style Assistant*, TECHCRUNCH (Nov. 6, 2014), https://techcrunch.com/2014/11/06/amazon-echo/.

[9] It was recently reported that Amazon employs workers to listen to, categorize, or transcribe voice queries users make to Amazon's Echo device. Matt Day, Giles Turner & Natalia Drozdiak, *Thousands of Amazon Workers Listen to Alexa User's Conversations*, TIME (Apr. 11, 2019), http://time.com/5568815/amazon-workers-listen-to-alexa/.

[10] Maria Korolov, *What is Biometrics? and Why Collecting Biometric Data is Risky*, CSO ONLINE (Feb. 12, 2019), https://www.csoonline.com/article/3339565/what-is-biometrics-and-why-collecting-biometric-data-is-risky.html.

data. For example, within the growing wearables industry,[11] IoT sensors within clothing, shoes,[12] jewelry, and even baby blankets[13] can track emotions, monitor anxiety levels in individuals, and assess posture and breathing to help avoid stress.[14] IoT also is a core feature in "smart city" initiatives across the country. For example, Sidewalk Labs, an Alphabet subsidiary, is building a smart micro-city that utilizes sensors and smart technologies to collect data to optimize everything from traffic, zoning, housing, gas emissions to health care.[15] IoT technology, and the resulting data collection is even making its way into consumables.[16] Such technological advancements have significantly expanded the volume and categories of first-party data accessible for big data purposes.

2. Second- and Third-Party Data Sharing

First-party data may be sold to or shared with known partners (sometimes referred to as second-party data), or collected or purchased by third parties, including high-volume data aggregators (third-party data). Data aggregators and data brokers may acquire data either through

[11] According to data from the International Data Corporation (IDC) Worldwide Quarterly Wearable Device Tracker. *See IDC Reports Strong Growth in the Worldwide Wearables Market, Led by Shipments of Smartwatches, Wrist Bands, and Ear-Worn Devices*, INT'L DATA COMPANY (Mar. 5, 2019), https://www.idc.com/getdoc.jsp?containerId=prUS44901819.

[12] *See* Dusan Johnson, *Smart Shoes: Tracking Fitness Through Your Feet*, GADGETS & WEARABLES (Apr. 1, 2019), https://gadgetsandwearables.com/2019/04/01/trackers-feet/.

[13] *See* Natasha Lomas, *Hugsy Snags $220k to Bring its Smart Baby Blanket to Market*, TECHCHRUNCH (June 19, 2017), https://techcrunch.com/2017/06/19/hugsy-snags-220k-to-bring-its-smart-baby-blanket-to-market/.

[14] *See Fitbits of the Future: What's Next for Biometric Data in Health?*, AETNA, https://www.aetnainternational.com/en/about-us/explore/future-health/fitbit-biometric-tech-health-care.html (Apr. 13, 2019).

[15] *See, Welcome to Sidewalk Toronto*, SIDEWALK TORONTO, https://sidewalktoronto.ca/; Sidney Fussell, *The City of the Future Is a Data-Collection Machine*, THE ATLANTIC (Nov. 21, 2018), https://www.theatlantic.com/technology/archive/2018/11/google-sidewalk-labs/575551/.

[16] *See* "FDA Approves Pill with Sensor that Digitally Tracks if Patients Have Ingested Their Medication," U.S. Food & Drug Admin., Nov. 13, 2017, https://www.fda.gov/newsevents/newsroom/pressannouncements/ucm584933.htm, accessed Apr. 13, 2019 (Data will be collected from the things we eat: the U.S. Food and Drug Administration has approved a drug used to treat schizophrenia and bipolar disorder that has an ingestible sensor embedded in the pill that records that the medication was taken.).

an arrangement with the company that collected it, or by entering into an arrangement to allow the third party to collect it directly from the company's customers. Alternatively, data aggregators may take available data from public-facing pages using a variety of means, such as bots, scrapers, or other technological measures.[17] Data aggregators may also combine data sets from disparate sources for resale. The ability to combine and analyze different data sources allows organizations to compile increasingly detailed portraits of consumers and their personal habits.

3. Storage, Infrastructure, and Access

Along with the proximity of the data to its origin, another key element of big data is the type of infrastructure used to derive value from the data. This includes, as briefly laid out below, high-capacity hardware specifically designed to store data in a secure and accessible format, the analytical software to extract meaningful insights from the data, and separate, specialized analytical hardware that can efficiently retrieve data from the storage hardware and execute the mathematical operations for the analytical software.

a. Storage Infrastructure

Big data storage and analysis infrastructure varies depending on numerous factors such as transparency and choice, cost, space, security, geographic localization, legal jurisdiction, and speed of access. Key factors, some of which are described in this report, include issues relating to consumer privacy (*e.g.*, the degree of consumer consent to the collection, use, and sharing of their data), data security (*e.g.*, how to effectively safeguard large data sets), as well as competition (*e.g.*, whether there is sufficient access to relevant data to allow competition and innovation).

[17] Some of these techniques have been and are the subject of the litigation. *See, e.g.*, Bradley Saacks, *Hedge funds are watching a key lawsuit involving LinkedIn to see if they can spend billions on web-scraped data*, BUS. INSIDER (Mar. 14, 2019, 9:48 AM) https://www.businessinsider.com/hedge-funds-watching-linkedin-lawsuit-on-web-scraped-data-2019-3.

Data infrastructure also must provide ready access to facilitate further processing of the data, either internally or through services provided by large technology vendors. Data storage and sourcing is one area that demonstrates the interplay between access and infrastructure.[18] When the relevant first-party and third-party data are in different formats, organizations need the ability to reach across these various data sets for the data to be actionable and understandable.

To manage data with different formats, some companies are creating "data lakes," or storage architectures where diverse data sets from disparate corporate silos can be maintained in their original format.[19] These big data analytics technologies allow for extraction of data from their original format, to permit meaningful analysis at scale from multiple sources.

Another form of infrastructure with particular relevance to big data is cloud computing. By using third-party cloud services, companies can forego large investments necessary build their own infrastructure. While this consolidation of resources enables clients to exploit valuable economies of scale, it potentially exposes large parts of the web to costly downtimes as numerous companies become dependent on the same infrastructure.

b. Analytical Software

The value of big data lies largely not in the data itself, but in the ability to use data analytics to make predictions or draw inferences from it. Yet, as the volume of data leveraged for big data purposes increases, so does the need for data governance, data integration, and analytical or

[18] For example, structured and unstructured data are generally stored in different forms. Structured data is typically highly organized, lending itself to immediate analysis and straightforward tabular storage (e.g., demographic information, purchase history), while unstructured data (*a*, text documents, videos) lacks an original inherent structure and typically cannot be analyzed without further processing and organization.

[19] *See* Jim Harris, *The Growing Importance of Big Data Quality*, SAS BLOGS, (Nov. 21, 2016), https://blogs.sas.com/content/datamanagement/2016/11/21/growing-import-big-data-quality/ (Harris defines data lakes as "the increasingly popular (and, arguably, increasingly necessary) storage repositories that hold a vast amount of raw data in its native format, including structured, semi-structured and unstructured data."); *Data Lakes and Analytics on AWS*, AMAZON WEB SERVS., https://aws.amazon.com/big-data/datalakes-and-analytics/ (Amazon Web Services offers data lakes along with their other products).

statistical methods to make sense of the data.[20] This demand for meaningful insights from compiled data sets has accelerated the need for sophisticated software.

Open source analytical software is one way to meet this demand. Unlike proprietary software, open source software generally provides access to source code, grants free redistribution rights, and allows the creation of derived works.[21] Once primarily the domain of non-profit organizations and loose collections of collaborative programmers, high-profile technology companies are now open-sourcing some of their most valuable software.[22] The reasons for pursuing an open-source model within a for-profit enterprise include the efficiencies with crowd-sourcing innovation, the ease of generating and coordinating software standards under transparency, the ability of unrelated entities to develop and release valuable extensions to in-house software, and the convenience of identifying and recruiting new employees that have already demonstrated familiarity with their software environment.[23]

4. Additional Data Taxonomies – Types and Source of Data

In addition to the first-, second-, and third-party modes of collection and disclosure and relevant types of infrastructure necessary to analyze big data, two additional taxonomies may apply

[20] Big data analytics is generally comprised of multiple categories, predictive analytics being the most prevalent one. Some of the other categories of big data analytics are text analytics, audio analytics, video analytics, and social media analytics. *See* Amir Gandomi & Murtaza Haider, *Beyond the Hype: Big Data Concepts, Methods, and Analytics*, 35 INT'L J. OF INFO. MGMT. 137, 140–143 (2015).

[21] *What Is Open Source Software*, THE LINUX FOUND., (Feb. 14, 2017), https://www.linuxfoundation.org/blog/2017/02/what-is-open-source-software/ ("Open Source Software … is software distributed under a license that meets certain criteria: (1) It is available in source code form without charge or at cost); (2) Open Source may be modified and redistributed without additional permission; [and] (3) [f]inally, other criteria may apply to its use and redistribution.").

[22] *See* Klint Finley, *Why 2018 Was a Breakout Year for Open Source Deals*, WIRED, (Dec. 23, 2018), *available at* https://www.wired.com/story/why-2018-breakout-year-open-source-deals/.

[23] *See Why Open Source?*, GOOGLE OPEN SOURCE, https://opensource.google.com/docs/why/.

when characterizing data and considering potential regulatory issues: (1) personal vs. non-personal information and (2) government vs. private data sources.

a. Personal vs. Non-Personal Information

The various legal regimes regulating data often focus on personal information, or various subcategories of personal information. The precise definition of personal information can vary significantly based on the relevant specific legal scheme(s), but, in general, these definitions include a wide range of information that may be linkable to an individual or device on its own or in combination with other information. For example, in some circumstances, personal information could include information about a particular individual compiled by collecting his or her social media postings or browser history. These data may contain sensitive personal information (*e.g.*, credit card numbers, details of medical treatment) or information from which it is possible to extrapolate sensitive insights (*e.g.*, rideshare information concerning a user's schedule). Individually identifiable information is often distinguished from anonymized data, though the distinction is not binary, and a significant amount of research has been devoted to methods of re-identifying data as well as assessing the risks of re-identification.[24]

b. Government vs. Private Data Sources

Similar to the private sector, government entities also may engage in the collection and generation of large amounts of data. However, the legal regimes governing government data can differ dramatically from those governing the private sector. While a full exploration of these differences is beyond the scope of this report, two key differences are worth noting.

First, Constitutional provisions that have significant implications concerning the collection and regulation of information—chiefly, the First and Fourth Amendments—generally only apply

[24] *See, e.g.*, Yves-Alexandre de Montjoye et al., *Unique in the Shopping Mall: On the Re-identifiability of Credit Card Metadata*, 347 SCI. 536 (2015).

to the government and not the private sector. Second, the fact that government-created information is, in some circumstances, a public good to which the government must provide free and open access means that governmental data is often treated differently than private sector data, typically protected as proprietary or a trade secret. For instance, Executive Order 13642 issued on May 9, 2013, and its accompanying Open Data Policy, expressly require that government information be "open and machine readable," in recognition that "[i]nformation is a valuable national resource and a strategic asset to the Federal Government, its partners, and the public."[25]

D. Insights provided by Data Analytics

1. Big Data Methodologies and Applications

Big data is "big" in two ways. First, it is "long" because it entails many observations (*e.g.*, millions of observations of an individual's choices regarding whether or not to buy a particular product after exposure to an ad). Second, it is "wide" because it contains many *different variables* for each observation (as one example, for a given individual, one may have the characteristics of different ads an individual has seen, the individual's past purchases, past searching behavior, demographic data, political data, location information, social media posts, etc.). Modern data analytic techniques are well suited for dealing with "wide" data because they can determine whether and how to use each new piece of data, or variable, to best predict the applicable outcome. Where traditional analytical techniques require that researchers select the relevant variables and model specifications, modern techniques can automatically select and synthesize important variables from very wide datasets to maximize accuracy.[26]

[25] Exec. Order No. 13,642, 78 Fed. Reg. 28,111 (May 14, 2013); Memorandum from the Office of Mgmt. and Budget to Heads of Exec. Depts. and Agencies (2013), *available at* https://project-open-data.cio.gov/policy-memo/.

[26] For example, researches used a dataset of 7.4 million Medicare beneficiaries, and more than 3,000 variables within Medicare claims, to predict mortality rates in the first 12 months after hip or knee replacement. Such a predictive model allows doctors to allocate scarce joint replacements to patients who are predicted to benefit

Modern data analytic methods span many approaches including, but not limited to, "artificial intelligence," "machine learning," and "neural networks," each of which can have distinct analytical applications and limitations. Various data analytics techniques can provide potential insights for a variety of purposes, such as pattern recognition, predictive analysis and causal inference.

Broadly, *pattern recognition* concerns the association of input data and a defined output concept. This process can pertain to concrete prediction tasks (e.g., weather forecasting) as well as situations without a defined, known output (e.g., the distribution of topics in a document database). These tasks can be narrow in focus (named entity recognition) or broad (machine translation).

Predictive analytics is a popular subset of pattern recognition. Here, a model aims to map historical input data to concrete, defined historical output data to serve predictions about future or unknown phenomena. These tasks can be temporal in nature (forecasting next quarter's sales) or non-temporal (identifying which product appears in customer images). Predictive analytics methods may or may not, by construction, identify a *causal inference*. In those that do, one can infer that the input data *causes* the value of the output data to transpire (e.g., more schooling causes higher wages on average). In those that do not, the input data is simply *correlated* with the output

most from the procedure. *See*, Jon Kleinberg et al., *Prediction Policy Problems*, 5 AM. ECON. REV.: PAPERS & PROCS. 105 (2015). In a more extreme case, image classification problems commonly involve datasets with hundreds of thousands of pixel values for each individual observation, yet successful analytical frameworks are able to highlight specific pixel groups that inform accurate predictions. *See* Alex Krizhevsky et al., *ImageNet Classification with Deep Convolutional Neural Networks*, ADVANCES IN NEURAL INFORMATION PROCESSING SYSTEMS (Fernando Pereira et al. eds., 25 ed. 2012). *But see, e.g.*, *Susan Athey on Machine Learning, Big Data, and Causation, Library of Economics and Liberty*, THE LIBR. OF ECON. AND LIBERTY (Sept. 12, 2016), http://www.econtalk.org/susan-athey-on-machine-learning-big-data-and-causation/ ("The problem with "data mining" [big data], or looking at lots of different models is that you might find, if you look hard enough you'll find the result you are looking for.").

data, but do not cause it (e.g., people who buy vanilla ice cream on average may also be more likely to buy yellow shirts).[27]

a. Artificial Intelligence

Artificial intelligence ("AI"), speaking broadly, is the overarching term for algorithmic-powered computer processes that learn to perform actions that correspond to and even surpass human abilities.[28] An AI process can include predictive models or decision-making algorithms that ensure that the entity "acts so as to achieve the best outcome or, when there is uncertainty, the best expected outcome."[29]

Primitive versions of "narrow" AI[30] could only "predict" the best outcome if they were programmed by humans through explicit instructions on how to respond.[31] For example, a program could play tic-tac-toe with a person using preset instructions for every possible move (*i.e.* predicting the best move in a given situation), or play checkers using a predetermined set of heuristics or algorithm.[32] While such a program is not necessarily deterministic, it cannot

[27] While not a traditional focus of artificial intelligence and machine learning, researchers in recent years have begun to apply machine learning models to conduct causal inference. The tasks focus on "causal inference for average treatment effects, optimal policy estimation, and estimation of the counterfactual effect of price changes in consumer choice models." *See* Susan Athey & Guido W. Imbens, *Machine Learning Methods Economists Should Know About*, 11 ANN. REV. ECON. (2019).

[28] There is no settled definition of artificial intelligence. In ARTIFICIAL INTELLIGENCE: A MODERN APPROACH (3d ed. 2010), Stuart Russel and Peter Norvig provide four approaches to the definition of artificial intelligence and discuss the limitations of each approach to fully define AI.

[29] *Id*. at 4. This definition uses the common "rational agent" understanding of AI. Multiple understandings/definitions of the field exist. *Id*. at 1–2.

[30] Artificial intelligence is characterized as being either "weak" or "narrow" AI and "strong" or "general" AI. Narrow AI is commonly developed today–it includes algorithms or models that, while complex, typically focus on a specific individual task. In contrast, strong or general AI, which remains relatively rare but is actively researched, is capable of addressing broad applications (and is more often the subject of science fiction). To date, it has not been possible to teach AI "common sense" or to program an AI system to make decisions about human ethical notions such as "fairness".

[31] *See* Arthur L. Samuel, *Some Studies in Machine Learning Using the Game of Checkers. II—Recent Progress*, In *Computer Games I* (Springer: New York, NY, 1988), pp. 366–400.

[32] *Id*.; David B. Fogel, *Using Evolutionary Programming to Create Neural Networks That are Capable of Playing Tic-Tac-Toe*, IEEE INT'L CONFERENCE ON NEURAL NETWORKS (1993).

systematically improve its performance unless reprogrammed with new strategies. More advanced AI applications are fundamentally different due to the concept of machine learning.

b. Machine Learning

Instead of relying solely on human instruction, some current AI programs incorporate machine learning to develop their algorithms. Machine learning occurs when a program can adapt in response to new observations. Using the example above, a checkers-playing program capable of machine learning would first "learn" by assessing a number of games (*e.g.*, the moves and outcomes), that is, "training" on datasets to discern and refine the best strategies for winning.[33] In another example, machine learning has facilitated automation of certain aspects of the legal discovery process, contributing to the rise of the e-discovery industry. In place of document discovery performed through exhaustive manual review, researchers have found that machine learning models trained using data generated by human reviewers can yield superior results as measured by recall and precision.[34] AI based on machine learning, once trained, can make determinations or decisions through algorithms that are driven by what has been learned by the data, rather than being dependent on programmed or preset inputs.

c. Neural Networks

Machine learning programs can learn through many different strategies.[35] One machine-learning model, often termed as a type of "deep learning" is the (artificial) neural network, which

[33] TOM MITCHELL, MACHINE LEARNING 2–3 (1997).

[34] *See*, *e.g.*, Maura R. Grossman & Gordon V. Cormack, *Technology-Assisted Review in E-Discovery Can Be More Effective and More Efficient than Exhaustive Manual Review*, 17 RICHMOND J. OF LAW & TECH. 1 (2011).

[35] These techniques address several different objectives and carry names like k-nearest neighbors, support vector machines, and latent Dirichlet allocation. Also consider supervised, unsupervised, and reinforcement learning as described in Ai Deng, *An Antitrust Lawyer's Guide to Machine Learning*, 32 ANTITRUST 82 (2018). *See also* Anita Banicevic et al., *Algorithms: Challenges and Opportunities for Antitrust Compliance*, 2018 A.B.A. SECTION OF ANTITRUST LAW: COMPLIANCE AND ETHICS SPOTLIGHT SPECIAL REPORT.

consists of a web of interconnected computer processing units that are inspired by biological neural or brain networks, and generally are not programmed with task-specific rules.[36] Neural networks are capable of consuming massive amounts of data both long and wide, as they autonomously discover complex relationships in data, without any explicit instructions from a researcher. For example, researchers are using neural network methods to process and analyze gigantic datasets used for cancer treatments, such as radiotherapy and chemotherapy.[37] Applying these methods, the researchers were able to develop standardized algorithms that improve tumor treatment for future patients.[38]

PART II: COSTS AND BENEFITS

A. Introduction

Big data, when combined with the capabilities of recent evolutions in data analytics, especially machine learning, promises substantial and wide-ranging benefits for consumers, businesses, other institutions, and society at large. At the same time, these technological advances raise legitimate concerns, including in relation to consumer privacy; the maintenance of competitive markets; equity and fairness; economic opportunity; the potential uses of such technologies for state surveillance; and the need to maintain human control over autonomous systems, among others.

[36] Regarding the term *neural network*, "[t]he study of artificial neural networks (ANNs) has been inspired in part by the observation that biological learning systems are built of very complex webs of interconnected neurons. In rough analogy, artificial neural networks are built out of a densely interconnected set of simple units, where each unit takes a number of real-valued inputs (possibly the outputs of other units) and produces a single real-valued output (which may become the input to many other units)." MITCHELL, *supra* note 33, at 82.

[37] *See* Hong Shen et al., *A Quantitative Quality Control Method of Big Data in Cancer Patients Using Artificial Neural Network*, IEEE 3RD INT'L CONFERENCE ON CLOUD COMPUTING AND INTELLIGENCE SYS. (2014).

[38] *Id*.

B. Data Analytics Opportunities and Challenges

This section identifies illustrative examples of the opportunities and challenges, and associated risks, with big data.

1. Opportunities

The ability to identify previously undetected patterns is at the heart of big data-triggered opportunities. Big data has already produced a number of key advances in how companies can serve their customers.

a. Customization and Convenience.

In addition to driving product personalization, machine learning and the expanding availability of consumer data have made it possible to deliver greater product convenience, adaptability, and ease of use. For example, a refrigerator with sensors and the ability to analyze data can recommend recipes based on the expiration of grocery ingredients in the refrigerator or can reorder needed items,[39] while voice-enabled virtual assistants can schedule appointments on a consumer's behalf,[40] or serve as virtual stylists or shoppers.[41]

b. More Effective Marketing

Data analytics (i.e., big data coupled with powerful analytics tools) enable businesses to target their marketing to consumers likely to be most interested in their goods and services, and to

[39] *See* Ry Crist, *LG InstaView ThinQ Alexa Fridge Adds Clever Kitchen Tricks*, CNET (Jan. 8, 2018, 7:31 AM), https://www.cnet.com/news/lg-instaview-thinq-alexa-fridge-clever-kitchen-tricks-ces-2018/.

[40] *See* Nick Statt, *AI is Google's Secret Weapon for Remaking its Oldest and Most Popular Apps*, THE VERGE (May 10, 2018, 2:00 PM), https://www.theverge.com/2018/5/10/17340004/google-ai-maps-news-secret-weapon-remaking-old-apps-products-io-2018.

[41] *See* Rob Smith, *Don't Know What to Wear? This AI Could Create an Outfit for You*, WORLD ECON. FORUM (Aug. 21, 2018), https://www.weforum.org/agenda/2018/08/this-ai-stylist-knows-what-you-like-to-wear-and-invents-new-outfits-for-you/.

measure the effectiveness of targeted advertisements and offers by recording "click-throughs" and purchases. This, in turn, can create a positive feedback loop to improve system performance.[42]

c. Production Efficiencies in Delivering Products and Services

Data analytics increases business process efficiency, while lowering costs, including manufacturing, production, and logistics processes. For example, a biopharmaceuticals manufacturer uses data analytics of its processes to increase its yield in vaccine production without additional capital expenditure.[43] Logistics businesses use data analytics to manage international networks efficiently and reduce shipping costs to consumers and businesses.[44] Agricultural scientists and engineers use machine learning to predict crop yields, detect crop disease before outbreaks, and recognize and eliminate weeds robotically.[45] By utilizing data on crop color, crop diameter, and textures, machine learning methods have predicted agricultural improvement with about 90 percent accuracy.[46] Similarly, financial institutions have been able to accurately detect fraudulent credit card activity in real-time by monitoring specific details about individual credit cards with 80 to 90 percent accuracy.[47]

[42] *See* Alison DeNisco-Rayome, *How Wayfair Used Big Data and Omnichannel Retail to Transform Shopping*, ZDNET (Sept. 1, 2017, 11:01 AM), https://www.zdnet.com/article/how-wayfair-used-big-data-and-omnichannel-retail-to-transform-shopping/.

[43] *See, e.g.*, Eric Auschitzky, et al., *How Big Data Can Improve Manufacturing*, MCKINSEY & CO. (July 2014), https://www.mckinsey.com/business-functions/operations/our-insights/how-big-data-can-improve-manufacturing.

[44] *See, e.g.*, Bernard Marr, *The Brilliant Ways UPS Uses Artificial Intelligence, Machine Learning And Big Data*, FORBES (June 15, 2018), https://www.forbes.com/sites/bernardmarr/2018/06/15/the-brilliant-ways-ups-uses-artificial-intelligence-machine-learning-and-big-data/; William B. Cassidy, *Tech-Savvy Truckers Search Big Data for Drivers*, JOC.COM (May 23, 2014, 3:59 PM), https://www.joc.com/trucking-logistics/labor/tech-savvy-truckers-search-big-data-drivers_20140523.html.

[45] *See* Konstantinos G. Liakos et al., *Machine Learning in Agriculture: A Review*, 18 SENSORS 2674 (Aug. 14, 2018).

[46] *Id.*

[47] *See* Siddhartha Bhattacharyya et al., *Data Mining for Credit Card Fraud: A Comparative Study*, 50 DECISION SUPPORT SYS. 602 (2011).

d. New Competition

As discussed below, accumulation of large quantities of data among some commercial players, and the value that flows from data analytics applied to it, is sometimes criticized as an impediment to competition, yet businesses can also leverage big data to develop products and to compete against incumbents by using it to identify market gaps or underserved market segments. As an example, new entrants to the mobile app space can license data on vehicle traffic patterns to, in turn, develop innovative routing technology that avoids traffic congestion or suggests alternative modes of travel such as public transit or ride-sharing services.[48] In doing so, they may obtain data from third parties, such as the traffic data from the services themselves, or from data brokers and other intermediaries. While such practices may raise privacy concerns, they also can enable smaller firms to develop products and services at increased scale and reduced cost.[49]

e. Medical Research and Patient Care.

Data analytics is enabling breakthroughs in health research and patient care. Researchers are using big data and machine learning for a range of purposes, including to analyze MRI scans to distinguish between Alzheimer's patients and those with milder forms of dementia;[50] and

[48] A number of services provide traffic data API access, such as MapQuest, TomTom, Bing, and ArcGIS. *See Traffic API*, MAPQUEST DEVELOPER, https://developer.mapquest.com/documentation/traffic-api/; *Traffic API*, TOMTOM FOR DEVELOPERS (May 16, 2019), https://developer.tomtom.com/traffic-api; *Traffic API*, MICROSOFT (Feb. 27, 2018), https://docs.microsoft.com/en-us/bingmaps/rest-services/traffic/; *Traffic Service*, ARCGIS FOR DEVELOPERS, https://developers.arcgis.com/rest/network/api-reference/traffic-service.htm.

[49] Data brokers provide access to vast stores of consumer data, and public sources of data are available from governments and nonprofit entities for free or at nominal cost. *See, e.g.*, *InfoBase®: The World's Most Powerful Consumer Insights*, ACXIOM, https://www.acxiom.com/what-we-do/infobase/ ("Comprehensive consumer data on approximately 250 million U.S. addressable consumers for powerful audience insights and targeting."). Likewise, cloud providers offer programmatic access to data analytics technologies that can allow startups to analyze large data sets and open-source software to build and test analytic models and theories are widely available. *See, e.g.*, *Machine Learning on AWS*, AWS, https://aws.amazon.com/machine-learning/; *AI and Machine Learning Products*, GOOGLE CLOUD, https://cloud.google.com/products/ai/; *Azure AI*, MICROSOFT AZURE, https://azure.microsoft.com/en-us/overview/ai-platform/.

[50] *See* Edd Gent, *Artificial Intelligence Could Help Catch Alzheimer's Early*, SCIENTIFIC AMERICAN (July 11, 2016), *available at* https://www.scientificamerican.com/article/artificial-intelligence-could-help-catch-alzheimer-s-early/.

develop more accurate information relating to heart attack risks,[51] cancerous tissue samples,[52] and tumor classification types.[53] Hospitals also are using machine learning to help prioritize patients for medical interventions based on predicted risk of complications, and healthcare providers are using it to predict the cost and quality-outcome of patients' medical procedures.[54]

f. Disaster Relief

Data analytics can help predict and address natural disasters. For example, the analysis of data from marine sensors that monitor waves and currents can help predict tsunamis.[55] Similarly, drone and street-level imagery fed to machine learning algorithms can detect buildings at risk of collapse in an earthquake, or can be used in combination with satellite imagery to identify impoverished and vulnerable neighborhoods to improve disaster risk management.[56]

g. Enhancing Human Capabilities

Data analytics may help augment human potential, enabling people to accomplish tasks that might otherwise be beyond their reach. Companies are utilizing data analytics to provide new products and services to disabled or impaired consumers. For example, businesses have developed

[51] *See* Eliza Strickland, *AI Predicts Heart Attacks and Strokes More Accurately Than Standard Doctor's Method*, IEEE SPECTRUM (May 1, 2017), https://spectrum.ieee.org/the-human-os/biomedical/diagnostics/ai-predicts-heart-attacks-more-accurately-than-standard-doctor-method.

[52] *See* Graham Templeton, *AI Beats Doctors at Visual Diagnosis, Observes Many Times More Lung Cancer Signals*, EXTREMETECH (Aug. 18, 2016), https://www.extremetech.com/extreme/233746-ai-beats-doctors-at-visual-diagnosis-observes-many-times-more-lung-cancer-signals.

[53] *See* Bing Zhang et al., *Radiomic Machine-Learning Classifiers for Prognostic Biomarkers of Advanced Nasopharyngeal Carcinoma*, 403 CANCER LETTERS 21 (Sept. 10, 2017).

[54] *See* Susan Athey, *Beyond Prediction: Using Big Data for Policy Problems*, SCI. MAG. (Feb. 3, 2017); Sherri Rose, *A Machine Learning Framework for Plan Payment Risk Adjustment*, HEALTH SERVS. RES. 21 (Dec. 2016).

[55] *See* BSA | THE SOFTWARE ALLIANCE, WHAT'S THE BIG DEAL WITH DATA? 11, *available at* https://data.bsa.org/wp-content/uploads/2015/12/bsadatastudy_en.pdf.

[56] *See* Giuseppe Molinario & Vivien Deparday, *Demystifying Machine Learning for Disaster Risk Management*, WORLD BANK DATA BLOGS (Mar. 6, 2019), https://blogs.worldbank.org/opendata/demystifying-machine-learning-disaster-risk-management.

mobile apps to assist the visually impaired in navigating the surrounding environment,[57] and speech recognition technology that uses algorithms trained on large volumes of data to convert speech to text, providing real-time translations for the hearing impaired.[58] Advancements also are extending human capabilities to traditional objects like the automobile. A suite of sensors, computer-vision technology, and data analysis software provides self-driving cars with the capacity to see and navigate the road,[59] which could significantly increase the mobility and independence of the elderly and disabled.[60]

h. Promoting Diversity and Inclusion

Though there are legitimate concerns that certain applications of data analytics will perpetuate inequalities, as discussed further below, it is also possible that data analytics can be used to help explain and reduce inequalities. For example, the National School Boards Association has used big data to identify factors indicating which students are at risk of not completing high school, some of which disproportionately affect African-Americans and help explain the

[57] *See, e.g.*, Steven Musil, *Google Developing Lookout App to Aid the Visually Impaired*, CNET (May 8, 2018, 5:17 PM), https://www.cnet.com/news/google-developing-lookout-app-to-aid-the-visually-impaired/; *Seeing AI*, MICROSOFT, https://www.microsoft.com/en-us/seeing-ai/.

[58] *See* Dieter Bohn, *Google Live Transcribe Could Be a Big Help for People Who are Deaf or Hard of Hearing*, THE VERGE (Feb. 4, 2019, 9:00 AM), https://www.theverge.com/2019/2/4/18209546/google-live-transcribe-sound-amplifier-accessibility-android-deaf-hard-hearing; John Roach, *AI Technology Helps Students Who Are Deaf Learn*, MICROSOFT (Apr. 5, 2018), https://blogs.microsoft.com/ai/ai-powered-captioning/.

[59] *See* Katie Burke, *How Does a Self-Driving Car See?,* NVIDIA (Apr. 15, 2019), https://blogs.nvidia.com/blog/2019/04/15/how-does-a-self-driving-car-see/.

[60] *See* Srikanth Saripalli, *Are Self-Driving Cars the Future of Mobility for Disabled People?*, SMITHSONIAN.COM (Oct. 11, 2017), https://www.smithsonianmag.com/innovation/are-self-driving-cars-future-mobility-disabled-people-180965241/.

discrepancy in graduation rates.[61] Companies also have leveraged data analytics to identify biases in their promotion processes and help address them.[62]

i. Expanded Economic Opportunity

Despite concerns of elevated unemployment due to automation,[63] data analytics can also help expand economic opportunity, such as through greater access to credit. Financial institutions are using AI with data from a wide variety of sources to determine the creditworthiness of consumers traditionally excluded from access to credit because they either had a "thin" credit file or lacked a credit file altogether.[64] This includes peer-to-peer ("P2P") lending through online platforms that leverage big data to match lenders with borrowers, including individuals or small or micro-businesses that traditional banking may overlook.[65] Some of these platform providers have advocated using their machine learning-based lending models as tools to reduce bias in lending decisions.[66]

[61] FUTURE OF PRIVACY FORUM & ANTI-DEFAMATION LEAGUE, BIG DATA: A TOOL FOR FIGHTING DISCRIMINATION AND EMPOWERING GROUPS 10 (2014), *available at* https://fpf.org/wp-content/uploads/Big-Data-A-Tool-for-Fighting-Discrimination-and-Empowering-Groups-FINAL.pdf.

[62] *See id.* at 2.

[63] *See, e.g.*, Darrell M. West, *Will Robots and AI Take Your Job? The Economic and Political Consequences of Automation*, BROOKINGS (Apr. 18, 2018), https://www.brookings.edu/blog/techtank/2018/04/18/will-robots-and-ai-take-your-job-the-economic-and-political-consequences-of-automation/.

[64] *See, e.g.*, Becky Yerak, *AI Helps Auto-Loan Company Handle Industry's Trickiest Turn*, WALL ST. J. (Jan. 3, 2019), https://www.wsj.com/articles/ai-helps-auto-loan-company-handle-industrys-trickiest-turn-11546516801; Penny Crosman, *BankMobile Deploys AI, Alternative Data to Lend to Thin-File Millennials*, AM. BANKER (Dec. 20, 2017, 3:50 PM), https://www.americanbanker.com/news/bankmobile-deploys-ai-alternative-data-to-lend-to-fico-poor-students.

[65] *See* Stijn Claessens et al., *Fintech credit markets around the world: size, drivers and policy issues*, BUS. Q. REV. (2018), *available at* https://www.bis.org/publ/qtrpdf/r_qt1809e.pdf.

[66] *See, e.g.*, Bruce Upbin, *ZAML Fair - Our New AI to Reduce Bias in Lending,* ZEST FINANCE (Mar. 19, 2019), https://www.zestfinance.com/blog/zaml-fair-our-new-ai-to-reduce-bias-in-lending ("Several mortgage lenders have tested ZAML Fair and, based on their results, ZAML Fair applied widely would eliminate 70% of the nation's gap in approval rates between Hispanic and white mortgage applicants, and cut the even larger gap between black and white borrowers by more than 40%.").

j. Enhanced Public Services

Data analytics has potential to deliver more efficient government services and to assist public safety and security, and are in use today at international, national, and local levels. For example, law enforcement is using machine learning to assist with fraud detection[67] and traffic control.[68] Data analytics is helping to mitigate the spread of HIV by identifying key people in social networks of the homeless youth population in Los Angeles.[69] In other examples, cities are using data analytics applied to big data to predict vacancies in state-funded residential properties, water-main break risks, the likelihood that illegally parked cars will block snowplows and emergency vehicles, and to predict outcomes of criminal justice measures (parole detail, bail recommendations, and sentencing).[70]

2. Challenges and Risks

The promise of immense benefits from big data and data analytics comes with challenges and risks to individuals, competition, and society, as many academics, government institutions, and other observers have warned.[71]

[67] *See, e.g.*, Christopher Rigano, *Using Artificial Intelligence to Address Criminal Justice Needs*, 280 NIJ JOURNAL (2019), *available at* https://www.ncjrs.gov/pdffiles1/nij/252038.pdf.

[68] *See, e.g.*, Francesca Baker, *The Technology that Could End Traffic Jams*, BBC (Dec. 12, 2018), http://www.bbc.com/future/story/20181212-can-artificial-intelligence-end-traffic-jams.

[69] *See* Joanna Clay, *USC Researcher, and AI, Give Homeless Youth a Helping Hand with HIV Education*, USC NEWS (July 14, 2017), https://news.usc.edu/124831/usc-researcher-and-ai-give-homeless-youth-a-helping-hand-with-hiv-education/.

[70] *Id*.

[71] *See, e.g.*, FED. TRADE COMM'N, BIG DATA: A TOOL FOR INCLUSION OR EXCLUSION? (2016), *available at* https://www.ftc.gov/system/files/documents/reports/big-data-tool-inclusion-or-exclusion-understanding-issues/160106big-data-rpt.pdf [hereinafter FTC BIG DATA REPORT]; WHITE HOUSE BIG DATA REPORT, *supra* note 2; Danielle Keats Citron & Frank Pasquale, *The Scored Society: Due Process for Automated Predictions*, 89 WASH. L. REV. 1 (2014); PAM DIXON & ROBERT GELLMAN, THE SCORING OF AMERICA: HOW SECRET CONSUMER SCORES THREATEN YOUR PRIVACY AND YOUR FUTURE (2014), *available at* http://www.worldprivacyforum.org/wp-content/uploads/2014/04/WPF_Scoring_of_America_April2014_fs.pdf.

At their most dire, these warnings raise concerns about the possibility that AI fueled by machine learning will escape human control and become the dominant form of intelligence and decision-making on earth.[72] Advances in big data and data analytics also raise more pragmatic concerns that they will increase the power of the state and of commercial entities vis-à-vis the individual, spawning more invasive forms of surveillance and societal control, and that they will produce a data analytics "arms race" with potentially significant individual privacy and national security implications. These broad concerns, however, are beyond the scope of this report.

Rather, this section predominantly focuses on the risks created by commercial applications of big data, such as the fact that big data sets will contain inaccurate or incomplete information, that it will not be representative of the relevant population, or that it will contain hidden biases. Some predictive analytics methods, including forms of machine learning, can be powerful and valuable because of the ability to uncover otherwise unknown correlations. However, correlation alone is distinct from causation, and thus there may be harmful results if such correlations lead to decisions or actions without any effort to determine a causal link between the inputs that are predicting the outputs of interest. Along with these equity concerns, big data and data analytics raise a host of privacy issues and challenges to existing privacy protection frameworks, particularly those relying on notice and choice. Likewise, there are questions about whether big data and data analytics pose challenges to competition and barriers to entry.[73] Commentators have identified the following potential impacts of failing to address these concerns.

[72] *See, e.g.*, NICK BOSTROM, SUPERINTELLIGENCE: PATHS, DANGERS, STRATEGIES (2014).

[73] *See, e.g.*, FTC BIG DATA REPORT, *supra* note 71, at 10–11; WHITE HOUSE BIG DATA REPORT, *supra* note 2, at 45–47.

a. Perpetuation of Societal Disparities

Data analytics can reinforce existing societal inequities. For example, big data has shown some organizations that employees who live closer to their jobs tend to have greater job longevity.[74] Applying this factor in hiring or promotion decisions could have a disparate impact on minority communities, whose members often have to travel longer distances to commercial hubs where jobs are often located.[75] Because data analytics methods can identify previously undetected patterns and correlations, they can give organizations new bases to exclude or undervalue already disadvantaged groups.[76]

Similarly, if minorities and less affluent consumers are disproportionality shown advertisements for for-profit colleges or subprime loans, this practice may solidify existing inequalities.[77] Whatever an organization's intent, an AI system is only as good as the underlying data and the system's intended objectives. The underlying data, or the construction of the algorithm itself, can reflect a lack of representation or hidden biases. Such problems have purportedly affected facial recognition technology[78] and the use of AI in hiring.[79]

[74] *See* John Sullivan, *How Commute Issues Can Dramatically Impact Employee Retention*, TLNT (Apr. 21, 2015), https://www.tlnt.com/how-commute-issues-can-dramatically-impact-employee-retention/ (citing case studies).

[75] *See* Edith Ramirez, Chairwoman, Fed. Trade Comm'n, Protecting Privacy in the Era of Big Data, Int'l Conference on Big Data from a Privacy Perspective (June 10, 2015), *available at* https://www.ftc.gov/system/files/documents/public_statements/671661/150610era_bigdata.pdf.

[76] *See* FTC BIG DATA REPORT, *supra* note 71, at 10–11; *see also* CATHY O'NEIL, WEAPONS OF MATH DESTRUCTION 158 (2016) (noting the use of spelling and capitalization errors to determine credit risk on loan applications).

[77] *See* FUTURE OF PRIVACY FORUM, UNFAIRNESS BY ALGORITHM: DISTILLING THE HARMS OF AUTOMATED DECISION-MAKING (2017), *available at* https://fpf.org/wp-content/uploads/2017/12/FPF-Automated-Decision-Making-Harms-and-Mitigation-Charts.pdf.

[78] *See, e.g.*, Steve Lohr, *Facial Recognition is Accurate, if You're a White Guy*, N.Y. TIMES (Feb. 9, 2018), https://www.nytimes.com/2018/02/09/technology/facial-recognition-race-artificial-intelligence.html.

[79] *See, e.g.*, James Vincent, *Amazon reportedly scraps internal AI recruiting tool that was biased against women*, THE VERGE (Oct. 10, 2018, 7:09 AM), https://www.theverge.com/2018/10/10/17958784/ai-recruiting-tool-bias-amazon-report.

Discrimination—intentional and unintentional—exists with human decision-making in the absence of AI. The one important difference, according to some commentators, is that human decision-making can evolve and progress, yet big data processes do not progress without human intervention to, for example, modify data used for algorithm training or remove factors correlated with protected characteristics.[80] As an illustration, author Cathy O'Neil observes that if a big data college application model had established itself in the early 1960s to predict college applicants who would be academically successful college students, many women might be denied admission to college because the model would likely have been trained largely on successful men.[81] Using big data in thoughtful, ethical ways to avoid such results is critical.

b. Black Box Decision-Making

Many machine-learning models allow the algorithm to choose among the available variables the best predictors of the outcome of interest. With such models, even the engineers and computer scientists overseeing the model may not understand the underlying dynamics that allow the prediction to be successful.[82] In especially complex models, they may not understand how the

[80] *See* O'NEIL, *supra* note 76, at 203–04. *See also* CAROLINE CRIADO PEREZ, INVISIBLE WOMEN: EXPOSING DATA BIAS IN A WORLD DESIGNED FOR MEN (2019).

[81] *See* O'NEIL, *supra* note 76, at 204. As a counterpoint, big data analytics can also provide solutions to correcting for bias in underlying data (*e.g.*, IBM Research developed a "methodology to reduce the bias that may be already present in a training dataset, such that any machine learning algorithm that later learns from that dataset will perpetuate as little inequity as possible.") 5 in 5 | Five Innovations that will Help Change Our Lives Within Five Years, IBM, https://www.research.ibm.com/5-in-5/ai-and-bias/.

[82] For example, researchers trained a machine learning model to recognize the differences between wolves and dogs from a training base of photographs. Ultimately it became clear that the algorithm's high level of success was because it was trained on a dataset that featured dogs typically on grass and wolves typically in the snow. Thus, the model had actually learned to distinguish grass from snow, not dogs from wolves. *See* Marco Tulio Ribeiro, Sameer Singh & Carlos Guestrin, *"Why Should I Trust You?": Explaining the Predictions of Any Classifiers,* PROCEEDINGS OF THE 22ND ACM SIGKDD INT'L CONFERENCE ON KNOWLEDGE DISCOVER AND DATA MINING (2016). In response, these researchers developed methods to improve the interpretability of *any* machine learning method, including neural networks. Increasing the transparency of otherwise opaque machine learning models remains a nascent but popular research topic, even securing government funding through DARPA's Explainable Artificial Intelligence initiative. *See*, Matt Turek, *Explainable Artificial Intelligence (XAI)*, DARPA, https://www.darpa.mil/program/explainable-artificial-intelligence.

models generate outputs.[83] It follows that machine learning-based decisions affecting individuals may lack transparency or ample procedural safeguards.[84]

c. Differential Access to Goods and Services

Data analytics can enable businesses to engage in "dynamic" pricing that varies by individual, zip code, or other factors for a product or service. Some commentators have raised concerns about potential effects of these pricing strategies on lower income communities, noting that there have been instances in which consumers in poorer neighborhoods appear to have been charged higher prices than consumers in affluent communities, perhaps due to the greater prevalence in the latter neighborhoods of competition from brick-and-mortar stores.[85] Additionally, data analytics raises the possibility that businesses may tailor prices in ways that correspond, even if unintentionally, to ethnicity or what would ordinarily be considered a protected characteristic.[86]

d. Narrowing of Choice and Information

Additional data analysis can have narrowing effects. In certain commercial contexts, for instance, big data and data analytics can limit individuals' choices and access to information than might otherwise be the case in the absence of big-data-enabled tailoring. For example, when firms present consumers with advertising based on their browsing history, the range of information to

[83] *See* Will Knight, *The Dark Secret at the Heart of AI*, MIT TECH. REV. (Apr. 11, 2017), https://www.technologyreview.com/s/604087/the-dark-secret-at-the-heart-of-ai/ ("The computers that run those services have programmed themselves, and they have done it in ways we cannot understand. Even the engineers who build these apps cannot fully explain their behavior.").

[84] *See generally* DIXON & GELLMAN, *supra* note 71.

[85] *See* FTC BIG DATA REPORT, *supra* note 71, at 11, n. 56–58 (discussing concerns of commenters). At the same time, experts have debated ways in which such practices can have pro-competitive or harmful economic effects. *See*, *e.g.*, Fed. Trade Comm'n, Transcript of Competition and Consumer Protection in the 21st Century at 27–32, 59–63 (2018), *available at* https://www.ftc.gov/system/files/documents/public_events/1418633/ftc_hearings_session_6_transcript_day_3_11-8-18.pdf [hereinafter FTC Hearings Transcript].

[86] *See* FUTURE OF PRIVACY FORUM, *supra* note 77.

which they are exposed may narrow. Likewise, personalized search results or streams of information in "feeds" in social media can lead to "filter bubbles" and "echo chambers" that create feedback loops that reaffirm and narrow an individual's thoughts and beliefs, with potentially broad ramifications for society.[87]

e. Privacy Challenges

Big data and data analytics present challenges to existing privacy frameworks, particularly those that rely heavily on notice and choice.[88] Consumers may agree to allow a company to collect non-sensitive information without appreciating the scope of data collection and that the company may use the information to infer sensitive information about the consumer.[89] For example, academic researchers were reportedly able to predict depression by analyzing social media data.[90] While such inferences can have benefits, such as early screening and prevention, they also can be used for less benign purposes, such as predatory marketing. In addition, even if consumers decline to consent to the collection or use of their data, businesses may use big data shared by other, similar consumers to draw inferences about them.[91]

[87] *See, e.g.*, Cynthia Dwork, et al., *It's Not Privacy, and It's Not Fair*, 66 STAN. L. REV. ONLINE 35 (Jan. 1, 2013).

[88] *See, e.g.*, EXECUTIVE OFFICE OF THE PRESIDENT, PRESIDENT'S COUNCIL OF ADVISORS ON SCIENCE AND TECHNOLOGY, REPORT TO THE PRESIDENT – BIG DATA AND PRIVACY: A TECHNOLOGICAL PERSPECTIVE (May 2014), *available at* https://obamawhitehouse.archives.gov/sites/default/files/microsites/ostp/PCAST/pcast_big_data_and_privacy_-_may_2014.pdf at 38; *see also*, ELECTRONIC PRIVACY INFORMATION CENTER, THE AI POLICY SOURCEBOOK (2019), *available at* https://epic.org/bookstore/ai2019/.

[89] *See, e.g.*, PRIVACY INTERNATIONAL, DATA IS POWER: PROFILING AND AUTOMATED DECISION-MAKING IN GDPR (2017), *available at* https://privacyinternational.org/sites/default/files/2018-04/Data%20Is%20Power-Profiling%20and%20Automated%20Decision-Making%20in%20GDPR.pdf; FTC BIG DATA REPORT, *supra* note 71, at 10, n.51.

[90] *See* Zeynep Tufekci, *Think You're Discreet Online? Think Again*, N.Y. TIMES (Apr. 21, 2019), https://www.nytimes.com/2019/04/21/opinion/computational-inference.html.

[91] *See* Solon Barocas et al., *Big Data's End Run Around Anonymity and Consent*, PRIVACY, BIG DATA, AND THE PUBLIC GOOD: FRAMEWORKS FOR ENGAGEMENT 44, 61–63 (Julia Lane et al. eds., 2014) (cited in FED. TRADE COMM'N, *supra* note 71, at 11, nn.59–60).

f. Foreclosing Competition and Entrenching Dominance

Big data and data analytics can create and amplify feedback effects. For example, more people using a product can mean that more data, and more diverse data, will be collected, allowing the company to both improve its products as well as potentially identify and offer new ones. This in turn can attract more customers, leading to a positive feedback loop, helping a company to grow and potentially dominate the market. Indeed, data-driven markets may "tip towards one or two products or platforms."[92] As these companies grow, the scale and scope of data necessary to compete effectively can become so large that it can serve as a significant barrier to market entry.[93] Because data is a critical input in many markets, especially online markets, a company that possesses this critical data can easily foreclose access to this input (or deny timely access to it) to existing and potential competitors, thus maintaining its dominance.[94] Likewise, Germany's Federal Cartel Office has held that if personal data is the de facto price of an ostensibly free service, a firm with market power may charge an excessively high price in terms of how much data is collected and how data is used.[95]

[92] *See* ALLEN GRUNES & MAURICE STUCKE, BIG DATA AND COMPETITION POLICY 163 (2016).

[93] *See United States v. Bazaarvoice, Inc.*, No. 13-cv-00133-WHO, 2014 U.S. Dist. LEXIS 3284 (N.D. Cal. 2014) (upholding challenge to acquisition of review site and noting a platform provider's dominance from network effects).

[94] *Id*. (Bazaarvoice's merger gave the company "complete ownership of the category and data asset," which was a reason why the merger was blocked).

[95] Bundeskartellamt Feb. 6, 2019, Facebook, Exploitative business terms pursuant to Section 19(1) GWB for inadequate data processing B6-22/16, 2019 (Ger.).

g. Collusive Pricing

Big data, combined with machine learning, creates a risk of collusion through algorithmic pricing. Although pricing algorithms can have beneficial effects to consumers and competition, commentators have raised concerns that they may facilitate explicit or tacit collusion.[96]

PART III: THE LAW AND AI

A. Introduction

As described in the prior chapters, big data and AI raise several issues specific to consumer protection, privacy, and data security. This section summarizes the legal and regulatory authority applicable to big data and AI, in both the U.S. and European Union, and compliance considerations and challenges facing companies seeking to leverage big data and AI technologies. As this Part will illustrate, existing laws regulate big data and AI in certain respects and, in other respects, have yet to address head-on key questions and challenges posed by these innovations.[97]

B. Considerations Under the FTC Act

As the nation's primary regulator of consumer privacy and data security, the Federal Trade Commission ("FTC"), is likely to play a significant role in shaping AI-related data practices. The FTC has held workshops and hearings on the subject and has issued a report on discriminatory practices resulting from the use of "big data." This section provides an overview of the FTC's

[96] *See* Ulrich Schwalbe, *Algorithms, Machine Learning, and Collusion*, 2 (2018), *available at* https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3232631; Emilio Calvano, et al., *Algorithmic Pricing: What Implications for Competition Policy?*, REV. OF INDUS. ORG. 1, 1–17 (2018); Steven Van Uytsel, *Artificial Intelligence and Collusion: A Literature Overview*, *in* ROBOTICS, AI AND THE FUTURE OF LAW 155, (Marcelo Corrales et al. eds., 2018).

[97] This Part focuses primarily on privacy, consumer protection, and equal opportunity laws and regulations. These legal regimes primarily focus on the collection, use, and disclosure of *personal* information. As a result, the considerations discussed in this Part might be irrelevant to other applications of data analytics, such as those involving agriculture and the design of industrial processes. However, separate legal frameworks might apply to those areas in general and their uses of data analytics in particular. A discussion of those considerations is beyond the scope of this Report.

consumer protection authority as it relates to personal data processing and discusses how this authority could apply specifically to data analytics methods and AI decision-making.

1. Overview of the FTC's Consumer Protection Authority

Although the FTC's consumer protection authority covers a broad spectrum of issues from advertising and financial practices to telemarketing, privacy and data security consistently rank among the FTC's top priorities.[98] This focus dates to the mid-1990s, when the FTC identified the collection and handling of personal information as key issues in the then-emerging digital economy.[99] Since then, in the absence of baseline federal legislation governing consumer privacy or data security, the FTC developed a privacy and data security enforcement program based on its consumer protection authority under Section 5 of the FTC Act.[100] To date, the FTC has brought more than 100 privacy and data security enforcement actions to stop allegedly unfair or deceptive data practices.[101] These actions have addressed a broad variety of practices, including deceptive statements about the collection and use of personal information, and unfair practices, such as the failure to reasonably secure personal information, and the collection and sharing of sensitive personal information without consumers' consent.[102]

[98] *See*, *e.g.*, Joseph Simons, Chairman, Fed. Trade Comm'n, Prepared Statement of the FTC on Oversight of the Federal Trade Commission, at 4 (July 18, 2018), https://www.ftc.gov/system/files/documents/public_statements/1394526/p180101_ftc_testimony_re_oversight_house_07182018.pdf. ("Year after year, privacy and data security top the list of consumer protection priorities at the Federal Trade Commission.").

[99] *See* FED. TRADE COMM'N, ANTICIPATING THE 21ST CENTURY: CONSUMER PROTECTION POLICY IN THE NEW HIGH-TECH, GLOBAL MARKETPLACE (1996), *available at* https://www.ftc.gov/reports/anticipating-21st-century-competition-consumer-protection-policy-new-high-tech-global.

[100] 15 U.S.C. § 45.

[101] *See*, *e.g.,* FED. TRADE COMM'N, PRIVACY & DATA SECURITY UPDATE: 2018 (2018), *available at* https://www.ftc.gov/system/files/documents/reports/privacy-data-security-update-2018/2018-privacy-data-security-report-508.pdf.

[102] *See, e.g.,* Complaint, FTC v. Vizio, Inc., Case No. 2-17-cv-00758, ¶ 33 (D.N.J. 2017), *available at* https://www.ftc.gov/system/files/documents/cases/170206_vizio_2017.02.06_complaint.pdf.

The FTC is principally a law enforcement agency that enforces Section 5's consumer protection provisions on a case-by-case basis. Except where authorized by specific statutes, such as the Children's Online Privacy Protection Act of 1998,[103] the FTC does not have notice-and-comment rulemaking authority that would allow the agency to prescribe detailed data use requirements through regulations under the Administrative Procedure Act. Instead, the FTC's enforcement actions are illustrative of the agency's views as to what it might view as unfair or deceptive in other settings.[104]

In addition, the FTC has conducted an extensive series of workshops and public hearings, and has issued multiple reports, to develop a set of best practices for consumer privacy protections and examine more specific policy questions and industry practices.[105] Of potential relevance to the use of data analytics and AI is the FTC's 2016 report on big data and a 2018 staff perspective on informational injury.[106] In November 2018, the FTC held a public hearing, as part of a broader series of hearings, that examined consumer protection and competition issues surrounding the use of "algorithms, artificial intelligence, and predictive analytics".[107] The purpose of the hearings,

[103] 15 U.S.C. §§ 6501–06.

[104] The scope of the FTC's unfairness authority remains a subject of public debate. For an argument urging the FTC to adopt a more expansive view of unfairness, see Rohit Chopra, Commissioner, Fed. Trade Comm'n, Comment (2018), *available at* https://www.ftc.gov/system/files/documents/public_statements/1408196/chopra_-_comment_to_hearing_1_9-6-18.pdf.

[105] *See, e.g.*, FED. TRADE COMM'N, PROTECTING CONSUMER PRIVACY IN AN ERA OF RAPID CHANGE 1–3 (2012), *available at* https://www.ftc.gov/sites/default/files/documents/reports/federal-trade-commission-report-protecting-consumer-privacy-era-rapid-change-recommendations/120326privacyreport.pdf (describing public engagement process leading to adoption of report) [hereinafter FTC PRIVACY REPORT].

[106] *See* FTC BIG DATA REPORT, *supra* note 71; FED. TRADE COMM'N, FTC INFORMATIONAL INJURY WORKSHOP: BE AND BCP STAFF PERSPECTIVE (2018), *available at* https://www.ftc.gov/reports/ftc-informational-injury-workshop-be-bcp-staff-perspective [hereinafter FTC STAFF PERSPECTIVE ON INFORMATIONAL INJURY].

[107] *See FTC Hearing #7: The Competition and Consumer Protection Issues of Algorithms, Artificial Intelligence, and Predictive Analytics*, FTC (2018), https://www.ftc.gov/news-events/events-calendar/ftc-hearing-7-competition-consumer-protection-21st-century.

according to the Commission, included "evaluation of the FTC's near- and long-term law enforcement policy agenda" and to identify "areas that warrant additional study."[108]

2. *How the FTC Act Could Apply to AI*

The FTC has not brought an enforcement action specifically related to AI decision-making generally, or machine learning specifically; however, the FTC's areas of focus with respect to its privacy and data security enforcement and policy may inform how the agency assesses consumer protection challenges that arise in connection with a company's data analytics practices, particularly with respect to transparency, concrete harms, and data security.

a. Transparency

To avoid deception under Section 5, statements to consumers about the collection, use, and disclosure of their personal information must be truthful and not misleading.[109] The FTC's privacy and data security enforcement frequently focuses on companies' alleged express or implied misrepresentations about their personal data use practices. The requirement that statements to consumers must be truthful and not misleading applies equally to a company's AI practices.[110]

Yet AI presents unique challenges to the FTC's deception authority. For example, a company's disclosures to consumers about its AI practices are unlikely to include a detailed description of the data sets used to train algorithms or machine-learning systems used for decision-making about consumers.[111] In turn, in many circumstances, it may be difficult for regulators to analyze and meaningfully understand this data in discovery or by civil investigative demand. In

[108] S*ee Hearings on Competition and Consumer Protection in the 21st Century*, FTC (2018), https://www.ftc.gov/policy/hearings-competition-consumer-protection.

[109] 15 U.S.C. § 45(a); FTC Policy Statement on Deception, 103 F.T.C. 110, 174 (1984) (appended to Cliffdale Assocs., Inc., 103 F.T.C. 110, 174 (1984)).

[110] *See, e.g.*, FTC BIG DATA REPORT, *supra* note 71, at 23

[111] *See, e.g.*, Yavar Bathaee, *The Artificial Intelligence Black Box and the Failure of Intent & Causation*, 31 HARV. J.L. & TECH. 890 (2018).

the absence of express or clearly implied discriminatory statements by a company, the FTC may find it challenging to use its deception authority to enforce a standard of algorithmic transparency that some advocates and policymakers have recommended.[112] Nonetheless, to the extent other laws, such as the California Consumer Protection Act (discussed *infra*), require businesses to provide access to personal information about a consumer (including in the data analytics context), a company's misrepresentation about how it satisfies those obligations might provide a basis for the FTC's deception authority.

b. Consumer Harm

Consumer harm—a requirement for enforcing Section 5's unfairness standard and a focus of the FTC's examination of data use in recent years—is frequently controversial as applied to privacy due to disagreement about the circumstances under which privacy harms constitute "substantial injury." Most stakeholders agree that injuries resulting from medical identity theft, doxing (the practice of researching and broadcasting private or identifying information about an individual or organization), disclosure of private information, and the erosion of trust are potentially significant injuries.[113] There is less agreement, however, on injuries that cause reputational injury or embarrassment, not to mention more controversial areas of potential

[112] *See*, Edith Ramirez, Chairwoman, Fed. Trade Comm'n, *Protecting Privacy in the Era of Big Data* (June 10, 2015), *available at https://www.ftc.gov/system/files/documents/public_statements/671661/150610era_bigdata.pdf* ("For example, the [Fair Credit Reporting Act] does not apply when businesses use their own in-house data analytics to make decisions about their customers or employees. And although [the Equal Credit Opportunity Act] would prohibit racial distinctions in terms of access to credit, it may not prohibit those distinctions in the types of advertisements served. Thus, a minority consumer may only see ads for subprime products and may never know about the availability of better credit offers. Finally, Section 5 does not require notices or choices about big data practices. This is why our efforts at the FTC must go beyond enforcement of existing laws.").

[113] *See* FTC STAFF PERSPECTIVE ON INFORMATIONAL INJURY, *supra* note 106, at 1–3 (Oct. 2018) (noting that "although no participants [in the FTC's December 2017 informational injury workshop] disputed that the potential for injuries described above are real, there was robust debate over whether and when governments should intervene to address these injuries, particularly in light of the benefits").

consumer harm, such as price discrimination based on machine learning methods.[114] Additionally, the FTC's recent informational hearings highlighted the risks of algorithmic bias leading to discrimination resulting from skewed data sets or improper weighting of factors in deep-learning neural networks.[115] As discussed in the FTC's 2016 report on big data, discriminatory outcomes could violate both civil rights laws and Section 5 of the FTC Act.[116]

c. Data Security

Two aspects of the FTC's approach to data security deserve close attention for AI purposes. First, the FTC's expectation that companies will use reasonable data security measures to protect personal information under their control applies in the context of AI. The long-standing view that any "company that keeps sensitive consumer information must take steps to ensure that the data is held in a secure manner"[117] carries no less force when personal information is part of an AI system based on a machine learning model. Machine learning data sets, which in some cases require vast amounts of potentially sensitive information, can pose serious security risks. Beyond this, researchers have demonstrated that "adversarial" inputs into AI systems can affect the decision-making models that the systems use, creating the potential for unexpected behavior that can cause

[114] *See, e.g.*, FTC Hearings Transcript, *supra* note 85, at 277–78 (observing that "we have discriminated on price . . . for generations); *id.* at 279–80 (arguing that the underlying issue with price discrimination is "taking advantage of people").

[115] *See id*. (discussing that Facebook "improved the state-of-the-art performance on the task of face verification by almost 20 percent, which was great news because it showed that there were effective techniques being employed using deep learning. However, if you look at [the targeted standard], you'll see that it is 78 percent male and 84 percent white. So if this is the gold standard we're using, we're giving ourselves a false sense of progress which can lead to misleading technology.").

[116] FTC Big Data Report, *supra* note 71, at 17–23.

[117] Press Release, Fed. Trade Comm'n, CardSystems Solutions Settles FTC Charges (Feb. 23, 2006), https://www.ftc.gov/news-events/press-releases/2006/02/cardsystems-solutions-settles-ftc-charges.

real-world harms such as fraud and physical injury.[118] This possibility creates a data security consideration that may not be part of many existing data security programs.

The second dimension of data security that deserves careful attention from AI developers and users is vendor management. Insufficient vetting of or disclosures about vendors used to deliver products and services have been at the center of several recent FTC privacy and data security enforcement actions.[119] Because machine learning development typically demands large amounts of data and processing power, as well as a concentration of technical expertise, many companies turn to third parties for machine learning-related services. As FTC enforcement expands into AI, it is likely to demand due diligence of a vendor's privacy and data security practices and the development of appropriate contractual protections and oversight of AI vendor selection and use. As applications of AI decision-making systems fueled by machine learning become part of more consequential actions, such as financial decisions and the control of automobiles, the FTC may require that these vendor assessment practices expand beyond traditional privacy and data security considerations.

[118] *See* Nicole Kobie, *To Cripple AI, Hackers are Turning Data Against Itself*, WIRED (Sept. 11, 2018), https://www.wired.co.uk/article/artificial-intelligence-hacking-machine-learning-adversarial ("Building adversarial examples is no easy task, often requiring access to technical details of a neural network, such as the model architecture, known as "white box" access. That said, robust attacks have been described that don't require detailed network information; those black-box attacks could prove more useful for outsiders to attack a system, as they're transferable across different neural networks.")

[119] *See, e.g.*, Decision and Order, BLU Products, Inc., FTC Case No. 172-3025 (Sept. 6, 2018), *available at* https://www.ftc.gov/system/files/documents/cases/172_3025_c4657_blu_decision_and_order_9-10-18.pdf; Decision and Order, GMR Transcription Servs. Inc., FTC Case No. 122-3095 (Aug. 14, 2014), *available at* https://www.ftc.gov/system/files/documents/cases/140821gmrdo.pdf.

C. The Intersection of Data Analytics and Eligibility Determinations

1. Big Data and Eligibility Determinations

The widespread use of smart devices and social media has greatly increased the amount of consumer data that firms can use to analyze and predict consumer behavior.[120] Through their interactions with smart devices, consumers leave behind a digital footprint containing information about their daily lives, including, among others, browsing history, shopping habits, location history, and social network.[121] Companies are increasingly using these large data sets that include consumer characteristics to make predictions about a consumer's future behavior, including to determine whether a consumer may be eligible for goods and services, and on what terms to offer them.

Big data increasingly affects consumer eligibility determinations for access to credit, employment, insurance, housing, and educational opportunities, among other things.[122] Companies make eligibility determinations by comparing a set of characteristics of an applicant to the characteristics of similarly-situated consumers about whom outcomes (*e.g.*, loan repayment) are known, on the assumption that these characteristics can predict the applicant's likely future behavior.[123]

Traditionally, companies use a limited set of consumer characteristics to determine eligibility. For example, credit bureaus use an individual's history of banking, borrowing, and

[120] *See, e.g.*, WHITE HOUSE BIG DATA REPORT, *supra* note 2; FTC BIG DATA REPORT, *supra* note 71; *Hearings on Competition and Consumer Protection in the 21st Century*, *supra* note 108; ROBINSON + YU, KNOWING THE SCORE: NEW DATA, UNDERWRITING, AND MARKETING IN THE CONSUMER CREDIT MARKETPLACE (2014), *available at* https://www.upturn.org/static/files/Knowing_the_Score_Oct_2014_v1_1.pdf.

[121] *See e.g.*, FTC BIG DATA REPORT, *supra* note 71, at 1.

[122] *See e.g.*, *id*. at 13.

[123] *See e.g.*, *id*. at 16.

repayment to calculate a credit score for that individual.[124] The credit score then is used by lenders to determine whether the individual is eligible for a loan and the amount and other terms of such loan. Credit bureaus may obtain the information used to calculate credit scores from "furnishers" such as commercial banks, credit card companies, and mortgage lenders.[125]

Participants in the alternative credit-scoring and underwriting industry now use non-traditional data to make repayment predictions, including data about consumers' shopping habits, location and web browsing history, and social media records.[126] In turn, this has led to a new market of data furnishers, such as social media companies and online tracking companies.[127]

Data collected from consumers' online activities tends to be voluminous and complex.[128] Because traditional techniques are not well suited for such large and complex data sets, companies are adopting statistical modeling techniques to support eligibility determinations.[129] For example, traditional credit scoring may rely on a weighted average of a handful of factors (*e.g.*, the length of credit history, payment history, outstanding debt, debt-to-credit ratio, and pursuit of new credit),[130] and eligibility determinations for insurance underwriting purposes traditionally rely on

[124] ROBINSON + YU, *supra* note 120, at 4.

[125] *See e.g.*, Mikella Hurley & Julius Adebayo, *Credit Scoring in the Era of Big Data*, 18 YALE J.L. & TECH. 148, 154 (2016); ROBINSON + YU, *supra* note 120, at 8.

[126] *See e.g.*, ROBINSON + YU, *supra* note 120, at 15; Hurley & Adebayo, *supra* note 125, at 164–165; Matthew Adam Bruckner, *The Promise and Perils of Algorithmic Lenders' Use of Big Data*, 93 CHI.-KENT L. REV. 3, 12–14 (2018).

[127] *See e.g.*, ROBINSON + YU, *supra* note 120, at 16.

[128] *See e.g.*, WHITE HOUSE BIG DATA REPORT, *supra* note 2, at 2–3.

[129] *See e.g.*, ROBINSON + YU, *supra* note 120, at 2; *see also*, FINANCIAL STABILITY BOARD, ARTIFICIAL INTELLIGENCE AND MACHINE LEARNING IN FINANCIAL SERVICES 13–14 (2017), *available at* https://www.fsb.org/wp-content/uploads/P011117.pdf.

[130] *See e.g.*, ROBINSON + YU, *supra* note 120, at 9.

a manual process by a human decision-maker.[131] Sources of alternative data, however, require different techniques because the data inputs that best predict certain consumer behavior may not be obvious, and flawed or biased inputs could result in inaccurate decision-making. New predictive analytics models may adopt machine learning techniques that "train" an algorithm to identify the data inputs that correlate with the consumer behavior of interest, thereby appropriately weighting to the data inputs.[132] If an algorithm is trained using a large and representative sample of the population, the resulting algorithm may be able to use a wide range of data inputs, assign the appropriate weight to each, and predict consumer behavior.[133]

2. Potential Advantages

The principal benefit of big data in eligibility determinations is that it can predict the behavior of more people, more accurately, and at lower cost.

a. Broader Reach

As described in the Costs and Benefits chapter of this report, new predictive analytics models using big data can leverage characteristics from a broader set of the population than traditional approaches, which in turn, can benefit traditionally underserved populations.[134]

b. Accuracy

Big data analytics models may also more accurately predict consumer behavior. A model that can incorporate myriad factors affecting consumer behavior theoretically should yield better predictions.[135] As such, a model's precision improves as it is "trained" on a larger number of

[131] Shanique Hall, *How Artificial Intelligence is Changing the Insurance Industry*, CIPR NEWSLETTER (Nat'l Ass'n of Ins. Commissioners, Ctr. for Ins. Pol. and Res., Kansas City, M.O.) Aug. 2017, at 6, *available at* https://www.naic.org/cipr_newsletter_archive/vol22_ai.pdf?63.

[132] *See e.g.*, Bruckner, *supra* note 126, at 16–17.

[133] *Id.*

[134] *See e.g.*, FTC BIG DATA REPORT, *supra* note 71, at 5–6.

[135] *See e.g.*, Bruckner, *supra* note 126, at 18.

consumers.[136] In addition, in areas such as the underwriting industry where a human reviewer makes eligibility determinations, big data analytics could reduce bias from human decision-making and produce consistent outcomes.[137]

c. Lower Cost

As the use of data analytics increasingly automates the process of determining eligibility, the associated costs may decline as well.[138] Additionally, as data analytics models generate increasingly accurate predictions as to likelihood of repayment, incidence of loan defaults may decrease, potentially reducing costs for lenders.[139]

3. Legal Risks: FCRA, ECOA, and FHA

The use of data analytics to make eligibility determinations is subject to several legal requirements, including fairness, transparency, accuracy, and non-discrimination standards under the Fair Credit Reporting Act ("FCRA"), the Equal Credit Opportunity Act ("ECOA"), and the Fair Housing Act ("FHA").[140] These laws, which play a critical role in consumer and civil rights protections in the United States, pre-date the big data economy by many years. As a result, it may be difficult to determine whether and how some of these requirements apply to AI development and use, and how the use of AI complies with such requirements in certain contexts. The following discussion highlights key accuracy, transparency, and non-discrimination considerations under the FCRA, ECOA, and FHA.

[136] For example, a model that considers (or is "trained" on) only the consumers living in Hawaii may not yield good predictions on what snow boots consumers in Alaska prefer.

[137] *See e.g.*, Bruckner, *supra* note 126, at 18–19.

[138] *See e.g.*, Bruckner, *supra* note 126, at 21.

[139] *Id*.

[140] 15 U.S.C. § 1681; 15 U.S.C. § 1691; 42 U.S.C. § 3601.

a. Accuracy

Accuracy is a key principle underlying the FCRA, which primarily applies to consumer reporting agencies ("CRAs") that compile and provide consumer reports used for certain eligibility determinations (e.g., credit, employment, and insurance).[141] Two important and related CRA obligations are to provide consumers with access to their own information and the ability to correct any error.[142] In addition, the FCRA requires furnishers (i.e., those who provide information to CRAs for inclusion on a consumer report) to provide, among other things, accurate information to CRAs and resolve consumer disputes about the accuracy of information they furnish.[143]

Non-traditional data that is furnished or used for eligibility determinations regulated by the FCRA is subject to the FCRA's accuracy obligations. The FTC has emphasized this point by issuing enforcement warning letters and bringing enforcement actions against companies that used non-traditional data (or provided their services through then-novel formats, such as mobile apps) and allegedly engaged in credit reporting without meeting their obligations under the FCRA.[144]

Big data may also pose practical challenges to fulfilling CRAs' and furnishers' accuracy obligations. For example, some online data sources track cookies embedded in web browsers instead of individual consumers. When multiple consumers share a device, it is difficult to

[141] *See* 15 U.S.C. §§ 1681a(f), 1681b(a). *See also id.* § 1681a(d)(1) (providing that consumer reports may include "any information . . . bearing on a consumer's credit worthiness, credit standing, credit capacity, character, general reputation, personal characteristics, or mode of living" that is "used or expected to be used or collected in whole or in part" to determine a consumer's eligibility for credit, employment, and certain other authorized purposes).

[142] 15 U.S.C. §§ 1681a(f), (d)

[143] *See* 15 U.S.C. §§ 1681s-2(a)(8), and 1681s-2(e).

[144] *See, e.g.*, Complaint, United States v. Spokeo, Inc., No. 2-12-cv-05001-MMM-SH (C.D. Cal. filed June 7, 2012), https://www.ftc.gov/sites/default/files/documents/cases/2012/06/120612spokeocmpt.pdf; Complaint, United States v. Instant Checkmate, Inc., No. 3:14-cv-00675-H-JMA (S.D. Cal. filed Mar. 24, 2014), https://www.ftc.gov/system/files/documents/cases/140409instantcheckmatecmpt.pdf.

ascertain who actually visited a website or made a particular online purchase.[145] Consequently, it may be impossible for a furnisher to assess the accuracy of such information, and for CRAs to develop procedures that assure the "maximum possible accuracy" of the information in its consumer reports, as it may be unclear to which consumer the information relates.[146] In other circumstances, such as where social media data is concerned, the furnisher may be unclear, and many data brokers who provide such data to CRAs do not consider themselves furnishers at all.[147]

b. Transparency and Explainability

The FCRA and ECOA also impose certain disclosure and transparency obligations that may present challenges to the development and use of data analytics. Under the FCRA, for example, a consumer report user that takes any "adverse action" (e.g., denial of credit or employment) based on a consumer report must provide the consumer with notice of the adverse action and information about the CRA that provided that consumer report, including instructions on how to obtain the report.[148] Should the consumer request the report, the CRA must disclose to the consumer all information in her file.[149] The consumer has the right to review the report and dispute any incorrect or incomplete information. If the consumer files a dispute, the CRA must "conduct a reasonable reinvestigation to determine whether the disputed information is

[145] *See e.g.*, Hurley & Adebayo, *supra* note 125, at 189.

[146] *See* McCready v. eBay, Inc., 453 F.3d 882 (7th Cir. 2006) (holding that "a consumer" within the meaning of 15 U.S.C. § 1681a(d) must at a minimum "be an identifiable person"); *see also* 15 U.S.C. § 1681e(b) (requiring CRAs to "follow reasonable procedures to assure maximum possible accuracy of the information concerning the individual about whom the report relates").

[147] NAT'L CONSUMER LAW CTR., BIG DATA: A BIG DISAPPOINTMENT FOR SCORING CONSUMER CREDITWORTHINESS 25–26 (2014), *available at* https://www.nclc.org/issues/big-data.html [hereinafter NAT'L CONSUMER LAW CTR. 2014 STUDY].

[148] 15 U.S.C. § 1681m(a).

[149] 15 U.S.C. § 1681g.

inaccurate"[150] and convey the dispute to the furnisher of that information.[151] When using big data to make such determinations, it may be difficult to produce a consumer report given the sheer volume of data involved.[152] Consumer reports that contain social media data can be particularly problematic because the data used for an eligibility determination might not directly relate to that consumer (*e.g.*, information on the consumer's Facebook friends).[153] In addition, consumer reports used in data analytics may be too voluminous for consumers to comprehend and identify any errors.[154]

As a dimension of transparency—and a further challenge to uses of machine learning—the FCRA and ECOA also require companies to provide some explanation of adverse actions. Specifically, ECOA requires consumer notifications of adverse actions relating to credit to explain, among other things, the reasons for such action.[155] Similarly, for decisions based in whole or in part on a credit score, the FCRA requires that the user provide a notice that sets forth information about that credit score, including the primary factors that affected the score.[156] ECOA requires companies to document and be able to justify the model that they used to make the decision, including "the factors considered, the weight applied to those factors, and the credit cutoff

[150] 15 U.S.C. § 1681i(a).

[151] *Id.*

[152] *See, e.g.,* NAT'L CONSUMER LAW CTR. 2014 STUDY, *supra* note 147.

[153] *Id.* at 24.

[154] *Id.*

[155] 12 CFR § 1002.9.

[156] 15 U.S.C. § 1681c(d).

determination, in order to demonstrate the empirical backing and legitimate business need of their specific model."[157]

Big data analytic tools can make it difficult to comply with these requirements. As described earlier in the report, the inner workings of algorithms can be opaque and can deliver inscrutable or unexplainable outputs.[158] While big data predictive analytic tools are capable of detecting correlations among a huge variety of data inputs, they do not provide explanations of why certain inputs were assigned greater weight than others. With machine learning processes, weighting can be particularly difficult to explain, and it may be impossible to explain how certain factors *caused* the predicted outcome. Consequently, explaining results derived from machine learning processes to consumers and regulators may be difficult.

c. Discrimination

Analytic tools for big data may raise concerns under ECOA and the FHA. Predictive analytics can pose risks of discrimination, even where the underlying models do not explicitly use protected characteristics, though there are techniques to correct for bias in underlying data.[159] Other big data models may rely on certain variables that are highly correlated with protected characteristics, and effectively serve as proxies for those characteristics. For example, literature has criticized the use of geographical characteristics such as zip codes, which can serve as a proxy for race.[160]

[157] Lawrence D. Kaplan, Gerald Sachs & Kristin S. Teager, *Addressing ECOA Risk in Marketplace Lending*, (Aug. 25, 2016), PAUL HASTINGS, https://www.paulhastings.com/publications-items/details/?id=802eea69-2334-6428-811c-ff00004cbded.

[158] WHITE HOUSE BIG DATA REPORT, *supra* note 2, at 7–8.

[159] *See* FTC BIG DATA REPORT, *supra* note 71.

[160] *See, e.g.*, FTC BIG DATA REPORT, *supra* note 71; FED. TRADE COMM'N, ANNUAL REPORT OF THE FEDERAL TRADE COMMISSION 47 (1980) (civil penalty settlement that alleged Amoco's use of consumer's zip codes as a factor in credit discriminated against blacks and Hispanic applicants), *available at* https://www.ftc.gov/sites/default/files/documents/reports_annual/annual-report-1980/ar1980_0.pdf.

Other problems arise where models used in predictive analytics contain bias that results in disparate treatment. Input bias, where a model contains biased data (such as under-representation of a particular subpopulation), can result in skewed outputs. Alternatively, learning algorithms can reflect the biases of human decision-makers. Attention to the outcomes of AI implementations may be necessary to detect and remedy any unwanted bias that may occur in practice.

D. Considerations Under the General Data Protection Regulation

The European Union's comprehensive General Data Protection Regulation ("GDPR") intersects with data analytics in several ways, and European regulators have expressed a keen interest in data analytics and AI developments.[161] This section describes the GDPR's baseline approach to data processing, as it applies to AI, and the regulation's restrictions on the use of systems for automated decision-making and associated exemptions.

1. GDPR and AI: The Fundamentals

Personal data "processing" under the GDPR broadly includes, among other things, the use of personal data to develop AI systems, and the use of AI for decision-making about data subjects.[162] The requirements applicable to such processing include the following.

[161] *See, e.g.*, Giovanni Buttarelli, European Data Protection Supervisor, Keynote Speech on Privacy, Data Protection and Cyber Security in the Era of AI at the Telecommunications and Media Forum: Artificial Intelligence and the Future Digital Single Market (Apr. 24, 2018), https://edps.europa.eu/sites/edp/files/publication/18-04-24_giovanni_buttarelli_keynote_speech_telecoms_forum_en.pdf.

[162] The GDPR's definitions of "personal data" and "processing" are broad. "Personal data" means "means any information relating to an identified or identifiable natural person ('data subject')." Council Regulation 2016/679 of the European Parliament and of the Council on the Protection of Natural Persons with Regard to the Processing of Personal Data and on the Free Movement of Such Data, and Repealing Directive 95/46/EC (General Data Protection Regulation), art. 4(1), 2016 O.J. (L 119) 1, 33 [hereinafter GDPR]. "Processing" means "any operation or set of operations which is performed on personal data or on sets of personal data, whether or not by automated means, ..." *Id.* art. 4(2).

a. Lawfulness, Fairness, and Transparency

Under the GDPR, all data processing must be lawful, fair, and transparent, to promote accountability and to facilitate control over the personal data of data subjects.[163] The large amounts of personal data, and the complex and potentially opaque processing operations typically involved in developing AI systems, present certain challenges when describing AI-related processing in a manner that is "easily accessible and easy to understand."[164]

b. Data Minimization and Purpose Specification

The GDPR requires both purpose specifications and data minimization for processing personal data, both of which create tension with many big data predictive analytics systems that depend on personal data.[165] Taken together, these principles require controllers to provide "explicit and legitimate purposes" for processing personal data, to limit their processing to these purposes, and to limit the amount of personal data they process to what is "adequate," "relevant," and necessary to achieve the stated purposes.[166] Yet, as described in this report, training machine learning models typically requires vast amounts of data.[167] As a result, to comply with GDPR, companies that process personal data for machine learning purposes must reasonably demonstrate that their machine learning activities align with the purposes communicated to data subjects at the

[163] *See id.* art. 5(1), at 35; Article 29 Working Party, Guidelines on Transparency Under Regulation 2016/679, at 5–6 (Apr. 2018), *available at* https://ec.europa.eu/newsroom/article29/document.cfm?action=display&doc_id=51025.

[164] *See* GDPR, *supra* note 162, ¶ 39, at 7.

[165] *See id.* art. 5(1)(c), (d), at 35.

[166] *Id.*

[167] *See Communication from the Commission to the European Parliament, the European Council, the European Economic and Social Committee, and the Committee of the Regions: Artificial Intelligence for Europe*, at 11, COM (2018) 237 final (Apr. 25, 2018) [hereinafter *AI Communication*].

time the data at issue was collected, and otherwise satisfy the GDPR's data processing principles.[168]

c. Lawful Basis for Processing

The GDPR requires a lawful basis to process personal data, including to develop AI systems or to apply AI to personal data in practice.[169] Legal bases recognized under the GDPR include consent of the data subject for one or more specific purposes, and the pursuit of the "legitimate interests" of the controller or a third party.[170] The EU data protection authorities have not clearly established how "specific" consent must be in connection with AI-related data processing. The "legitimate interest" basis also raises questions, as this legal basis must factor in "the interest or fundamental rights and freedoms of the data subject, which require protection of personal data."[171] Defining these limits and determining when they "override" the interests of a controller engaged in processing personal data for AI purposes remain unsettled issues.

d. Restrictions on Processing Special Categories of Information

The GDPR prohibits the processing of "special categories" of personal data unless a data subject gives "explicit consent" to the processing or another enumerated exception applies.[172] Special categories of personal data under GDPR include data "revealing racial or ethnic origin, political opinions, religious or philosophical beliefs, or trade union membership, and the processing of genetic data, biometric data for the purpose of uniquely identifying a natural person,

[168] *See* GDPR, *supra* note 162, art. 5, at 35.

[169] *See id.* art. 6(1), at 36.

[170] *See id.* art. 6(1)(a), (f), at 36. *See also id.* art. 4(11), at 34, ¶32, at 6 (requiring "consent" to be "freely given, specific, informed and unambiguous").

[171] *Id.* at art. 6(1)(f), at 36.

[172] *Id*. art. 6(1), (2), at 36 (establishing general prohibition on processing special categories of data and enumerating exceptions under which processing is permissible).

data concerning health, or data concerning a natural person's sex life or sexual orientation."[173] Companies that engage in machine learning using special categories of data must carefully consider their heightened obligations when processing this data.

e. Data Subject Rights

The GDPR provides individuals with several rights relating to their personal data, including the rights of access, rectification (correction), deletion, and data portability, and the right to restrict or otherwise object to processing.[174] These rights apply to personal data used to develop any analytic model as well as the use of these models to process personal data. Depending on how individuals exercise these rights, they could affect the development of the models. For example, if substantial numbers of individuals exercise their rights of deletion with respect to a specific controller, the resulting data set may not be representative of the relevant population, or insufficiently broad to train the machine learning algorithm on the full range of relevant situations.

It also is possible to reconstruct personal data used to train machine learning models from the models themselves.[175] For example, researchers have used "model inversion" and "membership inference" attacks to re-identify data from recommender systems and medical dosing systems.[176] If these types of attacks are feasible against the machine learning models that companies use in practice, the corresponding models could contain personal data and, as a result, be subject to the full set of data subject rights. Determining how to comply with applicable requests and assessing the impact that they may have on uses of machine learning remain open challenges under the GDPR.

[173] *Id.* art. 9(1), at 38.

[174] *See id*. arts. 15–21, at 43–46.

[175] *See* Michael Veale, Reuben Binns & Lilian Edwards, *Algorithms That Remember: Model Inversion Attacks and Data Protection Law*, 376 PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOC'Y A (2018).

[176] *See id.* 4–6 (reviewing technical literature).

f. Governance and Accountability

The GDPR also imposes various governance requirements that implicate AI systems. For example, the GDPR requires a controller to conduct a data protection impact assessment ("DPIA") prior to processing that creates a "high risk to the rights and freedoms of natural persons."[177] A DPIA also is required for "automated processing" that results in significant "legal" or similar effects on individuals, or includes large-scale processing of special categories of data.[178] The DPIA requirement does not restrict personal data processing for AI purposes, but requires a controller to conduct a systematic and in-depth review of its planned data processing operations to assess, among other things, the "necessity and proportionality" of those operations, risks to individuals, and ways to mitigate those risks.[179]

g. Privacy by Design and Default

The GDPR could affect the development of AI through the requirement to "implement appropriate technical and organizational measures" to protect individual rights.[180] The appropriateness of these measures may reflect not only the risks to individual privacy rights but also the "nature, scope, context and purposes of processing" as well as costs and the state of the technical art. Applying these factors to specific AI implementations is likely to require highly fact-specific planning and judgment.

[177] *See* GDPR, *supra* note 162, art. 35(1), at 53.

[178] *See id.* art. 35(3), at 53.

[179] *See id.* art. 35(7), at 54.

[180] *See id*. art. 25(1), at 48.

h. Cross-Border Data Transfers

The GDPR restricts cross-border transfers of personal data to countries that lack an "adequate" level of data protection, as determined by the European Commission.[181] AI development and use may depend on data collected globally. For example, applying AI to cybersecurity and healthcare challenges might require data from multiple countries or regions to understand the global nature of the underlying phenomena. Any transfer of personal data to train analytic models (or other purposes) to a country that is not viewed as adequate must take place under other conditions that meet the adequacy requirement, such as the EU-U.S. Privacy Shield, under standard contractual clauses or binding corporate rules, or under a specific derogation from the adequacy requirement.[182]

2. Automated Processing Under the GDPR

The GDPR could limit certain uses of data analytics through its prohibition on decisions "based solely on automated processing, including profiling," which have legal or "similarly significant[]" effects on the individual.[183] Regardless of whether such decisions involve machine learning, the GDPR permits the automated processing only under the following circumstances: if

[181] *See id.* art. 45(1), at 61.

[182] *See id.* art. 45(3), at 61; Commission Implementing Decision 2016/1250 of 12 July 2016 Pursuant to Directive 95/46/EC of the European Parliament and of the Council on the Adequacy of the Protection Provided by the EU-U.S. Privacy Shield, 2016 O.J. (L 207) 1; *Report from the Commission to the European Parliament and the Council on the Second Annual Review of the Functioning of the EU-U.S. Privacy Shield* COM (2018) 860 final (Dec. 19, 2018) (affirming 2016 Privacy Shield adequacy decision); *Standard Contractual Clauses (SCC)*, EUROPEAN COMM'N, https://ec.europa.eu/info/law/law-topic/data-protection/international-dimension-data-protection/standard-contractual-clauses-scc_en; *Binding Corporate Rules*, EUROPEAN COMM'N, https://ec.europa.eu/info/law/law-topic/data-protection/international-dimension-data-protection/binding-corporate-rules-bcr_en. *See also* GDPR, *supra* note 162, art. 49, at 64 (enumerating conditions under which transfers to a third country in the absence of an adequacy finding are permissible).

[183] GDPR *supra* note 162, art. 22(1), at 46. Decisions based solely on automated processing have "no human involvement in the decision process." *See* Article 29 Data Protection Working Party, Guidelines on Automated Individual Decision-Making and Profiling for Purposes of Regulation 2016/679, at 20 (Feb. 6, 2018) [hereinafter Art. 29 Working Party AI Guidelines].

the processing is performed with the explicit consent of the individual; is necessary for the performance of a contract between the individual and controller; or is specifically authorized by EU or Member State law. Controllers also must disclose, among other things, "meaningful information about the logic involved" in automated decision-making.[184] To the extent the Controller does not meet the applicable conditions, the GDPR requires some element of human intervention.[185] The intent of these provisions is to promote individual autonomy and empowerment[186] and the creation of opportunities for public engagement.[187]

The requirements applicable to automated decision-making have the potential to inhibit uses of autonomous or adaptive AI (i.e., AI systems that make decisions without human involvement), and raise questions regarding when a decision can be considered "solely" autonomous (e.g., where an AI system simply makes a recommendation to a human decision-maker). Moreover, the "legal effects" standard could significantly limit the range of uses of AI that require compliance with the conditions discussed above. The GDPR does not define such effects, though it cites consumer credit decisions and employment recruiting as examples.[188] The European Data Protection Board advises that "legal effects" concern an individual's legal rights, legal status, or rights under a contract, and "similarly significant" effects include those with a permanent impact or that lead to exclusion or discrimination.[189] Fully automated uses of AI warrant close review to determine whether they may produce these effects.

[184] *See* GDPR, *supra* note 162, arts. 15(1)(h) and 22(2), at 43, 46.

[185] *See id*. ¶72, art. 22(1), at 14, 46.

[186] *See* HIGH-LEVEL EXPERT GROUP ON ARTIFICIAL INTELLIGENCE, ETHICS GUIDELINES FOR TRUSTWORTHY AI 16 (2019).

[187] INT'L CONFERENCE OF DATA PROTECTION AND PRIVACY COMMISSIONERS, DECLARATION ON ETHICS AND DATA PROTECTION IN ARTIFICIAL INTELLIGENCE (2018).

[188] *See* GDPR *supra* note 162, ¶71, at 14.

[189] *See* Art. 29 Working Party AI Guidelines *supra* note 183, at 21.

3. *Exemptions and Mitigations*

Although the GDPR imposes comprehensive limits on data processing that may have significant effects on AI, there are a two significant exemptions and mitigations to consider.

a. Pseudonymization and Anonymization

The GDPR provides some incentives for controllers and processors to reduce or eliminate the identifiable nature of personal data, which may provide some flexibility for AI purposes. Specifically, the GDPR recognizes pseudonymized data, which "can no longer be attributed to a specific data subject without the use of additional information," provided certain other technical and organizational measures are in place.[190] Pseudonymized data is personal data under the GDPR, but pseudonymization can help provide appropriate safeguards for processing not based on consent, as well as comply with the GDPR data security and data protection by design requirements.[191]

Anonymous data—that is, data that does not relate to an individual—is not personal data subject to GDPR requirements.[192] As a result, controllers can use anonymous data for AI purposes without GDPR-imposed restrictions. Achieving true anonymization, however, may be challenging, and, depending on the AI application, the resulting anonymization may reduce or eliminate the data's utility.

b. Statistical Purpose Exemption

A second GDPR provision that could provide flexibility for AI is the statistical purpose exemption.[193] Specifically, processing of personal data for "statistical purposes" is compatible

[190] GDPR *supra* note 162, art. 4(5), at 33.

[191] *Id*. arts. 6(4)(e) (lawfulness of processing), 25(1) (data protection by design and default), and 32(1)(a) (security of processing), at 37, 48, 51.

[192] *Id.* 162, ¶26, at 5.

[193] *See id.* art. 89(1), at 84.

with the purpose for which such data originally was collected, provided "appropriate safeguards" are in place.[194] The parameters of this exemption, however, remain unclear, including the definition of a "statistical purpose" as it relates to the use of personal data in machine learning development and decision-making. The GDPR also does not describe the "appropriate safeguards" that must be in place to satisfy the exemption. The GDPR mentions pseudonymization as one such safeguard, but this appears to be illustrative rather than exhaustive and, as discussed above, pseudonymization itself is subject to some uncertainty.[195] Thus, while the statistical purpose exemption holds some promise, it may require additional clarification to become a practical source of flexibility for AI.

E. State Laws and AI: The California Consumer Privacy Act and Civil Liability

1. The California Consumer Privacy Act of 2018

Unlike Europe, the United States currently does not have a comprehensive privacy law at the federal level. Rather, U.S. privacy regulation has for the most part developed in a sector-specific manner, with the adoption of federal laws to regulate specific industries or data, such as the Health Insurance Portability and Accountability Act for health-related data.[196] In response to the GDPR and other recent developments, however, several U.S. states have recently introduced—and, in California's case, enacted—consumer privacy legislation that is the most comprehensive privacy law in the U.S. to date.

[194] *See id.* art. 5(1)(b), at 35 (providing that personal data processed for statistical purposes in accordance with art. 89(1) "shall not be considered be incompatible with the initial purposes"); *id.* ¶50 at 9 ("Further processing for . . . statistical purposes should be considered to be compatible lawful processing operations."). In addition, EU or Member State law may specify certain statistical purposes of processing; processing that occurs under such a law is exempt from certain GDPR requirements. *See id*. art. 89(2), at 85.

[195] *See id.* ¶156, at 29.

[196] *See* 45 C.F.R. Parts 160, 162, and 164.

The California Consumer Privacy Act of 2018 ("CCPA" or "Act"), which was enacted in June 2018, becomes effective on January 1, 2020.[197] Like the GDPR, the CCPA gives individuals certain rights with respect to their personal information. Under the CCPA, California consumers may prohibit the "sale" of their personal information by businesses, as well as request that businesses delete their personal information in certain cases.[198] Businesses regulated by the Act must disclose the information they have about individuals in response to verified requests from individuals.[199] The California Attorney General has primary enforcement authority, though individuals have a limited private right of action with respect to certain data breaches.[200] The CCPA applies only to personal information about California residents, but will have significant national impact given the scope of companies nationwide that meet the definition of a "business" within the statute. If additional states enact their own baseline privacy statutes—several states have considered such measures in the wake of the CCPA—additional compliance challenges could arise.

2. Understanding the CCPA's Scope: Personal Information

As a threshold matter, the CCPA's broad definition of "personal information" has potential implications for the development and use of AI. The Act applies to information "that identifies, relates to, describes, is capable of being associated with, or could reasonably be linked, directly or indirectly" to a person or household.[201] This can include identifiers, as well as network activity information like browsing history, geolocation data, biometric information, "thermal, olfactory, or

[197] Cal. Civ. Code §§ 1798.100 to 1798.199.

[198] *See* Cal. Civ. Code §§ 1798.105(a), 1798.120(a) (establishing rights of deletion and to opt out of sale, respectively).

[199] Cal. Civ. Code §§ 1798.100(a), (c), 1798.110(c)(5).

[200] *See* Cal. Civ. Code § 1798.150.

[201] Cal Civ. Code § 1798.140(o)(1).

similar information,"[202] and "inferences" drawn from data to create a "profile about a consumer."[203] This broad definition sweeps in data commonly used in AI tools. Specifically, data that is not "reasonably . . . linked,"[204] to a person in the offline environment can potentially be linked using AI. As a result, the CCPA considers as "personal information" a wide range of information potentially used for AI development.

AI may therefore expand the range of information that is considered personal information under the CCPA. Indeed, the capacity of AI to reveal associations between personal data elements may make it difficult for businesses that use AI to operate under the exemptions for aggregated and de-identified data. The Act provides that "aggregate consumer information" is information "from which individual consumer identities have been removed, that is not linked or reasonably linkable to any consumer of household, *including via a device*."[205] The Act's exception for de-identified information applies when a business using de-identified information (1) has implemented technical and process safeguards that prohibit re-identification; (2) has implemented businesses processes that specifically prohibit re-identification; (3) has implemented processes to prevent inadvertent release of de-identified information; and (4) makes no attempt to re-identify the information.[206] Aside from the difficulty of controlling what sophisticated AI "attempt[s]" to

[202] Cal. Civ. Code § 1798.140(o)(1)(H).

[203] Cal. Civ. Code § 1798.140(o)(1)(K).

[204] Cal. Civ. Code § 1798.140(o)(1).

[205] Cal. Civ. Code § 1798.140(a) (emphasis added).

[206] Cal. Civ. Code § 1798.140(h). The CCPA's de-identification standard is stricter than the three-part standard the FTC recommended in 2012. *See* FTC PRIVACY REPORT, *supra* note 105, at 21 ("[A]s long as (1) a given data set is not reasonably identifiable, (2) the company publicly commits not to re-identify it, and (3) the company requires any downstream users of the data to keep it in de-identified form, that data will fall outside the scope of the framework.").

do, AI's ability to discover connections between otherwise unlinked data may complicate compliance with provisions like these.

The Act also applies to certain publicly available information. While the CCPA exempts "publicly available" information from the definition of "personal information," the current definition of this term is narrow. "Publicly available" information currently is limited to information from federal state or local government records, and does not include information that is "used for a purpose that is not compatible with the purpose for which the information is maintained or made available."[207] Thus, the CCPA would regulate even the ability of AI developers to utilize publicly available datasets to train machine learning models.

3. Data Rights

The consumer rights under the CCPA have several implications for AI, particularly given the Act's broad "personal information" definition.

a. Right to Notice

The CCPA requires a covered business to notify consumers, "at or before the point of collection" about the categories of information collected from them and the purposes for which the business will use the data.[208] The Act defines "collection" broadly, to include most ways in which a business may acquire data.[209] It may be difficult for a company to specify at or before the point of collection the purposes for which the business will use the data in the context of analytics. Through their iterative processes, machine learning tools may learn and adapt and discover new purposes for the consumer data, creating transparency and potential compliance challenges.

[207] Cal. Civ. Code § 1798.140(o)(2). *But see* A.B. 874, 2019 Leg. (Cal. 2019), which would modify the definition of "personal information" to exclude "publicly available information," defined to mean "information that is lawfully made available from federal, state, or local government records."

[208] Cal. Civ. Code § 1798.100(b).

[209] *See* Cal. Civ. Code § 1798.140(e).

b. Right to Opt-Out of "Sales" of Personal Information

The CCPA gives customers the right to control and prohibit the "sale" and re-sale of their personal information.[210] Such restrictions may limit a company's ability to assemble and use the data sets necessary for algorithmic decision-making. The Act defines "sale" broadly to include the exchange of personal information for monetary or "other valuable consideration."[211] For example, exchanging personal information for access to a proprietary AI system could constitute a sale. The Act exempts from the "sale" definition instances where a business shares information with a service provider pursuant to appropriate notice, is necessary for a business purpose, and is not further sold.[212] The CCPA further prohibits third parties from selling personal information sold to them by a business unless the consumer receives explicit notice and an opportunity to opt out of the sale.[213]

c. Right to Access, Portability, and Deletion

The CCPA also grants individuals the right to access, transport, and delete their personal data, similar to the GDPR.[214] These rights may restrict a company's ability to leverage personal data for algorithmic decision-making, unless the data is anonymized or aggregated, such that it falls outside the definition of personal information.

d. Right to Non-Discrimination

The CCPA prohibits businesses from discriminating against consumers who exercise their rights by engaging in actions such as the denial of goods or services or charging different rates,

[210] Cal. Civ. Code § 1798.120(a).

[211] *See* Cal. Civ. Code § 1798.140(t).

[212] Cal. Civ. Code § 1798.140(t)(2)(C).

[213] Cal. Civ. Code § 1798.115(c).

[214] Cal. Civ. Code §§ 1798.100(a), (c), (d), 1798.105(a).

subject to exceptions.[215] The deletion right noted above, if exercised by a significant proportion of consumers, may be of particular concern to companies that develop and use AI decision-making from machine learning, because data deletion can skew the balance of representative data within the machine-learning models. Businesses otherwise might have sought to reward customers for not exercising their right to deletion or opt-out rights. While businesses will no doubt utilize the Act's multiple exceptions for these purposes, the provision potentially limits the ways in which businesses can do so.

e. Private Right of Action

The CCPA's private right of action for data breaches involving non-encrypted or non-redacted information, could expose companies that use large sets of sensitive personal information in connection with machine learning (e.g., to train models).[216] The Act contains minimum statutory damages for such actions, which may trigger class actions.[217] This risk is not unique to AI, however; any company covered by the CCPA has the same obligation to protect sensitive personal information.

F. Civil Liability

AI will undoubtedly pose challenging and significant questions about how the legal system should assign and apportion liability, identify relevant standards of care, and determine the scope of duty, among many others. The process of developing common law liability rules for AI is likely to take time. While the common law system readily adapts to moderate technological shifts, rapid and dramatic shifts strain the ability of courts to engage in the sort of analogic reasoning that fuels

[215] *See* Cal. Civ. Code § 1798.125.

[216] Cal. Civ. Code § 1798.150.

[217] *See id.* § 1798.150(a)(1)(A) (damages of not less than $100 and not more than $750 per customer per incident *or actual damages*, whichever is *greater*).

common law development,[218] often leading to a period of uncertainty before new consensus rules emerge.[219] For example, in an earlier era, courts considered whether defamation over radio or television constitute libel (which applies to writings and more stringent liability standards) or slander (which applies to the spoken word and is more difficult for plaintiffs to prove) and whether the damage from falling airplanes should be analogized to falling hot air balloons or automobiles.[220] Courts struggled with these questions for decades before resolving them (more or less) in favor of libel and automobiles.[221] With limited case law addressing the question of when and how fully automated technologies commit torts,[222] liability rules for actions taken on the basis of data analytics, including AI decisions supported by machine learning processes, remain in their infancy.

Moreover, it probably is not the case that the civil liability system will handle all autonomous technologies in one particular way. Robot doctors raise different concerns than robot

[218] *See generally,* Gary E. Marchant & Rachel A. Lindor, *The Coming Collision Between Autonomous Vehicles and the Liability System*, 52 SANTA CLARA L. REV. 1321 (2012).

[219] *See* Kyle Graham*, Of Frightened Horses and Autonomous Vehicles: Tort Law and its Assimilation of Innovations*, 52 Santa Clara L. Rev. 1241, 1241–42 (2012) ("In short, it takes time for an innovation, such as autonomous vehicles, to become fully assimilated within everyday tort law, and one rarely can anticipate the precise timetable for this process, or its final results.").

[220] *See id.* at 1252–56.

[221] *See id.*

[222] *See, e.g.,* U.S. CHAMBER INSTITUTE FOR LEGAL REFORM, TORTS OF THE FUTURE II: ADDRESSING THE LIABILITY AND REGULATORY IMPLICATIONS OF EMERGING TECHNOLOGIES 7 (2018), *available at* https://www.instituteforlegalreform.com/uploads/sites/1/tortsofthefuturepaperweb.pdf ("As autonomous robots and other products with AI make their way into the workplace, provide medical care in hospitals, operate on public highways, and serve us in our homes, hotels, and stores, they will be involved in incidents that result in personal injuries and other harms. Lawsuits inevitability will follow, and the legal system will wrestle with how to assign fault in order to compensate the injured person and allocate associated costs."); Donald G. Gifford, *Technological Triggers to Tort Revolutions: Steam Locomotives, Autonomous Vehicles, and Accident Compensation*, 11 J. TORT LAW at 64 ("Both industrial and surgical robots have caused injuries resulting in personal injury lawsuits against their manufacturers. Autonomous vehicles, however, pose the greatest challenge to preexisting tort law in the twenty-first century.").

vacuums, and it is therefore likely that the law will evolve, resulting in a variety of liability rules across different autonomous technologies.

Given these limitations, this report does not include an exhaustive discussion of how common law civil liability standards might regulate AI here. Rather, it highlights a few of the most important and interesting questions that might arise, using as an exemplar the next generation of a technology that has historically been of significant importance to the development of tort liability standards in the United States and is one of the most eagerly anticipated areas of autonomous technology deployment: the automobile.

The initial emergence of the automobile in the early 20th century is a paradigmatic example of how significant technological changes work their way through the civil liability system. As a leading scholar notes, "car wreck cases have produced many of the most significant tort cases of all time."[223] To give just two examples, in *MacPherson v. Buick Motor Co.*,[224] Justice Benjamin Cardozo was confronted with the mass production of automobiles and the sale of those automobiles through distributors who were not as well positioned to ensure the cars' safety as the manufacturers.[225] Cardozo thus "abolished the privity bar for product claims and opened the door" to the "entire field" of products liability.[226] Less than four years later, in a "second car wreck case out of the Empire State, [Cardozo] established the negligence per se doctrine for statutory violations."[227]

[223] Nora Freeman Engstrom, *When Cars Crash: The Automobiles Tort Law Legacy* 53 WAKE FOREST L. REV. 315, 293–336 (2018) ("In terms of doctrinal evolution, car wreck cases have produced many of the most significant tort cases of all time.").

[224] 217 N.Y. 382 (1916)

[225] *Id.*

[226] Engstrom, *supra* note 223, at 315.

[227] *Id.* (discussing *Martin v. Herzog*, 126 N.E. 814, 815 (N.Y. 1920)).

The anticipated prevalence of autonomous vehicles may very well produce a similar period of doctrinal development. Courts are just starting to confront cases concerning the civil liability of driverless cars, and as these cases become more frequent, they are likely to raise a number of challenging issues. [228]

An important and relevant question is how to apportion and assign liability for accidents caused (at least in part) by autonomous vehicles. Today, drivers are typically liable for their negligence, while manufacturers face product liability for design defects, manufacturing defects, and failure to warn. This arrangement fairly aligns with a world of semi-autonomous AI, and courts have apportioned liability in such cases.[229]

As AI takes over more functions, however, the distinction between operators and machines may blur.[230] Companies leading the development of autonomous vehicles have said that they will accept the liability if their cars crash when in "autonomous mode."[231] These representations, while perhaps clarifying, do not address the numerous questions that might arise about liability apportionment.

For example, what about when increasingly autonomous tools and humans interact? Will humans face liability because of their decisions to use a driverless car? Or, as autonomous vehicles

[228] *See, e.g.,* Nilsson v. General Motors LLC, No. 3:18-cv-00471 (N.D. Cal. 2018) (lawsuit brought by motorcyclist who was in an accident with a self-driving car, which was settled for an undisclosed amount); Holbrook v. Prodomax Automation Ltd., No. 1:17-cv-00219 (W.D. Mich. 2017) (lawsuit concerning the death of an auto parts worker killed by a robot in her workplace).

[229] *See* David C. Vladeck*, Machines Without Principles: Liability Rules and Artificial Intelligence*, 89 WASH. L. REV. 117, 120–21 (2014) (discussing surgical robots and airplane "autopilots").

[230] *See* Ryan Calo, *Robotics and the Lessons of Cyberlaw*, 103 CALIF. L. REV. 513, 515 (2015).

[231] *See* Nathan A. Greenblatt, *Self-Driving Cars Will Be Ready Before Our Laws Are*, IEEE SPECTRUM (Jan. 19, 2016, 12:00 PM), https://spectrum.ieee.org/transportation/advanced-cars/selfdriving-cars-will-be-ready-before-our-laws-are (discussing Volvo); Stephen Elmer, *Volvo, Google, and Mercedes to Accept Responsibility in Self-Driving Car Collisions*, AUTOGUIDE.COM (Oct. 7, 2015), https://www.autoguide.com/auto-news/2015/10/volvo-google-and-mercedes-to-accept-responsibility-in-self-driving-car-collisions.html.

improve, will there be circumstances in which *not* using a driverless car (or autonomous technology within a car) will expose an operator to liability? Conversely, will semi-autonomous vehicles become responsible for not preventing and anticipating human errors? Further, what about determining to what extent the designer of a tool is responsible for decisions the tool makes based on behaviors it "learns" from other sources? The fact that machine learning programming is often "the work of many hands" further complicates the task of apportioning liability and even distinguishing between various components of a product (*e.g.*, software and physical components).[232] How does the unpredictable course that AI decision-making based on machine learning models may take over time—as a result of limited training data, for example—affect assessments of foreseeability and overall liability under the negligence standard?

The pledge of autonomous vehicle manufacturers to assume liability also does not resolve a second important set of issues: how they will be liable. Many scholars argue that vehicle manufacturers could be held liable under a products liability theory (*e.g.*, design or manufacturing defect) that may carry with it strict liability.[233] Under such an approach, however, manufacturers of driverless cars would likely face greater exposure than would drivers facing negligence claims for similar accidents: products cases are more expensive to litigate (as they can require detailed analysis of design choices) and carry with them a greater risk of higher damages.[234] This could provide a disincentive for the development or adoption of autonomous vehicles, even if they are safer than vehicles driven by humans.[235]

[232] Jack B. Balkin, *The Path of Robotics Law*, 6 CALIF. L. REV. CIRCUIT 45, 53 (2015).

[233] *See, e.g.,* John Villasenor, *Products Liability and Driverless Cars: Issues and Guiding Principles for Legislation*, BROOKINGS (Apr. 24, 2014), https://www.brookings.edu/research/products-liability-and-driverless-cars-issues-and-guiding-principles-for-legislation/.

[234] *See* Greenblatt, *supra* note 231.

[235] *See generally id.* (arguing for a negligence standard for this reason); Ryan Abbott, *The Reasonable Computer: Disrupting the Paradigm of Tort Liability*, 86 GEO. WASH. L. REV. 1 (2018) (same).

Finally, if a negligence standard does prevail, this too will raise numerous questions. In traditional negligence, individuals are only responsible for the reasonably foreseeable consequences of their actions.[236] What does this concept mean as applied to autonomous decision-makers? What is foreseeable to them and, equally important, how will human decision-makers judge that? Will the "reasonable person" standard be replaced with that of a "reasonable robot"? In sum, from a civil liability standpoint, the questions about AI and legal enforcement are many and varied.

PART IV: USE OF SELF-REGULATION TO ADDRESS BIG DATA AND AI LEGAL AND ETHICAL CONCERNS

A. Introduction

Existing laws and regulations may address some privacy, consumer protection, and other concerns surrounding big data, but determining whether these technologies warrant additional, more specific regulation likely will involve a lengthy process. Entities are using data analytics in a broad range of contexts, based on myriad technological approaches, and across industries, to address a range of issues. Developing a self-regulatory framework that is flexible and promotes compliance within this broad ecosystem will take time.[237]

Still, consumers, advocates, and companies have identified concerns—some of which are discussed in prior chapters of this report—that they would prefer be addressed in the near term. Thus, to begin addressing these near-term concerns, even in a limited fashion, self-regulation may potentially serve as an alternative framework that provides companies with certain parameters that can be more quickly implemented and that do not unduly inhibit innovation.

[236] *See* Palsgraf v. Long Island R.R. Co., 162, N.E. 99, 100 (N.Y. 1928).

[237] Like Part III, this Part focuses on areas of data analytics that process personal information in relation to consumer-facing services. Applications of data analytics in other settings may involve different costs and benefits with self-regulation, and it is possible that self-regulation is not necessary or not viable in those settings.

B. The Potential Role of Self-Regulation

Self-regulation is a "concept that includes any attempt by an industry to moderate its conduct with the intent of improving marketplace behavior for the ultimate benefit of consumers."[238] As it relates to consumer protection and technology, self-regulation often is in the form of a "regulatory process whereby an industry-level organization (such as a trade association or a professional society), as opposed to a governmental or firm-level organization, sets and enforces rules and standards relating to the conduct of firms in the industry."[239]

In short, self-regulation complements existing laws by imposing supplemental rules to govern firm behavior. Self-regulation frameworks in the U.S. frequently are bolstered by federal or state regulators that provide a backstop to further promote accountability and enforce the self-regulatory regime, for example, by enforcing commitments that organizations make to follow certain practices. In addition to regulators, many self-regulatory organizations ("SROs") collaborate with other stakeholders, such as public interest groups, consumer advocates, and others, to develop and refine self-regulatory measures.

One example of a self-regulatory program from outside the data analytics arena includes the successful program for online behavioral advertising established by the Digital Advertising Alliance ("DAA").[240] In response to the FTC's call for self-regulation of online behavioral advertising ("OBA"), in 2009, the DAA adopted its code of conduct and online behavioral

[238] Maureen K. Ohlhausen, Chairwoman, Fed. Trade Comm'n, Success in Self-Regulation: Strategies to Bring to the Mobile and Global Era at the BBB Self-Regulation Conference (June 24, 2014), https://www.ftc.gov/system/files/documents/public_statements/410391/140624bbbself-regulation.pdf.

[239] DANIEL CASTRO, BENEFITS AND LIMITATIONS OF INDUSTRY SELF-REGULATION FOR ONLINE BEHAVIORAL ADVERTISING (2011), *available at* https://www.itif.org/files/2011-self-regulation-online-behavioral-advertising.pdf.

[240] DIG. ADVERT. ALL., SELF-REGULATORY PRINCIPLES FOR ONLINE BEHAVIORAL ADVERTISING (2009), *available at* http://www.aboutads.info/resource/download/seven-principles-07-01-09.pdf.

advertising principles, including the AdChoices icon to enhance consumers' ability to obtain more information about targeted ads based on their personal or device data, and choices regarding the use of their data.[241] As another example, in response to the anticipated growth in connected or smart-car technology, in 2014, 20 automakers pledged to meet or exceed commitments contained in the Automotive Consumer Privacy Protection Principles ("Principles"), which were developed to protect consumers' personal information collected through in-car technologies.[242]

The complex technology and speed of development of AI and big data may make the flexible and adaptable nature of industry self-regulation a particularly appropriate approach for regulating certain risks created by the technology. Similar to OBA and smart car technology, the complex technology and rapid development associated with big data and machine learning could be well served by a flexible and adaptable framework that supports compliance by industry participants and, in the process, helps foster trust and credibility with consumers and regulators.

Firms that use data analytic techniques, particularly those that rely on large amounts of consumer data for machine learning and decision-making, may recognize that their long-term commercial prospects depend on a more transparent relationship with consumers. In turn, this may motivate these firms to develop standards and principles that are welfare-enhancing overall and that engender greater protection of consumer data and privacy. A self-regulatory model that promotes transparency (as well as legal compliance and ethical uses of data) also enables consumers to make informed choices and assists government regulators in holding companies

[241] *See* YOURADCHOICES, https://youradchoices.com/.

[242] ALL. OF AUTO. MFRS., INC. & ASS'N OF GLOB. AUTOMAKERS, INC., CONSUMER PRIVACY PRINCIPLES: PRIVACY PRINCIPLES FOR VEHICLE TECHNOLOGIES AND SERVICES (2014), *available at* https://autoalliance.org/resources/consumer-privacy-protection-principles/

accountable.[243] Still, several factors relating to consumer protection and the viability of new entrants to the industry warrant caution when applying self-regulation to big data.

First and foremost, although there are certain risks that industry players may be better situated than governments to address, that is not true for all risks. As noted above, data analytics could have extremely significant implications for competition, the structure and nature of the workforce, national security, and even (at the extreme) maintain human control over technology. While industry can play a helpful role in addressing these issues, it is unlikely to be able to do so alone; the range of expertise needed to understand and assess the broadly distributed effects of data analytics are more likely to be available through traditional policy-making processes.

Second, companies may need to balance the opportunity to develop new lines of business with consumers' expectations and a broad view of consumer interests. Some companies may be inclined to advocate for a more high-level, principles-focused, and less enforcement-driven approach that may not align with regulators' expectations. Certain companies also may choose not to participate in the self-regulation framework altogether to avoid compliance costs, out of concern about potential antitrust liability, or to gain a commercial or competitive advantage by not committing to voluntary restrictions on their practices. Larger or embedded industry players may be inclined to advocate for self-regulatory regimes that create barriers to entry and impede competition, through rules for which compliance may be difficult for new entrants. Finally, in some instances, government regulators using a public notice and comment process may be better

[243] What constitutes an appropriate level of transparency could also vary sector-by-sector and may be best decided in a self-regulatory framework. For example, revealing the source code underlying an algorithm used to make sentencing recommendations in the criminal justice system may help ensure the protection of constitutional rights, the reduction of discrimination, and the guarantee of a fair process. However, revealing the source code underlying the detection of fake reviews on sites like Amazon may exacerbate the problem.

positioned to draft rules that balance the competing interests and diverse viewpoints of all stakeholders, including consumers, versus a narrower industry coalition.[244]

C. Developing and Applying Self-Regulatory Frameworks to Big Data and AI

Individual companies, industry organizations, consumer advocates, civil rights groups, academics, and governmental organizations, have put forth several sets of principles or ethical frameworks specific to AI. Although self-regulation in the big data and AI space could take many forms and may vary across industry sectors or within specific applications, the end products likely will reflect a common set of industry-agreed-upon principles. And while the key concepts that make up the core principles of current technology-focused self-regulatory frameworks—transparency, choice, accountability, accuracy, and security—address many of the primary areas of concern with AI and big data as well, AI self-regulatory efforts have some distinct features (including those discussed in the preceding sections) that warrant particular consideration.

1. AI Principles and Frameworks

In recent years, a number of stakeholders across the private sector, consumer protection, academia, and government have introduced AI principles or ethical frameworks to govern the use of AI in society.

[244] In Europe, proposals for co-regulation have gained some traction. *See, e.g.*, *Annex to the Communication from the Commission to the European Parliament, the European Council, the Council, the European Economic and Social Committee and the Committee of the Regions - Coordinated Plan on Artificial Intelligence*, at 8–9, COM (2018) 795 final (Dec. 7, 2018), (discussing the prospect of testing facilities for autonomous vehicles and other technologies, which might include “regulatory sandboxes (*i.e.*, areas where regulation is limited or favourable to testing new products and services) in selected areas where the law provides regulatory authorities with sufficient leeway, relaxing specific legal and regulatory requirements for the duration of the sandbox”).

a. Industry Groups

Members of the private sector have considered ethical issues with big data and AI.[245] For example, a number of industry groups have published statements on ethics in AI, including the Partnership on AI,[246] the Software & Information Industry Association ("SIIA"),[247] the Association for Computing Machinery ("ACM"),[248] and Institute of Electrical and Electronic Engineers ("IEEE") (which recently published *Ethically Aligned Design*).[249] These statements generally aim to identify the key ethical concerns arising from the development and use of AI systems, such as fairness, accountability, and transparency, and propose a voluntary set of principles to address such concerns.

Companies. In addition, certain companies have published principles or statements that apply to their own uses of AI and big data technology. These efforts and self-assessments may reflect companies' sense of responsibility toward consumers as well as their recognition that AI entails risks, including hostile public perception relating to certain AI investments, the potential for legislation, and concerns within companies' own workforces regarding the implications of AI for jobs and a variety of other topics. In particular, surveys reflect the concerns among senior leaders about the risk of machine learning contributing to the creation or spreading of false

[245] Thomas H. Davenport & Vivek Katyal, *Every Leader's Guide to the Ethics of AI*, MIT SLOAN MGMT. R. (Dec. 6, 2018), https://sloanreview.mit.edu/article/every-leaders-guide-to-the-ethics-of-ai/.

[246] *Tenets*, PARTNERSHIP ON AI, https://www.partnershiponai.org/tenets/.

[247] SOFTWARE & INFO. INDUS. ASS'N, ETHICAL PRINCIPLES FOR ARTIFICIAL INTELLIGENCE AND DATA ANALYTICS (2017), *available at* http://www.siia.net/Portals/0/pdf/Policy/Ethical%20Principles%20for%20Artificial%20Intelligence%20and%20Data%20Analytics%20SIIA%20Issue%20Brief.pdf?ver=2017-11-06-160346-990.

[248] Ass'n for Computing Machinery U.S. Pub. Policy Council, Statement on Algorithmic Transparency and Accountability (Jan. 12, 2017), *available at* http://www.acm.org/binaries/content/assets/public-policy/2017_usacm_statement_algorithms.pdf.

[249] IEEE, ETHICALLY ALIGNED DESIGN (1st ed. 2019).

information (or "fake news"), bias, unintended consequences, misuse of personal data, and the difficulty in explaining decisions by machine learning algorithms.[250]

b. Civil Rights Groups, Consumer Advocates, and Academics

Various civil rights groups, consumer advocates, and academics have published ethics principles and guidance for AI, including the Future of Life Institute (which published the Asilomar AI Principles),[251] the Center for Democracy & Technology ("CDT"),[252] The Leadership Conference,[253] Fairness, Accountability, and Transparency in Machine Learning ("FAT/ML"),[254] Privacy International,[255] and The Public Voice.[256] These principles similarly identify and seek to address ethical concerns arising from the development and use of AI, but they approach the issues from the perspective of individuals rather than industry. For example, they generally place a greater emphasis on individual rights and big data's overall impact on society.

Machine learning ethics theorists also have drawn on international human rights law.[257] For example, the EU's Independent High-Level Expert Group on Artificial Intelligence ("AI

[250] Of 1,400 U.S. executives knowledgeable about AI surveyed by Deloitte in 2018, 32% ranked ethical issues as one of the top three risks. *State of AI in the Enterprise*, DELOITTE (Oct. 22, 2018), *available at* https://www2.deloitte.com/insights/us/en/focus/cognitive-technologies/state-of-ai-and-intelligent-automation-in-business-survey.html.

[251] *Asilomar AI Principles*, FUTURE OF LIFE INSTITUTE, https://futureoflife.org/ai-principles/.

[252] *Digital Decisions*, CTR. FOR DEMOCRACY & TECH., https://cdt.org/issue/privacy-data/digital-decisions/.

[253] *Civil Rights Principles for the Era of Big Data*, THE LEADERSHIP CONFERENCE ON CIVIL & HUMAN RIGHTS (Feb. 27, 2014), https://civilrights.org/civil-rights-principles-era-big-data/.

[254] *Principles for Accountable Algorithms and a Social Impact Statement for Algorithms*, FAIRNESS, ACCOUNTABILITY, AND TRANSPARENCY IN MACHINE LEARNING, https://www.fatml.org/resources/principles-for-accountable-algorithms.

[255] PRIVACY INT'L, PRIVACY AND FREEDOM OF EXPRESSION IN THE AGE OF ARTIFICIAL INTELLIGENCE (2018), *available at* https://privacyinternational.org/sites/default/files/2018-04/Privacy%20and%20Freedom%20of%20Expression%20%20In%20the%20Age%20of%20Artificial%20Intelligence.pdf

[256] PUBLIC VOICE, UNIVERSAL GUIDELINES FOR AI (2018), *available at* https://thepublicvoice.org/ai-universal-guidelines/.

[257] The IEEE drew from: The Universal Declaration of Human Rights, the International Covenant on Civil and Political Rights, the Convention on the Rights of the Child, the Convention on the Elimination of all forms of

HLEG") has looked to fundamental rights enshrined in the EU Treaties,[258] the EU Charter, and international human rights law as the basis for its approach to AI ethics. These sources focus on core political and social rights or values, including respect for human autonomy, prevention of harms, and fairness, as the foundation of ethical AI principles.[259]

c. Government

Governmental data analytics frameworks typically address ethical and legal issues as well as innovation, investment, research, talent and workforce development, and economic growth.[260] Recent initiatives include legislative efforts and regulator-driven reports, including by the Obama and Trump White Houses,[261] European Commission,[262] EU data protection authorities,[263] the Independent High-Level Group on Artificial Intelligence convened by the European Commission,

Discrimination against Women, the Convention on the Rights of Persons with Disabilities, and the Geneva Conventions. *See* IEEE, *supra* note 249, at 19. *See, e.g.,* AMNESTY INTERNATIONAL & ACCESS NOW, THE TORONTO DECLARATION: PROTECTING THE RIGHT TO EQUALITY AND NON-DISCRIMINATION IN MACHINE LEARNING SYSTEMS (2018); HIGH-LEVEL EXPERT GROUP ON ARTIFICIAL INTELLIGENCE, *supra* note 186, at 6. *See, e.g.*, Karen Hao, *Establishing an AI Code of Ethics will be Harder than People Think*, MIT TECH. R. (Oct. 21, 2018), https://www.technologyreview.com/s/612318/establishing-an-ai-code-of-ethics-will-be-harder-than-people-think/; MARK LATONERO, GOVERNING ARTIFICIAL INTELLIGENCE: UPHOLDING HUMAN RIGHTS AND DIGNITY (2018), *available at* https://datasociety.net/wp-content/uploads/2018/10/DataSociety_Governing_Artificial_Intelligence_Upholding_Human_Rights.pdf.

[258] The EU constitutional documents commit to protect the fundamental and indivisible rights of human beings, to ensure respect for the rule of law, to foster democratic freedom and promote the common good. These rights are reflected in Articles 2 and 3 of the Treaty on European Union, and in the Charter of Fundamental Rights of the EU.

[259] HIGH-LEVEL EXPERT GROUP ON ARTIFICIAL INTELLIGENCE, *supra* note 186, at 12–13.

[260] *See* Tim Dutton, *An Overview of National AI Strategies*, MEDIUM (June 28, 2018), https://medium.com/politics-ai/an-overview-of-national-ai-strategies-2a70ec6edfd.

[261] *See, e.g.,* WHITE HOUSE BIG DATA REPORT, *supra* note 2; EXEC. OFFICE OF THE PRESIDENT, Executive Order on Artificial Intelligence, American AI Initiative (Feb. 11, 2019), *available at* https://www.whitehouse.gov/ai/.

[262] *AI Communication*, *supra* note 167.

[263] *See, e.g.*, INFO. COMM'R'S OFFICE, BIG DATA, ARTIFICIAL INTELLIGENCE, MACHINE LEARNING AND DATA PROTECTION, UK INFORMATION COMMISSIONER'S OFFICE (2017), *available at* https://ico.org.uk/media/for-organisations/documents/2013559/big-data-ai-ml-and-data-protection.pdf; COMMISSION NATIONALE DE L'INFORMATIQUE ET DES LIBERTES, HOW CAN HUMANS KEEP THE UPPER HAND? THE ETHICAL MATTERS RAISED BY ALGORITHMS AND ARTIFICIAL INTELLIGENCE (2017), *available at* https://www.cnil.fr/sites/default/files/atoms/files/cnil_rapport_ai_gb_web.pdf.

which included input from non-governmental stakeholders,[264]and the Organization for Economic Co-operation and Development.[265]

D. Summary of Common Principles

Although there is variation among the AI principles and frameworks released thus far, these documents generally share certain core principles, which we summarize here briefly:

Fairness. As described in other chapters, there is a concern that AI processes can lead to decisions and outcomes that are biased or discriminatory. AI principles recommend several measures to protect against these outcomes.[266] Testing is essential to detect where bias may be arising in an AI system for different protected groups, and testing techniques are increasingly used and useful, including for certification purposes.[267] In addition, technical tools, impact assessments, human intervention, or auditing may be used to address potential bias or review AI system outputs and decisions, as discussed in the next section.

Transparency/Explainability. The "opacity"[268] or "black box" aspect of machine learning algorithms[269] makes machine learning results difficult to explain in clear and comprehensible

[264] EUROPEAN COMM'N EUROPEAN GRP. ON ETHICS IN SCI. AND NEW TECHS., STATEMENT ON ARTIFICIAL INTELLIGENCE, ROBOTICS AND 'AUTONOMOUS' SYSTEMS (2018), *available at* https://ec.europa.eu/research/ege/pdf/ege_ai_statement_2018.pdf.

[265] ORG. FOR ECONOMIC CO-OPERATION AND DEV., RECOMMENDATION OF THE COUNCIL ON ARTIFICIAL INTELLIGENCE (May 21, 2019), *available at* https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0449.

[266] Moritz Hardt, Eric Price & Nathan Srebro, *Equality of Opportunity in Supervised Learning*, 30TH CONF. ON NEURAL INFO. PROCESSING SYS. (2017); Alexandra Chouldechova, *Fair Prediction with Disparate Impact: A Study of Bias in Recidivism Prediction Instruments* (Oct. 24, 2016) (unpublished manuscript), *available at* https://arxiv.org/abs/1610.07524.

[267] *See* Michael Feldman et al., *Certifying and Removing Disparate Impact*, PROCEEDINGS OF THE 21TH ACM SIGKDD INT'L CONF. ON KNOWLEDGE DISCOVERY AND DATA MINING (2015), *available at* http://sorelle.friedler.net/papers/kdd_disparate_impact.pdf.

[268] Jenna Burrell, *How the Machine 'Thinks': Understanding Opacity in Machine Learning Algorithms*, BIG DATA & SOC. (2016).

[269] *See* FRANK PASQUALE, THE BLACK BOX SOCIETY: THE SECRET ALGORITHMS THAT CONTROL MONEY AND INFORMATION (2015).

terms, which can present challenges to people who make decisions based on machine learning as well as individuals who are affected by machine learning-driven decisions.[270] This situation has led to a focus on creating greater transparency in how AI systems based on machine learning work while raising questions about how these goals can be achieved in a meaningful manner and consistent with appropriate protection for confidential business information.

Auditability. AI systems can and should be subject to internal or external audits to assess compliance with laws and adherence to ethical principles of fairness and non-discrimination, among other things. The party leading the audit might seek documentation to confirm the purpose and intent of the AI system, its architecture and performance, the datasets on which it is trained, records of outcomes, and the entity's internal processes to monitor for biases or test for compliance.

Accountability. Companies that develop or use AI will need to consider what types of accountability various AI systems offer. One approach, for example, in cases where AI-enabled decision-making has material impacts on individuals, is to provide individuals affected by AI decision-making the opportunity to review or appeal the outcome.

Accuracy/Reliability/Safety. A large volume of structured and unstructured data, compiled from disparate direct and indirect sources, can raise issues of data accuracy and relevance over time, with the potential to undermine the validity and credibility of the outcomes. The consequences of unreliable outcomes from using inaccurate data as an input are especially acute in certain applications within particular industries (e.g., AI-enabled medical decision-making, AI-enabled transportation technologies). Certifications or industry standards can play an important

[270] GOOGLE, PERSPECTIVES ON ISSUES IN AI GOVERNANCE, 8 *available at* https://ai.google/static/documents/perspectives-on-issues-in-ai-governance.pdf. While ultimately everything may be explainable, not everything can be explained using reasonable resources both to generate and understand explanations at scale.

role in establishing a baseline for reliability, for example, to validate that an AI system has sufficient procedures in place to ensure that data is appropriately used and retired.[271] Google, for instance, has proposed disclosing in certain contexts how confident users should be in the accuracy of an AI system's output, which could help users better assess the reliability of any results and understand the overall functioning of the AI system.[272]

Privacy and data security.[273] As described in in this report, fundamental privacy concepts rooted in existing laws or self-regulatory frameworks are applicable to AI and big data. However, it may be difficult to reconcile some of these concepts and (in some jurisdictions) legal requirements. For example, the concept of data minimization is in some tension with accumulating large volumes of data to train machine learning algorithms. Other concepts, such as access and correction, may be difficult to implement with the data used to develop and implement machine learning algorithms. Self-regulatory programs may be helpful in developing practical approaches to implementing privacy and data protection principles in these analytic models.

E. Practical Measures to Address Legal and Ethical Concerns

Companies that commit to a set of principles or a self-regulatory regime, or that more broadly seek to comply with applicable laws and regulations, will need to create internal policies and procedures, along with other measures as part of a formal compliance program.

[271] *Id.* at 18.

[272] *Id.* at 10. Such explanations also could change user behavior. For example, Google noted that "users told that a result was only 70% likely to be correct would be more careful in acting on it than if told it was 98% likely."

[273] *Microsoft AI Principles*, MICROSOFT, https://www.microsoft.com/en-us/ai/our-approach-to-ai; GOOGLE, *supra* note 270; *Asilomar AI Principles*, *supra* note 259; BOOZ | ALLEN | HAMILTON, ASSESSING THE ETHICAL RISKS OF ARTIFICIAL INTELLIGENCE (2017), *available at* https://www.boozallen.com/content/dam/boozallen_site/sig/pdf/publications/assessing-the-ethical-risks-of-artificial-intelligence.pdf.

Internal policies, and procedures. Companies can formalize their commitment to certain principles by creating internal policies and procedures that map the practical steps that a company will undertake to comply with the associated framework or legal obligations. As one example, in the consumer privacy space, companies may adopt a "privacy by design" approach that requires privacy as a core consideration at each stage of the product or service development life cycle, rather than treating privacy as an afterthought.[274] Companies that leverage big data can adopt a similar approach; many of the privacy, cybersecurity, ethics, transparency, fairness, and governance considerations should likewise be built into the "design" of data analytics systems.

Employee training. Companies will need to train applicable employees on the policies and procedures referenced above to educate and reinforce the need for compliance.

Audits or impact assessments. Internal audits and impact assessments are critical to gauge whether a company is complying with its policies and procedures, as well as whether the policies and procedures are operating effectively and in full compliance with evolving regulatory regimes and changes in technology. Such audits can be internally managed, or conducted by a third party using tools and frameworks tailored to the industry.[275] One tool that may be useful in the course of an audit is an impact assessment, which provides a formal, evidence-driven approach to evaluate the consequences of using a particular system, including data analytics implementations.[276]

Contractual protections. Companies that commit to a self-regulatory framework should consider whether and how to require their vendor partners to make the same commitments for

[274] *See* FTC PRIVACY REPORT, *supra* note 105 (advocating for a privacy by design model).

[275] THE INST. OF INTERNAL AUDITORS, GLOBAL PERSPECTIVES AND INSIGHTS: THE IIA'S ARTIFICIAL INTELLIGENCE AUDITING FRAMEWORK (2018), *available at* https://na.theiia.org/periodicals/Public%20Documents/GPI-Artificial-Intelligence-Part-III.pdf.

[276] DILLON REISMAN ET AL., ALGORITHMIC IMPACT ASSESSMENTS: A PRACTICAL FRAMEWORK FOR PUBLIC AGENCY ACCOUNTABILITY (2018), *available at* https://ainowinstitute.org/aiareport2018.pdf; *Principles for Accountable Algorithms and a Social Impact Statement for Algorithms*, *supra* note 254.

services that they provide to the company. To take an analogy from the online behavioral advertising space, NAI requires that its members contractually obligate their partners to certain notice and transparency commitments.[277] Companies should also ensure that the contracts clearly outline the parameters of use for data analytics purposes, what data can be leveraged for machine-learning purposes, ownership of the data, including the results of the machine-learning, and who is responsible for applicable privacy and consumer protections.

Due diligence and oversight. Vendor or supplier oversight also extends to data analytics partnerships, vendor relationships, acquisitions, and supply chain considerations. Companies looking to expand their capabilities likely will need to introduce products and technologies from different AI or data companies. For example, a company that uses AI to support its interactions with customers, such as through an AI chatbot, may purchase an "AI as a Service" (AIaaS) product from a separate provider. This same company may also hire a data analytics company to use big data techniques to detect sales patterns and assess the effectiveness of its marketing. Companies should put in place programs to assess their AI or big data suppliers, partners, acquisitions, and vendors, including due diligence prior to the engagement and ongoing oversight.

Advisory boards. Companies may create internal or external boards, committees, or councils made up of company employees and senior management, academics, and researchers to advise the company on trends that impact compliance strategies. Internal boards may be better positioned to influence company leadership and may have greater access to confidential information. External boards, on the other hand, may be more independent and objective. Both internal and external boards can be involved in high-level establishment of principles, pre-

[277] NETWORK ADVERTISING INITIATIVE, 2020 NAI CODE OF CONDUCT 11–12 (2019), *available at* https://www.networkadvertising.org/sites/default/files/nai_code2020.pdf.

deployment reviews of datasets or models, or post-deployment audits of systems, depending on the needs of the company. Ethics boards may also be able to provide on-demand ethical guidance during the data analytics development process.

Bias-detection and "explainability" tools. To address the fairness and explainability concerns described in the prior sections, several stakeholders have created tools to help companies assess and develop a more comprehensive understanding of their data analytics implementations. The development of such tools may help standardize and improve upon mechanisms for self-assessments and third-party assessments data.