

Rethinking Privacy in the AI Era

Jennifer King
Caroline Meinhardt

Policy Provocations for a Data-Centric World



Authors

Jennifer King is the Privacy and Data Policy Fellow at the Stanford University Institute for Human-Centered Artificial Intelligence (HAI). An internationally recognized expert in information privacy, her research examines the public's understanding and expectations of online privacy as well as the policy implications of emerging technologies, including artificial intelligence. Her recent research explores alternatives to notice and consent (with the World Economic Forum), the impact of California's new privacy laws, and manipulative design (dark patterns). She also co-directs the [Dark Patterns Tip Line](#) repository at Stanford. Prior to joining HAI, she was the Director of Consumer Privacy at the Center for Internet and Society at Stanford Law School from 2018 to 2020. Dr. King completed her doctorate in information management and systems (information science) at the University of California, Berkeley School of Information.

Caroline Meinhardt is the policy research manager at the Stanford Institute for Human-Centered Artificial Intelligence (HAI), where she develops and oversees policy research initiatives. She is passionate about harnessing AI governance research to inform policies that ensure the safe and responsible development of AI around the world—with a focus on research on the privacy implications of AI development, the implementation challenges of AI regulation, and the governance of large-scale AI models. Prior to joining HAI, Caroline worked as a China-focused consultant and analyst, managing and delivering in-depth research and strategic advice regarding China's development and regulation of emerging technologies including AI. She holds a Master's in International Policy from Stanford University, where her research focused on global governance solutions for AI, and a Bachelor's in Chinese Studies from the University of Cambridge.

Acknowledgments


The authors would like to thank Brenda Leong, Cobun Zweifel-Keegan, Justin West, Kevin Klyman, and Daniel Zhang for their valuable feedback, Nicole Tong and Cole Ford for research assistance, and Jeanina Casusi, Joe Hinman, Nancy King, Shana Lynch, Carolyn Lehman, and Michi Turner for preparing the publication.

Disclaimer

The Stanford Institute for Human-Centered Artificial Intelligence (HAI) is a nonpartisan research institute, representing a range of voices. The views expressed in this White Paper reflect the views of the authors.

Table of Contents

Authors	2
Acknowledgments	2
Table of Contents	3
Executive Summary	4
Chapter 1: Introduction	5
Chapter 2: Data Protection and Privacy: Key Concepts and Regulatory Landscape	7
a. Fair Information Practice Principles: The framework behind data protection and privacy	9
b. General Data Protection Regulation: The “global standard” for data protection	10
c. U.S. State Privacy Laws: Filling the federal privacy vacuum	12
d. Predictive AI vs. Generative AI: An inflection point for data protection regulation	14
Chapter 3: Provocations and Predictions	17
a. Data is the foundation of AI systems, which will demand ever greater amounts of data	17
b. AI systems pose unique risks to both individual and societal privacy that require new approaches to regulation	19
c. Data protection principles in existing privacy laws will have an implicit, but limited, impact on AI development	22
d. The explicit algorithmic and AI-based provisions in existing laws do not sufficiently address privacy risks	25
e. Closing thoughts	29
Chapter 4: Suggestions for Mitigating the Privacy Harms of AI	31
Suggestion 1: Denormalize data collection by default	33
Suggestion 2: Focus on the AI data supply chain to improve privacy and data protection	36
Suggestion 3: Flip the script on the management of personal data	41
Chapter 5: Conclusion	45
Endnotes	46



Executive Summary

In this paper, we present a series of arguments and predictions about how existing and future privacy and data protection regulation will impact the development and deployment of AI systems.

Data is the foundation of all AI systems. Going forward, AI development will continue to increase developers' hunger for training data, fueling an even greater race for data acquisition than we have already seen in past decades.

Largely unrestrained data collection poses unique risks to privacy that extend beyond the individual level—they aggregate to pose societal-level harms that cannot be addressed through the exercise of individual data rights alone.

While existing and proposed privacy legislation, grounded in the globally accepted Fair Information Practices (FIPs), implicitly regulate AI development, they are not sufficient to address the data acquisition race as well as the resulting individual and systemic privacy harms.

Even legislation that contains explicit provisions on algorithmic decision-making and other forms of AI does not provide the data governance measures needed to meaningfully regulate the data used in AI systems.

We present [three suggestions](#) for how to mitigate the risks to data privacy posed by the development and adoption of AI:

1. Denormalize data collection by default by shifting away from opt-out to opt-in data collection.

Data collectors must facilitate true data minimization through “privacy by default” strategies and adopt technical standards and infrastructure for meaningful consent mechanisms.

2. Focus on the AI data supply chain to improve privacy and data protection. Ensuring dataset transparency and accountability across the entire life cycle must be a focus of any regulatory system that addresses data privacy.

3. Flip the script on the creation and management of personal data. Policymakers should support the development of new governance mechanisms and technical infrastructure (e.g., data intermediaries and data permissioning infrastructure) to support and automate the exercise of individual data rights and preferences.

Chapter 1: Introduction

In the opening months of 2024, artificial intelligence (AI) is squarely in the sights of regulators around the globe. The European Union is set to finalize its AI Act later this year. Other parts of the world, from the United Kingdom to China, are also contemplating and, in some cases already implementing, wide-ranging AI regulation. In the United States, a recent milestone Executive Order on AI marked the clearest signal yet that the Biden administration is poised to take a comprehensive approach to AI governance.¹ With federal legislation to regulate AI yet to pass, a growing number of federal agencies and state legislators are clarifying how existing regulation relates to AI within their jurisdictional areas and proposing AI-specific regulation.²

While much of the discussion in the AI regulatory space has centered on developing new legislation to directly regulate AI, there has been comparatively little discourse on the laws and regulations that already impact many forms of commercial AI. In this white paper, we focus on the intersection of AI regulation with two specific areas: privacy and data protection legislation. The connective tissue between privacy and AI is data: Nearly all forms of AI require large amounts of training data to develop classification or decisional capabilities. Whether or not an AI system processes or renders decisions about individuals, if a system includes personal information, particularly identifiable personal information, as part of its training data, it is likely to be subject—at least in part—to privacy and data protection regulations.

We make a set of arguments and predictions about how existing and future privacy and data protection regulations in the United States and the EU will impact the development and deployment of AI systems. We

start with the fundamental assumption that AI systems require data—massive amounts of it—for training purposes. It is this need for data, as best evidenced by data-hungry generative AI systems such as ChatGPT, that we predict will fuel an even greater race for data acquisition than we’ve witnessed over the last decades of the “Big Data” era. This need will in turn impact both individual and societal information privacy—not just through the demand for data, but also by the impacts this need will have on specific issues such as consent, provenance, and the entire data supply pipeline and life cycle more generally.³

We move on to examining AI’s unique risks to consumer and personal privacy, which—unlike many technology-fueled privacy harms that primarily impact individuals—aggregate to pose societal-level risks that existing regulatory privacy frameworks are not designed to address. We argue that existing governance approaches, which are based predominantly on the globally accepted Fair Information Practices (FIPs), will not be sufficient to address these systemic privacy risks. Finally, we close with suggested solutions for mitigating these risks while also offering new directions for regulation in this area.

What’s at Stake: The Future of Both Privacy and AI

Data is a key component for all AI systems—to date, the most significant improvements in AI systems have been tied to access to very large amounts of training data. This fact does not necessarily mean that all advancements in AI will require massive amounts of data; as we discuss later, some researchers are observing quality versus quantity trade-offs

that indicate more may not reliably mean better. Regardless, we are presently at an inflection point where there is considerable pressure for companies to build massive training datasets to maintain their competitive advantage.

A primary concern motivating this paper is that despite the fact that existing and proposed privacy and data protection laws on both sides of the Atlantic will have an impact on AI, they will not sufficiently regulate the data sources that AI systems require in a way that will substantively preserve, or even improve, our data privacy. In this paper, we explore several related concerns:

1. The framework that underlies data protection laws has weaknesses that will not give individuals the tools they need to preserve their data privacy as AI advances;
2. It also fails to address societal-level privacy risks;
3. Policymakers must expand the scope of how we approach privacy and data protection to address these weaknesses and bolster data privacy in an increasingly AI dominant world.

We start from the assumption that for most of us the current state of our data privacy ranges from suboptimal to dismal. In the United States, polls have shown that the public largely feels as if they have no control over the data that is collected about them online;⁴ that the benefits they receive in exchange for their data are not always worth the bargain of free access; and that in most data relationships, consumers have no ability to negotiate more favorable terms—and in many instances, believe they are locked in or have few if any alternatives.⁵

In short, as we move toward a future in which AI development continues to increase demands for data, data protection regulation that at best maintains the status quo does not inspire confidence that the data rights we have will preserve our data privacy as the technology advances. In fact, we believe that continuing to build an AI ecosystem atop this foundation will jeopardize what little data privacy we have today.

This paper focuses on the core issues that we believe require the most attention to address this state of affairs. It does not claim to address or solve everything. But we do believe that if these issues aren't sufficiently acknowledged and addressed through regulation and enforcement, we leave ourselves open to a situation where privacy protection continues to deteriorate. There are many worries attached to how our world will change as it continues to embrace AI. Concerns related to bias and discrimination have already generated extensive debate and discussion, and we argue that a substantial loss of data privacy is another major risk that deserves our heightened concern.

Chapter 2: Data Protection and Privacy: Key Concepts and Regulatory Landscape

The last two years have seen groundbreaking advances in AI, a period in which generative AI tools became widely available, inspiring and alarming millions of people around the world. Large language models (LLMs) such as GPT-4, PaLM, and Llama, as well as AI image generation systems such as Midjourney and DALL-E, have made a tremendous public splash, while many other less headline-grabbing forms of AI also continued to advance at breakneck speed.

While recognizing the recent dominance of LLMs in public discourse, in this paper we consider the data privacy and protection implications of a wider array of AI systems, defined more broadly as “engineered or machine-based system[s] that can, for a given set of objectives, generate outputs such as predictions, recommendations, or decisions influencing real or virtual environments.”⁶ For example, we consider a range of predictive AI systems, such as those based on machine learning, that analyze vast amounts of data to make classifications and predictions, ranging from facial recognition systems to hiring algorithms, criminal sentencing algorithms, behavioral advertising and profiling, and emotion recognition tools, to name a few. These systems operate with varying levels of autonomy, with “automated decision-making” referring to AI systems making decisions (such as awarding a loan or hiring a new employee) without any, or minimal, human involvement.⁷ While generative AI systems also rely on predictive processes, those systems ultimately focus on creating new content ranging from text to images, video, and audio as their output.

While some policymakers are keen to demonstrate that they are assuaging the public’s growing concerns about the rapid development and deployment of AI by introducing new legislation, there is a growing debate over whether existing laws provide sufficient protection and oversight of AI systems.

In response to these widely publicized developments, both policymakers and the general public have called for regulating AI technologies. Since 2020, countries around the world have begun passing AI-specific legislation.⁸ While the EU finalizes the parameters of its AI Act, the bloc’s attempt to provide overarching regulation of AI technologies, the United States presently lacks a generalized approach to AI regulation, though multiple federal agencies have released policy statements asserting their authority over AI systems that produce outputs in violation of existing law, such as civil rights and consumer protection statutes.⁹ Several U.S. states and municipalities have also tackled general consumer regulation of AI systems.¹⁰

While some policymakers are keen to demonstrate that they are assuaging the public’s growing concerns about the rapid development and deployment of AI by introducing new legislation, there is a growing debate over whether existing laws provide sufficient protection and oversight of AI systems. As we discuss in this white paper, privacy and data protection laws in the United States and the EU already do the work of regulating some—though not all—aspects of AI. Whether these existing laws, and proposed ones based on these frameworks, are adequate to anticipate and

respond to emergent forms of AI while also addressing privacy risks and harms is a question we will address later in this paper.

Before we delve into the details of our arguments, we provide a brief overview of the present state of data protection and privacy regulations in the EU and the United States that impact AI systems, starting with the foundational Fair Information Practices (FIPs). *Those familiar with these regulations may wish to skip ahead to the next chapter.*

Data Privacy and Data Protection

Data privacy and data protection are sometimes used interchangeably in casual conversation. While these terms are related and have some overlap, they differ in significant ways.

Data privacy is primarily concerned with who has authorized access to collect, process, and potentially share one’s personal data, and the extent to which one can exercise control over that access, including by opting out of data collection. The term’s scope is fairly broad, as it pertains not just to personal data but to any kind of data that, if accessed by others, would be seen as infringing on one’s right to a private life and personal autonomy.

Privacy is often described in terms of personal control over one’s information, though this conception has been challenged by the increasing loss of control that many have over their data. But it is this notion of personal control that underlies both existing privacy regulations and frameworks. What is considered “private” is also contextually contingent, in that data shared in one context may be viewed as appropriate by an individual or data subject (e.g., sharing one’s real time location data with a friend) but not in another (e.g., a third party collecting one’s real time location data and using it for advertising purposes without explicit permission). The relational nature of data has also challenged the idea of privacy as personal control, as data that is social in nature (e.g., shared social media posts) or data that can reveal both biological ties and ethnic identities (e.g., genetic data) continue to grow.

Data Privacy and Data Protection (cont'd)

Data protection refers to the act of safeguarding individuals' personal information using a set of procedural rights, which includes ensuring that data is processed fairly, for specified purposes, and collected on the basis of one of six accepted bases for processing.¹¹ Consent is the strictest basis and allows individuals to withdraw it after the fact. By contrast, legitimate interest provides the greatest latitude—this legal ground for processing data allows processors to justify data processing on the basis of this data being needed to carry out tasks related to their business activity. Data processors must still respect individuals' fundamental data protection rights, such as providing notice when data is collected, giving access to one's collected information, providing the means to correct errors, delete, or transfer it (data portability) to other processors, and affording the right to object to the processing itself. But there is a bias toward accepting as a given the collectibility of some forms of personal data by default.

The EU formally distinguishes between personal privacy (i.e., respect for an individual's private life) and data protection, enshrining each in its European Charter of Fundamental Rights. Nevertheless, there are areas of overlap and the concepts complement each other. When data protection principles do not apply because the collected information is not personal data (e.g., anonymized body scanner data), the fundamental right to privacy applies as the collection of bodily information affects a person's individual autonomy. Conversely, data protection principles can ensure limits on personal data processing, even when such processing is not thought to infringe upon privacy.¹²

a. Fair Information Practice Principles: The framework behind data protection and privacy

Most modern privacy legislation, at its core, is based on the Fair Information Practices (FIPs), a 50-plus-year-old set of principles that are accepted around the globe as the fundamental framework for providing individuals with due process rights for their personal data.¹³ Proposed as a U.S. federal code of fair information practices for automated personal data systems in the early 1970s, the FIPs introduced five

safeguard requirements regarding personal privacy as a means of ensuring “informational due process.”¹⁴ They focus on the obligations of record-keeping organizations to allow individuals to know about, prevent alternative uses of, and correct information collected about them.¹⁵ As policy expert Mark MacCarthy describes, “All these measures worked together as a coherent whole to enforce the rights of individuals to control the collection and use of information about themselves.”¹⁶

Rather than framing information privacy as a fundamental human right, as both the United Nations Universal Declaration of Human Rights and the

European Charter of Fundamental Rights do with a more general conception of privacy, the FIPs outline a set of rules and obligations between the individual (data subject) and the record-keeper (data processor).¹⁷ The FIPs were drafted around a core assumption that the state has a legitimate need to collect data about its citizens for administrative and record-keeping purposes.¹⁸ This assumption—that data collection is necessary and appropriate for the workings of the modern state but must be done fairly and with procedural safeguards in place—was incorporated into subsequent revisions of the FIPs, even as they were increasingly applied to the private sector.

The most internationally influential version, developed by the Organisation for Economic Cooperation and Development (OECD) in 1980 and amended in 2013, consolidates and expands the original FIPs into eight principles covering collection limitation, data quality, purpose specification, use limitation, security safeguards, openness, individual participation, and accountability.¹⁹ The guidelines reflect a broad international consensus on how to approach privacy protection that has translated into a policy convergence around enshrining the FIPs as a core part of information privacy legislation around the world.²⁰

Despite having been conceived long before the emergence of the commercial internet, let alone social media platforms and generative AI tools, core components of the FIPs, such as data minimization and purpose limitation²¹, directly impact today’s AI systems by limiting how broadly companies can repurpose data collected for one context or purpose to create or train new AI systems. The EU’s General Data Protection Regulation (GDPR), as well as California’s privacy regulations and the proposed American Data Privacy and Protection Act (ADPPA), relies heavily on these principles. These regulations’ attempts to clarify

The FIPs were drafted around a core assumption that the state has a legitimate need to collect data about its citizens for administrative and record-keeping purposes.

the application of the FIPs to privacy controls amid exponentially increasing volumes of online consumers and commercial data shed further light on the impact of privacy regulation on AI.

b. General Data Protection Regulation: The “global standard” for data protection

Passed in 2016 and in effect as of 2018, the General Data Protection Regulation is the EU’s attempt to both update the 1995 Data Protection Directive and harmonize the previous patchwork of fragmented national data privacy regimes across EU member countries and to enable stronger enforcement of Europeans’ data rights.²² At its core, the GDPR is centered on personal data, which is defined as “any information relating to an identified or identifiable natural person.”²³ It grants individuals (“data subjects”) rights regarding the processing of their personal data, such as the right to be informed and a limited right to be forgotten, and guides how businesses can process personal information. It is arguably the most significant data protection legislation in the world today, spurring copycat legislation and impacting the framing of data protection around the globe. As a result of the GDPR’s direct applicability to AI and its dominance across

the globe, data protection and privacy concerns are largely absent from the EU’s AI Act.

The GDPR contains several provisions that apply to AI systems, even though it does not specifically include the term “artificial intelligence.” Instead, Article 22 provides protections to individuals against decisions “based solely on automated processing” of personal data without human intervention, also called automated decision-making (ADM).²⁴ It enshrines the right of individuals not to be subject to ADM where these decisions could produce an adverse legal or similarly significant effect on them. Given the widespread use of ADM as it relates to health, loan approvals, job applications, law enforcement, and other fields, the article plays a crucial role in enforcing

a minimum degree of human involvement in such decision-making processes.

Beyond Article 22, the GDPR also puts in place several key data protection principles that affect AI systems (see table). Most notably, the *purpose limitation* principle forbids the processing of personal data for purposes other than those specified at collection, and the *data minimization* principle restricts the collection and retention of data to that which is absolutely necessary. These principles, in theory, curb unfettered personal data collection (or data mining) that is common for data-intensive AI applications. Despite the commonly held assumption that more data always makes for better AI, and that such constraints on data collection and use will hamper progress in AI, there is

Core Data Protection Principles

Data Protection Principle	Summary of Relevance
Data Minimization	Defined in Article 5 of the GDPR as ensuring that collected data is “adequate, relevant and limited to what is necessary in relation to the purposes for which they are processed.” This principle prescribes proportionality: Data processors should not collect as much data as possible, particularly out of the context provided for collection. The intent is to prevent data collectors from engaging in indiscriminate data collection.
Purpose Limitation	Defined in Article 5 as data “collected for specified, explicit and legitimate purposes and not further processed in a manner that is incompatible with those purposes.” This principle emphasizes the importance of context, restricting uses of data beyond the explicit purpose given at collection. If a data processor wishes to repurpose collected data, they need to seek consent for that new use.
Consent	Defined in Article 7 and Recital 32 as a key requirement for data processing. Consent must be “given by a clear affirmative act establishing a freely given, specific, informed and unambiguous indication of the data subject’s agreement to the processing of personal data relating to him or her, such as by a written statement, including by electronic means, or an oral statement.” Notably, consent is required for all processing, including if data is collected for multiple purposes. Recital 42 describes the burden of proof data processors must meet to prove data subject consent, noting that “[c]onsent should not be regarded as freely given if the data subject has no genuine or free choice or is unable to refuse or withdraw consent without detriment.”

extensive research demonstrating that building ADM systems within these constraints is feasible and even desirable.²⁵

The GDPR also enshrines transparency obligations in the form of rules about giving notice to individuals when their personal information is processed for the purpose of profiling or ADM.²⁶ It further establishes rules granting individuals the *right to access* their own data and ensure the accuracy of the data processing. Finally, it introduces Data Protection Impact Assessments (DPIA)—an accountability measure that requires the collecting organization assess the potential risks and harms of data processing activities (as they pertain to the relevant organization but also potential societal-level harms) prior to conducting them.²⁷

c. U.S. State Privacy Laws: Filling the federal privacy vacuum

As of 2024, the United States still lacks a federal omnibus consumer privacy law similar to the GDPR. The closest it has come to passing consumer privacy regulation is the American Data Privacy and Protection Act (ADPPA), which was introduced in the House in 2022 but did not advance to a floor vote in that session and has yet to be reintroduced.²⁸ Similar to the GDPR, the ADPPA would have imposed limits on the “collection, use, and sharing of personal information, requiring that such a process be “necessary and proportionate.” It would acknowledge the connection between information privacy and civil rights, strengthening relevant civil rights laws and essentially enacting the privacy section of the Biden administration’s subsequent “Blueprint for an AI Bill of Rights.”²⁹ ADPPA was the result of lengthy bipartisan negotiations and future privacy legislation is likely to hew closely to the original 2022 bill.

[The purpose limitation and data minimization] principles, in theory, curb unfettered personal data collection (or data mining) that is common for data-intensive AI applications.

In the absence of consumer-specific federal legislation, several sectoral laws have created a patchwork of privacy protections over the decades, such as the Family Educational Rights and Privacy Act (FERPA), the Children’s Online Privacy Protection Act (COPPA), the Health Insurance Portability and Accountability Act (HIPAA), and even the Video Privacy Protection Act (VPPA), to name a few. In this splintered landscape, U.S. states have been passing their own consumer privacy laws. As of 2023, 12 states have passed consumer privacy regulations, though California’s Consumer Privacy Act (CCPA) remains the most far-reaching.³⁰ For that reason, we will focus on the CCPA for discussion purposes.

Sometimes dubbed California’s version of the GDPR, the CCPA—together with its 2022 update, the California Privacy Rights Act (CPRA)—is arguably the most significant state-level effort so far to enact both stringent and broad consumer privacy protections.³¹ While some scholars have argued that the CCPA consciously creates a fundamentally different data privacy regime for California than the GDPR, it nevertheless marks a landmark shift in the U.S. privacy regulation debate.³²

The initial version of the CCPA created rights of data access, deletion, and portability, as well as a right to opt out of sales of personal data for two-year cycles, and a purpose limitation provision. Businesses are obliged to provide notice of the types of data they collect, to obtain opt-in consent for data collection from children ages 13 to 16, and to abide by purpose limitations when collecting and using or reusing data, which must be consistent with individuals' general expectations and the purpose specified upon collection. The subsequent CPRA, passed as a ballot proposition (Proposition 24), amends the CCPA to add a data minimization prong as well as a right to correct personal data, a right to opt out of processing categories of sensitive personal data, and—similar to the GDPR—a right to opt out of some forms of ADM (those with significant effects, such as on housing and employment), which in draft regulations has been interpreted by California's privacy regulator

to include AI systems.³³ Businesses must conduct privacy risk assessments and cybersecurity audits, offer alternatives for accessing services for those who opt out, and cannot discriminate against consumers for exercising these rights.

A notable difference between California's privacy regime and other states is that California remains the only state to have created an enforcement agency (the California Privacy Protection Agency, or CPPA) with rulemaking authority, rather than delegating this function to the state's attorney general's office, as many such laws do. In practice, this may mean that the CPPA has more in-house expertise than most state attorneys general and latitude to both engage in proactive enforcement via published guidance and tackle complex and emergent issues at the intersection of AI and personal data.

Beyond the EU and United States: Data Protection in China

In 2021, China's legislature followed the EU's example by promulgating a comprehensive and stringent data privacy law. Heavily inspired by the GDPR, China's Personal Information Protection Law (PIPL) was designed to give Chinese citizens control over their personal and sensitive data by delineating who can access, process, and share their information.³⁴ As such, it incorporates many elements of the FIPs, including data collection limitations, purpose specification requirements, and use limitations.

Despite commonly being referred to as a privacy law, the PIPL never directly mentions privacy but instead focuses on curbing the abuse and mishandling of personal information—theoretically by both corporate and state actors, though practically the state's ability to surveil its citizens remains unchecked.³⁵ Like the GDPR and the CCPA, the law contains explicit provisions banning automated decision-making that enables differential treatment of consumers, including price discrimination. More broadly, it introduces limits on what was largely unfettered data collection by data-hungry AI companies, requiring informed consent for all kinds of data-processing activities and granting individuals key rights over their data, including the right to amend, delete, and request copies of information collected about them.

Since the PIPL predominantly acts as a framework law that sets out broad principles and requirements, it was followed by a string of more granular implementing regulations, which have been directly impacting

Beyond the EU and United States: Data Protection in China (cont'd)

AI companies, particularly those with facial recognition products.³⁶ However, the true impact of the PIPL on China's AI ecosystem remains hard to assess given the government's tendency to use it as a political tool. For example, in 2022 when China's ride-hailing giant Didi was fined by the government following a comprehensive cybersecurity review, the regulatory decision cited the PIPL and Didi's illegal collection of data, including facial recognition data.³⁷ However, the unprecedented size of the fine and opaque application of a variety of laws and regulations may point to the PIPL being used as a tool to control the country's tech giants.³⁸

d. Predictive AI vs. Generative AI: An inflection point for data protection regulation

Until generative AI systems broke through the public and policymaker consciousness in late 2022, discussions about AI regulation were focused on predictive AI systems that use data to classify, sort, and predict outcomes. Within the scope of predictive AI, concerns focused primarily on the outputs produced by these systems, with less focus on the data used to train them. Both policy discussions and proposed regulation for AI were primarily concerned with algorithmic audits³⁹ and impact assessments,⁴⁰ transparency and explainability,⁴¹ and enforcing civil rights⁴² as a means of ensuring decisional outputs were fair and unbiased.⁴³ To the extent that privacy played a role in these discussions, concerns were typically related to the growing awareness of our main argument in this paper—that existing privacy laws such as the GDPR would impact aspects of AI development and that passing AI regulation without comprehensive privacy legislation, as would currently be the case in the United States, would be a job half-finished.⁴⁴

It is not an overstatement to say that generative AI substantially shifted the terms of the debate. Awe

over the capabilities of image generators such as DALL-E or Midjourney and LLMs such as ChatGPT simultaneously raised questions about how these systems were built and what data was used to power them. As it became more widely understood that generative systems are built predominantly on data scraped from across the internet, concerns mounted about exactly what data—and whose data—was powering these systems.⁴⁵

These weren't novel concerns. Facial recognition software company Clearview AI had already raised the ire of privacy and civil liberties advocates, as well as European policymakers, for their aggressive acquisition of facial images to power their predictive criminal suspect identification app. Clearview built their software by scraping image data from across the internet, including from online services that explicitly prohibit such scraping. But given Clearview's niche product (available only to law enforcement organizations) and targeted impact (used to identify criminal suspects), their data use wasn't widely discussed, despite extensive reporting on the company by Kashmir Hill of *The New York Times*.⁴⁶ Clearview has virtually been shut out of the EU marketplace after its data-gathering practices were found to be in gross violation of the GDPR.⁴⁷ In the United States, a 2020 lawsuit by the American Civil Liberties Union

leveraging the state of Illinois' Biometric Information Privacy Act resulted in a settlement that prohibits the company from making its products available to individuals and companies across the country, as well as also prohibiting use of its products by law enforcement agencies in Illinois.⁴⁸

Meanwhile, as generative AI systems gained greater exposure, privacy regulators around the world scrambled to understand the impacts of these systems on the public and whether they violated existing laws.⁴⁹ The G7 data protection authorities went so far as to issue a group statement summarizing their concerns—specifically calling out the legal authority generative systems may have for processing personal information, especially related to children; the potential for generative systems to be used for attacks to extract personal information; and the need to produce compliance documentation about the life cycle of the data used to develop and train their models. The statement also called for “privacy by design,” the practice of taking privacy into account throughout all stages of system development, while reiterating the

need for developers to respect data protection rights and the data minimization principle.⁵⁰

The Italian data protection authority went so far as to ban ChatGPT until OpenAI, its creator, put specific practices in place (see below). The fact that many generative systems are built at least in part on scraped data raises questions about whether and under what contexts data-scraping practices can be compliant with the GDPR, particularly when personally identifiable data is scraped and included in training data, even if that data is publicly available. In particular, it may place consent and legitimate interest at odds, as companies like Clearview argue (albeit unsuccessfully in this instance) that they do not need consent for publicly accessible data.⁵¹ Generative systems raise other crucial questions about training data, such as the extent to which procedural data rights will apply to them, if individuals can request to delete their data from training datasets or object to this form of processing, and whether any of this will depend on the context of use of the generative application in making these determinations.

Italy Scrutinizes ChatGPT's Data Practices

On March 20, 2023, the Italian Data Protection Authority (the Garante) received a report that OpenAI—the company that developed GPT-4, the AI model which is the basis for ChatGPT—experienced a breach of user data. The Garante swiftly launched an investigation that found OpenAI was collecting user-generated data to train its AI model, including “users’ conversations and information on payments by subscribers to the service.”⁵⁶ It deemed the collection of this data to train ChatGPT’s language model unlawful under the GDPR.

On March 31, 2023, the Garante demanded that OpenAI block Italian users from having access to ChatGPT. It further required OpenAI to disclose how it utilizes user data to train its AI model, to address concerns that ChatGPT produced inaccurate information about individuals, and to create an age verification mechanism within a month—or risk being fined 20 million euros or 4% of the company’s annual turnover.⁵⁷

Italy Scrutinizes ChatGPT's Data Practices (cont'd)

Throughout April, OpenAI implemented changes to meet the Garante's demands, including a new information notice describing how personal data is used to train its AI model, as well as a new, ad-hoc form that allows users to opt out from having their data processed to train the algorithms. They also added an age verification system and gave users the ability to erase personal information they deem inaccurate. However, OpenAI stated that "it is technically impossible, as of now, to rectify inaccuracies."⁵⁸

The Garante accepted OpenAI's changes and allowed Italians to access the chatbot again. Yet the regulator continued its investigations into the developer's data practices, concluding on January 29, 2024, that ChatGPT is in breach of the GDPR and giving OpenAI 30 days to respond with a defense against the alleged breaches.⁵⁹

In the United States, discussions about the permission needed for data used to build generative AI have tended to shift toward copyright given that, in the absence of a federal consumer privacy law, copyright has offered the clearest path for content creators to demand that companies remove their data from training datasets.⁵² This approach yields mixed results, given the challenge of reverse engineering the existence of a particular item of content in a system's training data absent any transparency obligations by the companies to share how and with what they trained their models. It is also a poor approach for resolving privacy issues other than those that may implicate copyrightable content.

In July 2023, the Federal Trade Commission (FTC) issued a civil investigative demand to Open AI with detailed requests concerning their training data.⁵³ This highly specific focus on obtaining information about a company's training data is not without precedent; the FTC has settled multiple investigations with companies that used AI in their product offerings, demanding that the companies delete their model and the associated

data because the data used to train it was improperly acquired.⁵⁴ Lina Khan, chair of the FTC, argued in a *New York Times* op-ed that "exploitative collection or use of personal data" falls within the agency's authority to prohibit "unfair or deceptive trade practices."⁵⁵

These events demonstrate that both EU and U.S. regulators have some flexibility and regulatory tools at their disposal to adapt enforcement to changes in technology. Nonetheless, relying only on existing legislation, especially in the United States, is akin to bringing a knife to a gunfight. While the GDPR is settled law, as of early 2024 the CCPA remains a work in progress that is unlikely to be finalized until later in the year. As we discuss in the next chapter, incorporating automated decision-making into these regulations provides the necessary latitude for regulators to include AI in their oversight of algorithmic systems, and to potentially broaden their scope to focus on AI-specific issues, such as training data.

Chapter 3: Provocations and Predictions

In this chapter, we present a set of four provocations and predictions that we believe highlight the key issues that must be confronted as we continue with regulating both privacy and AI.

First, we predict that continued AI development will continue to increase developers’ hunger for data—the foundation of AI systems. **Second**, we stress that the privacy harms caused by largely unrestrained data collection extend beyond the individual level to the group and societal levels and that these harms cannot be addressed through the exercise of individual data rights alone. **Third**, we argue that while existing and proposed privacy legislation based on the FIPs will implicitly regulate AI development, they are not sufficient to address societal level privacy harms. **Fourth**, even legislation that contains explicit provisions on algorithmic decision-making and other forms of AI is limited and does not provide the data governance measures needed to meaningfully regulate the data used in AI systems.

a. Data is the foundation of AI systems, which will demand ever greater amounts of data

The era of “Big Data”—the exponentially increased amount of data collected, created, and stored as the internet expanded and people’s online activities grew to encompass virtually every aspect of their lives—created one of the preconditions for the explosive growth of AI. Companies now know more about our personal lives than we ever thought they would: who we are, what we like, where we go, what we do, whom we do it with, and what we think and even feel.

We predict that the expansion of AI systems across the globe will continue to increase the demand for data among developers.

We predict that the expansion of AI systems across the globe will continue to increase the demand for data among developers. This growing demand will heighten the pressure on the entire existing data ecosystem to increase the amount and types of data collected from consumers, as well as incentivize companies to violate the principles of data minimization and purpose limitation in its pursuit of ever more data. Both the totality of data, and the surface areas by which data is generated and collected, such as embedded sensors in household objects, smart appliances, and biometric cameras in public spaces, will continue to expand. AI’s appetite for data currently knows few bounds. According to the Global Partnership of AI, “[b]uilding an AI system typically involves sourcing large amounts of data and creating datasets for training, testing and evaluation, and then deployment. This process is iterative in the sense that it may require several rounds of training, testing and evaluation until the desired outcome is achieved and data plays an important role at each step.”⁶⁰ None of the AI advances achieved over the past decade would have happened without this broad availability combined with the massively more powerful computers, processing capacity, and cloud storage that developed at the same time. As Mark

MacCarthy describes, “artificial intelligence, machine learning, cloud computing, big data analytics and the Internet of Things rest firmly on the ubiquity of data collection, the collapse of data storage costs, and the astonishing power of new analytic techniques to derive novel insights that can improve decision-making in all areas of economic, social and political life.”⁶¹

Companies have not been incentivized to curb their collection of consumer data, in part due to competitive pressures to maximize targeted, highly personalized services, a task that requires data collection for analytical purposes even if the initial purpose and value of collecting the data is speculative. As commercial sector AI development increased, so did companies’ demands for data—the result, in part, of testing AI systems that generally show improvements in the accuracy and validity of outputs when exposed to greater amounts of sufficiently representative training data.⁶² Predictive AI in particular demands large datasets in order to complete advanced pattern analysis, where almost any variable could potentially hold the key to reliable correlations or associations between inputs and outputs. However, a growing body of research is increasingly challenging the assumption that more data is better by showing that similar performance levels can be achieved using comparatively less data overall when it is selected with more intentionality and specificity.⁶³

While not all applications of AI require consumer data, the largest technology companies which have been building massive stores of consumer data for at least fifteen years and in some cases longer, have emerged with a marketplace advantage in the development of AI in part because of their ready access to these immense datasets. Newer AI developers like Anthropic or OpenAI have had to turn to other data sources to acquire the data to build and train their systems.⁶⁴

Companies have not been incentivized to curb their collection of consumer data, in part due to competitive pressures to maximize targeted, highly personalized services.

While most forms of predictive (machine learning-based) AI are data-dependent for their development, it is the recent emergence of powerful generative AI systems that best illustrates the magnitude of data required for model training. Generative systems such as LLMs (like GPT-4) and user-facing tools built on top of them (like ChatGPT), as well as image generation systems like Stable Diffusion or Midjourney have dazzled the public with their practical as well as entertaining applications. At the same time, as we discussed in Chapter 2, their high visibility has raised questions about how such systems operate, including what data they are trained on, and the potential privacy and other risks of interacting with these systems.⁶⁵

There are presently no transparency mandates requiring companies to detail where and how they acquire their training data outside of the EU AI Act, and those requirements only apply to systems designated as high-risk.⁶⁶ Many of the largest companies building generative AI systems have not been responsive to public inquiries into where they source their data and what procedures they use to strip their training data of personally identifiable information and other sensitive

aspects.⁶⁷ Of course, legal jurisdictions also matter; web scraping that captures personal information that is legal in the United States may not be permissible under the GDPR, and companies are increasingly forced to navigate territorial issues, both between the United States and the EU and others following the GDPR model.

b. AI systems pose unique risks to both individual and societal privacy that require new approaches to regulation

Existing and proposed privacy regulations are largely a retrospective answer to the past twenty years of technological change and increasing threats to our individual data privacy. However, the rise in the breadth and amount of data collected from individuals across all aspects of their online interactions, as well as new threats posed by AI systems, require that we think prospectively and ensure that we have the tools in place to grapple with the changes ahead. Further, beyond the documented harms to individuals, AI systems also pose considerable societal privacy risks that existing regulations are ill equipped to address.

Risks and Harms to Individuals from AI Systems

Information privacy can be a difficult concept to specify as it is both multidimensional and highly contextual. Law professors Danielle Citron and Daniel Solove created a taxonomy of information privacy harms, which include physical, economic, reputational, emotional, and relational harms to individuals.⁶⁸ Citron and Solove also call out discrimination and vulnerability-based harms—those that can occur due to information asymmetries between individuals and data collectors. There are also the harms that the FIPs were intended to address:

harms to one's autonomy, the inability to make informed choices, the inability to correct data, and a general lack of control over how one's information is gathered and used. All of these are relevant to AI based systems as they were to the technological developments of the past three decades of internet expansion.

While these harms predate the application of AI to the consumer sector, commercial AI systems will cause them, exacerbate them, and even pose new ones. For example, recent research based on Solove's own privacy taxonomy⁶⁹ identified not only existing privacy risks that AI exacerbates but also those the authors argue AI creates, such as new forms of identity-based risks, data aggregation and inference risks, personality and emotional state inferences via applications of phrenology and physiognomy, exposure of previously unavailable or redacted sensitive personal information, misidentification and defamation.⁷⁰ Technical advances in AI are also creating new avenues for privacy harm, such as the harms caused by generative AI systems inferring personal information about individuals or providing users with the ability to target individuals by creating content about them that is defamatory or impersonates them.

In addition to traditional concerns about individual privacy and personal data, these systems generate predictive or creative output that, through relational inferences, can even impact people whose data was not included in the training datasets or who may never have been users of the systems themselves. When personal data is included in the training dataset, research has demonstrated that these systems can memorize the data and then expose it to other users as part of the outputs.⁷¹ While most generative AI systems advise that individuals not include personal data in prompts or other inputs, many people still do, and when users of these systems input personal

information, including confidential or legally protected data, these systems may store this data for future uses, including model retraining, or share it with other users as part of the system outputs.

The evolution of risks to online information privacy over the past two decades is a history of ever-increasing consumer surveillance and individual profiling, primarily driven by the goal of targeting consumers with advertisements and offers based on their behavior both on- and offline. As social media platforms proliferated and grew, they too became an avenue for consumer surveillance, expanding the realm of information that could be collected about consumers across a growing set of contexts. Mobile devices and apps, smart speakers, smart home devices—each new technological development added another layer of information that could be collected beyond the initial ambit of online shopping.

Shoshana Zuboff terms the practice of extracting value from and about individuals *surveillance capitalism*, which “unilaterally claims human experience as free raw material for translation into behavioral data.”⁷² Today it is exceedingly difficult, if not impossible, for an individual using online or connected products or services to escape systematic digital surveillance across most facets of their life. The collection of personal data occurs not only in instances where individuals make the choice to engage directly with an app or a service; in many cases, it also occurs silently by invisible third parties tracking individuals’ actions in browsers and mobile apps without giving affirmative notification or securing their consent.

The focus on capturing consumer behavior and using it for predictive purposes expanded with the proliferation of sources for data collection. Individual profiling and inference-making became indispensable

The evolution of risks to online information privacy over the past two decades is a history of ever-increasing consumer surveillance and individual profiling.

for a broadening range of contexts beyond merely serving ads. Profiling for determining credit, insurance, employment, housing, and medicine are but a few examples. Over the past five years, emergent AI systems have increasingly been deployed in these contexts as well, as their predictive capabilities are even greater than previous big data applications due to the computational capacities of AI.

The Future of Privacy Forum categorizes the harms to individuals from automated systems into four areas: losses of opportunity, liberty, economic losses, and social detriments.⁷³ These can result in harms such as discrimination in housing, employment, education, and other areas; surveillance and incarceration; denials of credit, differential pricing, and an overall narrowing of available choices; and harms to dignity due to bias or opportunity losses, as well as algorithm-based social sorting and filtering that can influence what or whom you connect with in digitally mediated social environments. Profiling in particular increases the scope and scale of data collected about individuals and the related inference-building across a variety of contexts.

There is also a lack of transparency about how automated systems function, making it challenging for individuals to alter or limit their impact. AI

systems can automate many forms of decision-making and classification, exacerbating the privacy risks and harms already present in our “pre-AI” data ecosystem.⁷⁴ The potential harms resulting from such privacy infringements aren’t limited to the consumer marketplace, where today companies can not only tailor advertisements to you with fine-grained precision, but in some cases also use the data they have collected and inferred about you for manipulative or discriminatory commercial uses.⁷⁵ The result is an exploitation of data that undermines societal norms and values by removing the structural and contextual barriers that previously acted as safeguards against its widespread access.⁷⁶ This is a means of collecting data against which FIPs-based, individual due process rights offer little protection or recourse. Individuals cannot use the FIPs effectively to protect themselves from this form of data collection, especially when it happens without notice, through inferences, and even from sources we may be unaware of (including when data scrapers obtain data from services without permission, such as from social media or photo-sharing sites).⁷⁷

As nearly all facets of our lives are increasingly mediated through technology, the risks increase for AI systems to perpetuate biases, stereotypes, and errors, manipulate consumers, and enable discrimination, particularly in the absence of regulation or transparency measures designed to keep these harms in check. Already, the scope and scale of our “data relationships”—with the companies that collect our data directly (first party) and those that do so indirectly (third party)—are too numerous for individuals to manage in any reasonable way, assuming we even know who is collecting our data. Under existing or proposed privacy laws, the incentives for companies to collect as much data as they possibly can is unlikely to diminish. As generative AI systems continue to

proliferate, many built with online data scraped from the internet without consent, individuals stand little chance of addressing these privacy risks themselves through opt-out, correction, or deletion rights.

Societal Risks and Harms to Privacy from AI

The privacy risks and harms posed by AI systems are not limited to individuals; they also threaten groups and society at large in ways that cannot be mitigated through the exercise of individual data rights. Returning to the Future of Privacy Forum’s taxonomy, the societal-level harms from automated systems based on group membership include differential access to opportunities such as jobs, housing, education, credit, goods and services; increased surveillance and disproportionate incarceration of specific groups; and reinforcement of negative stereotypes and biases.⁷⁸ AI systems create the capacity for large-scale societal risks precisely because they operate at scale, analyzing tremendous amounts of data and in turn making connections and predictions previously not possible through other means. This capacity can result in classifying and applying decisional outcomes to large swaths of the population based on group affiliation—thereby amplifying social biases for particular groups. Harms at the societal level can also pose threats to democracy, as well as impact the benefits that privacy affords individuals, which in turn impact the development of autonomy necessary for cultural and societal flourishing.⁷⁹

From a privacy perspective, a specific concern is that profiling at a societal level contributes to a widespread erosion of privacy norms and expectations. The expectation that your data will be gathered at every turn, the powerlessness of being unable to do anything about it, and the lack of transparency about how one’s data is used or decisions are made about you all feed a growing sense of inevitability that data

privacy has already been lost.⁸⁰ This is not simply a reflection of changing norms about online sharing and publicness, as Mark Zuckerberg disingenuously argued in 2009 when forcing Facebook users' data to be public by default—and setting the stage for the Cambridge Analytica scandal.⁸¹ The growth of generative AI has drawn attention to the pervasiveness of data collection, and the sources of that data, as the connection between scraped data and the ability of generative systems to create their wondrous outputs has raised questions about exactly where the data is coming from. The more we mine the public sphere for data, the more we erode the sense that we should have a right to exist in public, whether a digitally mediated space or a physical one, with any degree of privacy or anonymity.

This shift toward using AI in contexts with civil rights implications, such as hiring,⁸² criminal justice,⁸³ and policing⁸⁴ have profound implications for both individuals and society at large. Individuals interact with systems they may not think of as highly technical (such as applying for a job), or to which they haven't signed up as users, but within which AI calculations are applied to them through inferences—to, say, predict their health outcomes, calculate their insurance rates, or determine whether their employment application gets reviewed. As these systems proliferate, they can amplify existing biases and inequities. At their most extreme, they can be used by governments as tools of social control.

c. Data protection principles in existing privacy laws will have an implicit, but limited, impact on AI development

The application of specific fair information practices in existing regulations, such as requirements for data

The question we raise is whether [existing data protection] principles are sufficient for tackling the privacy risks and harms posed by AI.

minimization and purpose limitation, will impact AI development. The question we raise is whether these principles are sufficient for tackling the privacy risks and harms posed by AI. In the United States, lawmakers are increasingly arguing that passing federal privacy legislation, similar to the GDPR, is a necessary precondition to any regulation that explicitly targets AI systems.⁸⁵ Existing privacy and data protection laws in both the EU and the United States (at the state level) will regulate AI systems that rely on personal data for training purposes or that ingest it as part of the service they offer—but only up to a point. Even if the United States adopts a GDPR-esque law that provides FIPs-based rights, this approach will not be sufficient on its own to address the risks and harms we discussed above.

Indirect Regulation Through Data Minimization and Purpose Limitation

Both the CCPA and the GDPR, as well as other similar federal agency and state-level regulations in the United States and EU member state regulations, impact the development and deployment of AI-based systems by limiting the personal data that companies can collect and use to train and retrain AI models ad infinitum.⁸⁶ Specifically, the principles of data minimization and purpose limitation, if clearly delineated and enforced, should limit how much personal information is collected and how it can be used and reused for AI

systems. Companies need to justify how data collected from consumers in one context for a particular use could be reused in an entirely different context or for a new purpose. However, the degree of protection varies considerably between jurisdictions. With the GDPR as their foundation since 2018, EU member states have a stronger and broader set of enforcement powers than do the minority of U.S. states that have passed data privacy laws.

In response, researchers and industry practitioners have already developed, tested, and deployed a wide array of techniques to meet data minimization and purpose limitation requirements—without compromising performance.⁸⁷ During the training phase of AI models, privacy-preserving methods (including federated learning) have been employed to minimize data.⁸⁸ During the model inference phase, experts point to the conversion of personal data into less “human readable” formats, the anonymization of queries, and data shuffling among other privacy-preserving techniques.⁸⁹ Still, more research into data minimization and purpose limitation compliance in AI systems is greatly needed.⁹⁰

Existing privacy laws do address the use of data collected or generated directly by an AI system (e.g., a user’s prompts to a chatbot or other generative AI system, or data processed by a predictive AI system, such as a recruiting and hiring software). To the extent that these systems directly ingest or process personal data, or make predictions or inferences about individuals based on this collected data, privacy regulations implicitly regulate their operation by requiring compliance with individuals’ rights to access, correct, and delete personal data, to request a copy of their data, or to opt out of future sales or sharing of their data. In many AI use cases, companies must conduct the same privacy or data protection

impact and/or risk assessments that they would even if not utilizing AI in order to demonstrate they have adequately considered the risks to individuals by collecting and using personal data as part of deploying their systems.

Limitations of the FIPs-based Framework

The FIPs provide the substantive framework for existing privacy and data protection laws around the globe, based on principles that were developed over 50 years ago. Many policymakers view them as a model for future privacy legislation; even China has adopted a version of the FIPs in its own privacy legislation, largely viewed as modeled after the GDPR (see Chapter 2). However, both the FIPs and the laws based upon them have their critics. Law professor Woodrow Hartzog in particular has criticized the FIPs as inadequate but invaluable, noting that in a modern society awash in data and data collection, “control does not scale.”⁹¹

Data Minimization and Purpose Limitation

Enforcement of the data minimization and purpose limitation principles should, in theory, translate to more conservative and thoughtful personal data collection. However, these approaches as practiced today fail to address many of the fundamental weaknesses of our current data ecosystem. For example, they do not address the inequitable power dynamics of a data ecosystem in which the data collectors and processors, most of which are powerful private tech companies, hold far more market power over personal data collection than do individuals. Further, it may be reasonably straightforward to hold a company to account if its use of data doesn’t match the purpose it gave at collection. However, in the absence of an agreement as to what constitutes too much data, it will be a challenge for regulators to operationalize whether a company is sufficiently practicing data minimization outside of egregious violations. Today, the pursuit of

quality (i.e., data that is reliable, relevant, and collected ethically) is still mostly overridden by a pursuit of *quantity* (i.e., collecting vast amounts of data cheaply and at scale, by any means necessary)—especially in markets that lack robust privacy legislation like the United States.⁹²

The Limits of Privacy Self-Management

A core weakness with the FIPs framework is that individuals are assumed to have a level of control and power equal to that held by companies and institutions collecting and processing their data.⁹³ However, this is not the case; often individuals cannot simply choose an alternate product or service with more privacy protective data collection practices. Monopolistic practices in the tech sector, consumer lock-in, and a general incentive for businesses to collect as much data as possible undermine privacy as a competitive factor except in a few cases. Privacy law expert Daniel Solove named the burden on individuals to manage and exercise their rights to curb data collection “privacy self-management.” As our use of digital products and services has increased, privacy self-management has failed to give individuals the tools they need if they want to prevent, or at least reduce, the amount of data collected about them.⁹⁴

Thus, while the FIPs are a necessary baseline to ensure that individuals have due process rights with respect to their personal information, they fail to empower individuals to have a meaningful impact on their privacy in the age of AI. FIPs-based regulations may be designed to constrain companies from collecting and processing data for AI systems, but they ultimately don’t solve the core problem of how to prevent data collection in the first place in a society where it is difficult, if not impossible, for the majority of people to avoid interacting with technology.

The FIPs (...) fail to empower individuals to have a meaningful impact on their privacy in the age of AI.

Data Collection by Default: Opt In or Opt Out?

The expansion of the FIPs from their original application to governmental data collection in the early 1970s to the private sector reinforced the approach of allowing data collection by default. There are legitimate reasons to allow governments to collect data in many circumstances without requiring individuals to give their explicit consent: tax collection, census taking, and provisioning public benefits are but a few examples. But applying this rationale to the private sector normalized the idea that individuals should have to opt *out*, rather than choose to opt *in*.

The GDPR tries but does not fully resolve this dilemma. As Mark MacCarthy notes, the “GDPR provides procedural, not substantive protections. Its goal is not to limit any specific use of information but to ensure that all uses are subject to certain fair procedures to ensure the protection of the rights of data subjects.”⁹⁵ The GDPR threads the needle between always requiring consent and allowing collection without it by providing six bases for processing data; the two most salient for this discussion are *consent* and *legitimate interest*.⁹⁶ No matter which basis is used, data processors must inform individuals about the processing when collecting their data, and the processing must not “seriously impact” individuals rights and freedoms.⁹⁷

Arguably, asking for *consent* (opt in) is the most straightforward basis for data processing, as individuals maintain the right to withdraw it. When it comes to *legitimate interest*, the U.K. Information Commissioner’s Office describes it as the most “flexible” of the bases.⁹⁸ The office suggests that legitimate interest may be appropriate when processing offers a clear benefit “to you (the processor) or others, if there is limited privacy impact on the individual, the use matches an individual’s reasonable expectations, or the controller cannot or does not want to give the individual “full upfront control (i.e., consent) or bother them with disruptive consent requests when they are unlikely to object to the processing.”⁹⁹ Legitimate interest can allow for opt-out-based data collection, though controllers are cautioned that it cannot be used as a basis for all data processing, and controllers must have a clear justification for using it for their particular context.

Legitimate interest has been criticized for allowing data collection practices that some argue violate the basis and act as an opt-out rather than an opt-in basis. For example, in March 2023, in response to a ruling by the European Data Protection Board that denied Meta’s use of its sites’ terms of service as a basis for using behavioral targeting in advertising, Meta then switched to the legitimate interest basis,¹⁰⁰ which activist group Noyb argued was a violation of users’ fundamental rights.¹⁰¹ The switch requires Facebook and Instagram users to submit an online form to register their objection to the use of their behavioral product usage for targeting; unless they object, however, Meta will proceed with the targeting, placing the burden on individual users to sort out the details.¹⁰²

While the FIPs provide crucial procedural data protection rights, they fundamentally do not curb data collection by default. Instead, the focus is on

rights one can exercise after data has already been collected, leaving the burden of managing one’s privacy on individuals who may have little time or inclination to actively participate in this work. They do not provide a right of refusal, or a clear, convenient, non-fatiguing means to interact with digital products or services without having to give up some personal information.¹⁰³ While the EU’s 2002 ePrivacy Directive (discussed in Chapter 4) attempted to curb cookie setting by default, and GDPR has positively impacted the design and simplicity of cookie consents, they remain the hallmark of how *not* to implement opt in. As we will argue later, there are better ways to implement an opt in approach.

d. The explicit algorithmic and AI-based provisions in existing laws do not sufficiently address privacy risks

In addition to the implicit impacts of FIPs-based laws discussed above, both existing and proposed privacy regulations include specific provisions targeted at algorithmic systems in such a way that will include AI. These include provisions in the GDPR and U.S. state laws, such as California’s, that address automated decision-making and profiling and that require data protection impact assessments to obligate companies to identify uses of data that pose risks to their customers.¹⁰⁴

These provisions are intended to ensure that privacy and data protection regulations cover specific data-intensive practices that implicate individual data privacy. They use a risk-based framework to place obligations on data processors to incorporate risk mitigation into their data governance practices that includes risks to their customers, not just to the

business. These measures will have some impact on AI systems as we discuss below. However, these explicit regulations do not address the limitations of the FIPs framework, nor do they sufficiently focus on broader data governance measures needed to regulate the data used for AI development. Addressing these challenges will require additional policy measures, which we discuss in Chapter 4.

Automated Decision-Making and AI

The term automated decision-making (ADM) is not a recent invention. Concerns with delegating decision-making about individuals using personal information to automated systems dates back at least to the origination of the FIPs in the early 1970s,¹⁰⁵ though the FIPs themselves do not address the issue of automation. The U.K. Information Commissioner’s Office defines ADM as “the process of making a decision by automated means without any human involvement. These decisions can be based on factual data, as well as on digitally created profiles or inferred data.”¹⁰⁶ The GDPR incorporates the concept in Article 22, noting that “[t]he data subject shall have the right not to be subject to a decision based solely on automated processing, including profiling, which produces legal effects concerning him or her or similarly significantly affects him or her.”¹⁰⁷ The inclusion of the term “profiling”—defined as “any form of automated processing of personal data consisting of the use of personal data to evaluate certain personal aspects relating to a natural person, in particular to analyse or predict aspects concerning that natural person’s performance at work, economic situation, health, personal preferences, interests, reliability, behaviour, location or movements”—is a specific call-out of data mining and prediction practices that implicate privacy through their use of personal data and their focus on individuals.¹⁰⁸

These explicit regulations do not address the limitations of the FIPs framework, nor do they sufficiently focus on broader data governance measures needed to regulate the data used for AI development.

ADM can arguably be construed as including any form of AI that trains on or ingests personal data, or that makes predictions or decisions about individuals, though existing laws narrow the applicability to ADM with significant or legal impacts to avoid the overinclusion of low to no privacy risk ADM (e.g., an algorithm that takes in an address to find a nearby store but does not store that location data, or a clothing sizing algorithm that asks the customer for their size in specific brands of clothing to calculate a more accurate sizing estimate). The term itself is not specific to any particular technology, describing a process that can be accomplished using rule-based algorithms as well as forms of AI, such as predictive AI.¹⁰⁹

The GDPR provides data subjects the rights to contest and withdraw from automated processing, creating guardrails to prevent a Kafkaesque landscape of ADM that cannot be contested.¹¹⁰ In this context, the GDPR expressly connects automated processing to the practice of profiling as a threat to privacy. The scope of the GDPR’s automated decision-making provision arguably impacts AI to the extent that a system renders a decision, or makes a prediction, that results in significant impact on individuals’ lives and relies on processing personal data to do so. Regulations that

specifically call out ADM follow the GDPR’s lead with focusing on systems with “legal effects” or similar impact, such as extending or denying credit, hiring, housing eligibility, and so on.

The present scope of ADM regulations in the EU and the United States focuses on providing notice to consumers that automated processing is occurring, giving them opt-out rights in qualifying contexts (e.g., with significant impacts or legal effects) and requiring ADM systems to provide information about the “logic” of the system design: its purpose, how it renders decisions, the potential safeguards in place, and clarifying the extent of human oversight over the system. For example, the crafters of the CCPA¹¹¹ and the subsequent update to the law (the California Privacy Rights Act of 2020¹¹²) tasked the new California Privacy Protection Agency with regulating access and opt-out rights for businesses’ use of ADM that processes personal information (including training data¹¹³) or otherwise poses a risk to privacy.¹¹⁴ Colorado’s 2023 privacy law also includes ADM regulations, distinguishing between solely automated, human-reviewed, or human-involved automated processing, and sets obligations in accordance with the level of human involvement.¹¹⁵ These types of measures closely follow the FIPs (i.e., notice, data access, data correction) in providing procedural rights and protections.

Privacy and Data Protection Impact Assessments

Impact assessments for privacy and data protection have their roots in the growth of environmental protection regulation that emerged in the 1960s.¹¹⁶ In the privacy and data protection sectors, they are used to guide both public and private sector organizations toward proactive risk assessment when planning a new product or service that utilizes personal data.

In the United States, Section 208 of the e-Government Act of 2002 obligates federal agencies to conduct privacy impact assessments (PIAs) when “developing or procuring information technology that collects, maintains, or disseminates information that is in an identifiable form.”¹¹⁷ The GDPR requires data protection impact assessments (DPIAs) that are triggered “whenever processing is likely to result in a high risk to the rights and freedoms of individuals,”¹¹⁸ such as large-scale uses of sensitive data or public surveillance, systematic individual profiling, and automated decision-making without human involvement. Once the regulatory process is completed in 2024,¹¹⁹ the CCPA may require DPIAs from companies whose processing of data poses a substantial risk to privacy, including selling or sharing personal information; processing sensitive personal information; using ADM in specific ways (including decisions with “legal or similarly significant effects”); profiling employees, job applicants, students, and consumers in publicly accessible places; behavioral advertising profiling; and processing the personal information of children under 16.¹²⁰

PIAs and DPIAs are tools to prompt organizations to engage in sufficient planning and self-reflection to foresee potential risks and to integrate mitigations into their design and planning processes (or in the case of startups, to compel them to adopt such processes) by considering both the data types and the processing activities that pose high risks to individuals. Algorithmic impact assessments have been proposed as a tool for the general oversight of AI systems, and though they may make mention of privacy and data, they are not focused exclusively on these topics.¹²¹

How Do These Explicit Provisions Fall Short?

The scope of ADM provisions in both the GDPR and the CCPA seek to strike a balance between

preventing runaway ADM scenarios versus casting such a broad net that every form of ADM requires the application of the full set of notice and opt-out rights. To be sure, an opt-out right is a potentially powerful deterrent against the over-application of ADM. The prospect of creating an alternate non-ADM process for businesses to comply with opt-out requests may keep privacy and data protection lawyers up at night. Consider, for example, a large company that receives thousands of job applications per month—requiring a non-automated process for rendering judgments for applicants could be daunting.

However, the underlying logic of both opt-outs and notice requirements doubles down on the privacy self-management approach, placing the burden on individuals to understand what automated decision-making is and why they may wish to opt out of it. The notice and consent approach for privacy already places a significant burden on individuals not only to exercise these rights but also to comprehend why one might want to do so.¹²² Given the complexity of understanding AI systems and how one's data may interact with them, this presents an even heavier lift. And, it's unclear what it may accomplish for consumers.

The requirement to label ADM systems for the public may raise more questions for consumers than answers. Paradoxically, it's conceivable that people might elect to opt out of ADM systems in favor of a non-automated process that could turn out to be even more arbitrary or biased than an ADM-based one. In essence, this becomes a form of labeling not unlike the practice of genetically modified ingredient (GMO) labeling on foods; in the absence of clear scientific evidence determining whether the consumption of GMO-based foods is harmful, the decision to consume them or not is punted to the consumer who may have little to no understanding of the issue, leading them

The underlying logic of both opt-outs and notice requirements doubles down on the privacy self-management approach, placing the burden on individuals to understand what automated decision-making is and why they may wish to opt out of it.

to make uninformed choices that could help or harm them. This is not to say that labeling or providing notice of the use of ADM has no benefit; certainly, to the extent that there are individuals who want to exercise the right not to be subject to ADM, this opt-out right is crucial. But it becomes an exercise of personal preference that has the potential to harm an individual given their unique circumstances. And while existing federal and state laws may prohibit AI systems that produce discriminatory or biased outputs, this approach leaves a loophole in regard to systems that may have negative implications for one's data privacy. A framework that would place strict limits on the use of AI systems that had negative privacy impacts both for individuals and at a societal level would provide a more consistent approach. Finally, the ADM track may miss most uses of generative AI systems to the extent that they are not used for decision-making purposes, leaving open the question of whether they could be implemented in a way that consumers may not know that AI is being utilized but not subject to notice requirements.

While data protection impact assessments are a necessary and useful regulatory tool for protecting data privacy, they are not a steadfast guarantee against either government or private companies implementing harmful technologies. For one, they depend on a regulatory or institutional structure that has sufficient authority to act when DPIAs or PIAs are done poorly or fail to anticipate risks.¹²³ Without such structural support, they are little more than a bureaucratic hurdle with no teeth. In the United States, for example, several of the large technology companies developing AI systems have elected to blow past the advice of their own risk-adverse legal staff and responsible innovation teams and market AI tools without fully understanding the risks to their users, let alone the larger public.¹²⁴ Another issue is that if there isn't a standard by which impact assessments can be assessed, businesses can turn the process into an exercise of grading their own homework by setting their own internal standards.

While the GDPR's DPIA requirements recognize the heightened risks from some forms of data processing, including from fully automated decision-making technologies, this approach assumes one is acting on data that has already been collected. It is possible that the process of anticipating as well as completing a DPIA could dissuade or prevent an organization from electing to launch a product or service due to the risks it identifies. But the fact that these tools presently do not direct organizations to engage in these processes before product creation—perhaps even before collecting training data—means that their ability to surface risky data collection or data management decisions, or to prevent such actions, is lessened. It is possible in light of the importance of training data on the outputs of AI systems that the entire data development pipeline should be subject to data or privacy impact assessments. Mehtab Khan

and Alex Hanna review the stages of dataset creation and discuss some of the documentation interventions that are increasingly being suggested to add greater accountability to the dataset development process (we will discuss these in more depth in Chapter 4).¹²⁵

The proposed California regulations do attempt to tackle some of these issues. For example, §7154 places disclosure obligations on businesses that process personal information to train ADM systems or AI, requiring that they disclose to downstream users the appropriate uses of the technology, as well as conduct their own risk assessment that addresses “any safeguards the business has implemented or will implement to ensure that the automated decision-making technology or artificial intelligence is used for appropriate purposes by other persons.”¹²⁶ Further, for businesses required to conduct a risk assessment as described above, §7155 prohibits businesses from processing personal information if the “risks to consumers’ privacy outweigh the benefits resulting from processing to the consumer, the business, other stakeholders, and the public.”¹²⁷ In order for our existing frameworks to fully grapple with AI-based privacy threats regulators will need to keep refining and expanding provisions like these, and more.

e. Closing thoughts

Overall, FIPs-based privacy and data protection laws have not anticipated the growth of AI systems or variants such as generative AI. Despite a decade-plus exposure to the emergence and growth of big data, these regulatory frameworks are not prepared to respond to and oversee the data-intensive aspects of AI systems. Ultimately, existing FIPs-based privacy regulations cannot sufficiently regulate the data that feeds AI development in a way that sustains our

existing state of privacy today or, better yet, improves it. A privacy and data protection framework that places the primary responsibility on individuals to manage their data across hundreds, even thousands, of digital relationships and channels fundamentally does not scale, and thus will not succeed in protecting individual privacy. Nor will it solve population or societal level risks and harms to privacy. As legal theorist Salomé Viljoen notes, “responding adequately to the economic imperatives and social effects of data production will require moving past proposals for individualist data-subject *rights* and toward theorizing the collective *institutional forms* required for responsible data governance.”¹²⁸ These challenges have become more visible following the explosion of generative AI systems, built primarily from data scraped online. It will be very difficult, if not impossible, for individuals to shoulder the burden of exercising their deletion and correction rights with these massive and nontransparent systems, let alone to proactively prevent their data from inclusion. Privacy self-management approaches that force individuals to bear the burden of systemic privacy challenges will not substantively improve individual privacy.

Unfortunately, passing more FIPs-based regulations will not resolve individual privacy challenges or systemic risks posed by AI systems. Even if the United States were to pass the 2022 version of ADPPA, neither it nor the GDPR provides sufficient oversight of the data used to develop and train AI. While these laws nibble at the edges, they do not confront the bias toward collecting data first and asking questions later. They do not adequately address consent. They do not provide sufficient methods for people to engage with technological systems without ubiquitous data collection. They do not address societal-level privacy harms. And they do not provide a framework for addressing the privacy issues raised by AI training

data, whether they be from proprietary datasets, open source or public datasets, or data scraped from the internet. The computer science maxim of “garbage in, garbage out” is as relevant as ever when it comes to building AI systems. Whether they are trained on curated but biased datasets, or on data scraped from questionable parts of the internet, the impact assessment process must direct attention to the antecedents of AI products and not simply their outputs.

In the next chapter, we make three suggestions we believe should be adopted in order to address these issues. They alone may not be sufficient to address the issues we have raised above, but we believe they are an important starting point for moving forward in the right direction.

Chapter 4: Suggestions for Mitigating the Privacy Harms of AI

The adoption of AI can bring benefits across many different societal contexts if—and only if—AI systems are designed to center human needs and values. But AI systems’ requirements for data, combined with applications that will generate or consume an extraordinary amount of personal information, raise several crucial questions: Is data privacy compatible with the growth of AI? Can we have a widespread adoption of AI and still preserve our information privacy, even at the minimum state it exists today? Can we do better?

The rapid growth and adoption of AI raises legitimate concerns about its possible risks to humanity. At the same time that we are debating questions of whether and how we want to live in a world that utilizes AI, some are questioning whether governments should adopt bright line rules that forbid particular applications of AI completely.¹²⁹ We suggest that when evaluating these issues, policymakers must also consider that a side effect of AI’s adoption could be a world with substantially diminished data privacy for all of us unless we specifically take measures to protect it. The suggestions we make below are motivated by this question: What will it take for both data privacy and AI to coexist?

What will it take for both data privacy and AI to coexist?

As we note in the introduction, a significant assumption in the framing of our questions is that, especially in the United States but also in the EU (and many countries around the world), the present state of data privacy is suboptimal. Individual data rights are both necessary and insufficient for protecting data privacy in a world with AI. Even in countries and states with data rights in place, the burden continues to be on individuals to exercise their rights after data collection rather than for their preferences to be respected at or before the initiation of any collection. This approach, as we have argued, also neglects societal risks and threatens our collective data privacy.

While data privacy is the focus of this paper, it isn’t the only lens or priority when considering how to regulate AI. For example, the extent to which others are able to gather data about you and potentially make inferences about you has direct implications for issues of bias and discrimination, whether those others are private companies or the government. The reality that non-personal information as well as others’ personal information can also be used to make inferences about both individuals and groups is yet another reason policymaking on data privacy must move beyond individual control to set clear rules on data collection and use more broadly.

Another key aspect to this debate is that possessing data gives rise to considerable market power. In the AI land grab presently underway, the actors who already possess large datasets have a significant advantage over developers that do not have stores of data and must gather, purchase, or license it.¹³⁰ Additionally,

one must also have the resources to pay for curating and labeling data, to transform it into a quality resource for training AI systems. While advocating for data quality is an important issue in this debate, as higher quality data can help address issues of bias and discrimination, we must acknowledge that it can also be a source of power and advantage as large enterprises can expend the resources to improve their data or license quality data from others.

We are also pushing back against a strain of technological determinism in these debates. Much like the arguments that privacy is dead and we should all acquiesce to a total loss of control over our data in exchange for the bounty of free online services, in many of the discussions around AI today, and generative AI specifically, there are assumptions that there are no limits on what data should be included in AI models, particularly foundation models. When the scrapable data on the entire public and publicly accessible internet appears up for grabs, we can be forgiven for assuming that this path is inevitable. This line of thinking in particular appears driven by computer scientists and others in AI development who are focused on the lure of quantity over quality and do not consider the sociotechnical context in which data resides. As we point out in Chapter 3, some researchers are already questioning whether bigger will always be better, even with regards to foundation models, given the trade-offs between capabilities and output quality. A LLM that can make inferences and reason in a human-like way is only useful if the model produces accurate and reliable outputs. Otherwise, the technology may be nothing more than a “stochastic parrot,” mimicking human language without connection to meaning.¹³¹

Finally, it’s important to note that these are concerns that span both commercial and governmental

contexts. The primary focus of this paper has been commercial data collection. But governments can (and do) purchase data from the private sector and direct governmental deployment of AI systems that are trained on or that process personal information, which raises concerning questions about the potential for surveillance and the impact on civil liberties.¹³² To date, several of the public sector uses of AI in the United States that have garnered concern have focused on predictive tools for criminal sentencing,¹³³ or assignment of public benefits¹³⁴ that have perpetuated biases or raised questions about fairness outcomes. Governments building AI systems using administrative data, for example, pose risks that are out of the scope of this paper to explore in depth. But one cannot regulate the commercial sector’s data practices and turn a blind eye to how governments may adopt and use this technology, including when it is procured from the private sector—which, in turn, implicates the sources of the data used to train such systems. In other words, the line between the training data used for private and public uses of AI can easily become blurred. Neither usage exists in a vacuum, which points to the need for data provenance as well as downstream data privacy impacts to be centered in these debates.¹³⁵

With these concerns in mind, we offer the following three suggestions that we believe will aid in mitigating the risks to data privacy posed by the adoption of AI. To corrupt a famous quote: “It’s the data, stupid.”¹³⁶ Any problem-solving about the impacts of AI on data privacy must look beyond individual data rights to include strategies that include the governance and management of data as a resource in a privacy-respecting and preserving manner, as well as a focus on societal impacts and human rights.

Suggestion 1: Denormalize data collection by default

Shift away from opt-out to opt-in data collection by facilitating true data minimization and adopting technological standards to support it.

As discussed earlier, the FIPs provide a crucial framework of rights for our collected data. But the principles of data minimization, purpose limitation, and even consent have been operationalized in ways that normalize data collection by default in many contexts. This normalization can be traced in part to the FIPs' original focus of providing due process rights for government-based data collection rather than for the commercial sector. The FIPs do not include a right to refuse data collection, for example. There is also an assumption of exclusivity and intentionality: that an individual has a one-to-one, known relationship with a data collector with whom one intends (or is required) to interact. The architects of the FIPs did not anticipate the ubiquitous, always-on digital surveillance and data collection enabled by digital networks and mobile devices that emerged in the 2000s. Nor did they foresee that our data would be collected by third parties that have no direct relationship with us. These assumptions have led to practices that prioritize the frictionless operation of the market over adherence to principles.

The one major example of an experiment with both adding friction and surfacing the principle of consent into data collection has not gone well: browser cookie consent dialogs. Cookie consents are a prime example of consent fatigue and how **not** to denormalize data collection. European regulators put consent for data collection by websites front and center with the adoption of the EU's 2002 ePrivacy Directive, the key regulation governing browser cookie consents.¹³⁷ This approach quickly backfired: Requiring individuals

to accept or reject website cookies with every visit inserted too much friction, causing annoyance and confusion for the public. Browser cookies are not only the mechanism that allows websites to identify their visitors, they also allow data collectors to engage in cross-site tracking and profiling. Even today, many internet users struggle to understand what cookies are and how their collection may undermine their privacy. The implementation of the ePrivacy Directive demonstrated the problem with requiring individuals to manage consent on a continual, site-by-site basis in a manner that treats a wide spectrum of possible risks with the same level of notice and choice: The approach does not scale in a world where consumers have relationships with many different online providers. Recent changes sparked by GDPR consent requirements have improved the format of consent notices (e.g., consumers now have a "reject all" option), but they still remain a source of friction and annoyance.

In contrast, the approach taken by the U.S. Federal Trade Commission, which endorsed a watered-down version of the FIPs in 1998, is particularly weak.¹³⁸ Called "notice and consent," the FTC's implementation allows companies to post a notice of their practices (a privacy policy) and to assert that a consumer's use of the service, or perfunctory acceptance of a service's terms and conditions, constitutes adequate consent for whatever data collection practices the company engages in.¹³⁹ Most famously, privacy policies don't actually have to protect consumers' privacy—they must only state what a company will do with the data it collects, which can include selling it to or sharing it with whomever it wishes. The FTC has recently stepped up enforcement of unfair or excessively abusive privacy practices¹⁴⁰ and is signaling an intent to revisit the notice and consent paradigm.¹⁴¹ But the agency's rulemaking process is lengthy (on the order

of five years or longer), and adopting federal legislation that resolves these issues would provide a quicker solution.

While the U.S. policy approach illustrates the outcomes with industry self-regulation as the guiding standard, the EU's demonstrates what happens when tech policy is made without regard to how people actually use technology. Neither way is a win for data privacy. Unfortunately, we continue to be in a stalemate with regard to resolving data collection by default. The EU has delayed the renegotiation of the ePrivacy Directive,¹⁴² and neither the GDPR nor the ePrivacy Directive has curbed all third party data collection practices.¹⁴³ Similarly, if ADPPA were enacted in the United States, it too would not curb data collection by default; the law would add affirmative consent obligations for collecting sensitive data but continues the practice of placing the burden on consumers to opt out of data collection after the fact. There is some progress on this front regarding children—the GDPR prohibits data collection for children under 16 without parental consent, and the Federal Trade Commission is proposing updates to the Children's Online Privacy Protection Act (COPPA) rules¹⁴⁴—but not for adults.

The challenge moving forward in a world with greater demands for data is how to mitigate excess data collection without adding too much friction with excessive consent requests. Digital services need consumer data to operate, and not all such requests are excessive. Some demographic data will be required to assess whether AI systems are biased or discriminatory, though this can be accomplished within the scope of purpose limitation rules.¹⁴⁵ But in the absence of clear rules, the incentives are such that companies will try to maximize data collection, especially if they are concerned their competitors will do so even if they do not. One emergent example in the United States is the

The challenge moving forward in a world with greater demands for data is how to mitigate excess data collection without adding too much friction with excessive consent requests.

collection of consumer mobile phone numbers, which, due to portability regulations, have become a form of persistent identifier similar to social security numbers. Many online services are now requiring that a customer provide a mobile phone number at sign-up even if phone numbers are not necessary for the provisioning of the service. This is a clear overreach and one that should be addressed with both data minimization and purpose limitation rules. If we don't address how data escapes into the larger data ecosystem at the source, we will neither be able to gain a foothold on improving our data privacy, nor exercise adequate control over how our data feeds AI systems.

Data Minimization by Default, Through Defaults

We need to operationalize the principle of data minimization to prevent the collection of excess data. One approach is to apply the privacy by default strategy in recent children's privacy legislation to adults.¹⁴⁶ Digital services would have to set all user accounts by default to the strongest privacy option; it would be up to users whether to change these defaults. Regulators would need to provide guidance

to determine which data use practices companies would be allowed to ask consumers to opt in, or set clear limits through duties of loyalty. For example, companies could be required to ask whether a customer's data could be used for training purposes, and the specificity by which one could be asked (e.g., for any training purpose, or for a specific product?). The opt-in mechanisms themselves would require strict oversight. For example, the CCPA requires that companies not use dark patterns when consumers exercise their right to opt out of the sale of their personal information, preventing companies from manipulating or coercing consumers when making their choice.¹⁴⁷

One example of this approach is demonstrated by Apple's rollout of app tracking transparency (ATT) in iOS 14.5.¹⁴⁸ Users are asked when they first open an app whether they wish to allow an app to track their activity across other apps and websites. The setting prohibits apps from using third party tracking methods unless the user approves it on a per-app basis. ATT was available in a previous release¹⁴⁹ as a setting that users had to proactively find and switch on, but the iOS 14.5 update put the option directly in front of users. It is reported to have significant uptake, with one source reporting the industry-wide average at 70% of users opting out of tracking, and even higher rates in the United States and in Europe.¹⁵⁰ It remains one of the best examples to date of how consumers will choose pro-privacy defaults when asked simply, clearly, and at a sensible decision point in their task flow without adding excessive friction.

Automate Consumer Preferences

The ATT example above succeeds in part because the delegation of these choices is managed by iOS. We need to expand this approach by making it possible for consumers to delegate their data preferences

and access permissions to software-based agents beyond a single operating system or browser. In order to do so we must create the technical standards and infrastructure that allow consumers to engage with digital products and services, at least on a limited basis, without giving up their data or relying on a consent paradigm that forces them to make constant one-off, case-by-case decisions. There is ample research from the field of human-computer interaction that illustrates why such approaches are ineffective and burdensome.¹⁵¹

California has opened the door to this approach with its explicit acceptance of a browser opt-out signal, Global Privacy Control (GPC), that can function as an automatic opt-out for the CCPA's "do not sell my personal information" provision. GPC is an example of the direction we must move in if we wish to curb data collection by default: delegating and enforcing data collection preferences to software rather than to individuals on a constant basis. Transitioning to GPC or any other software-based solution cannot occur until one is adopted as an official W3C standard and laws enforce its acceptance.¹⁵² Even so, GPC is but one example of the direction we must take; we are overdue for a re-architecting of the personal data ecosystem, which we discuss below.

While California's (and Colorado's) laws include universal opt-outs for the sale of personal information via an automated browser signal¹⁵³ like Global Privacy Control,¹⁵⁴ opt-outs are limited to the use of data for targeted advertising purposes, and/or information that is "sold" by the collector as defined by the statute. The opt-outs are not required to be enabled by default, they only apply to internet browsers, and not all browsers support them.¹⁵⁵ The CCPA, in particular, gives companies a means to not adopt an opt-out signal if they offer multiple means for consumers to use

non-automated opt-outs, but this approach places the burden on consumers to exercise these rights, which few are likely to do given the burden of doing so.

Today, the incentives are such that when there are no restraints placed on data collection and use, companies will collect as much data as they possibly can using any means possible. This is the lesson from the U.S. experiment in self-regulation of the data economy, and it is a lesson in failure from the perspective of privacy. Without regulation there is no reason for these incentives to change, but even the GDPR demonstrates that data protection isn't necessarily enough as long as there continue to be weaknesses in the framework. In order for the data minimization and purpose limitation aspects of the FIPs to be fully effective, the consent piece of the framework must also function to give individuals meaningful, not de minimis, consent. If we fail to address the weaknesses in the consent model, and do not provide a right to refuse or mitigate data collection, particularly by third party collectors, we risk losing even more ground on privacy and any semblance of control over our personal data in the face of rising demand from AI.

Some will respond to this argument that adopting these measures will throw a wrench in the data economy and stall, if not destroy, progress in developing AI. We disagree. Pushing back against the status quo of ubiquitous data collection will not cripple data intensive industries. It might slow the rate of some progress in AI, though this may be a feature, not a bug. Many have raised objections about the rate at which AI is presently developing, with considerable concerns about safety. Though data is not a tangible resource and may not have the same short-term material impacts on the environment when exploited,¹⁵⁶ it does impact both humans and human

rights. The utilization of data can beget real world harms. To date, personal data has been treated like an inexhaustible resource to be harvested at will and exploited as desired. As more features of human life are mediated through technology and more people have access to the internet, we have seen data's impacts on society increase significantly over the past two decades. Online data collection and use is no longer limited to influencing purchasing behavior; data use by governments and private companies can impact not only civil rights, but also the functioning of democratic governments. As both technology use and AI continue to grow and spread, the need to address these fundamental flaws in the FIPs framework only become more urgent.

Suggestion 2: Focus on the AI data supply chain to improve privacy and data protection

Summary: Regulating the AI data supply chain must be a focus of any regulatory system that addresses data privacy.

The emergence of AI, particularly generative AI, presents a test for the FIPs-based privacy and data protection frameworks, especially the GDPR. While existing regulations may provide some oversight of the data collected and processed directly by an AI system, there is less clarity regarding oversight of the training data pipeline. For example, there are presently no transparency or documentation requirements in the EU¹⁵⁷ or the United States for companies to report the provenance of their training data, to document how it conforms with the principles of data minimization and purpose specification, or to otherwise address quality issues that could lead to downstream negative privacy outcomes such as a leakage of personally identifiable information in a system's outputs.¹⁵⁸

While existing regulations may provide some oversight of the data collected and processed directly by an AI system, there is less clarity regarding oversight of the training data pipeline.

Furthermore, existing regulations do not address how individuals can learn whether their data is included in a company’s training datasets, what to do if that data is inadvertently revealed by an AI system in its outputs, or whether individuals can request deletion of it and then verify its removal. With generative AI systems, there is the additional complexity of untangling whether an output of personal identifiable information was hallucinated or based on inferences made by the model without the actual data in the training dataset.¹⁵⁹

Understanding the data development pipeline becomes even more complex when companies source AI “models-as-a-service” from other businesses. In such instances, the relationship between the consumer and both the training data and model is a step removed. But as the U.S. Federal Trade Commission recently warned, that arm’s-length relationship does not provide exemption from liability with regard to AI companies’ privacy commitments to the consumers whose data power the models, and potentially to the business utilizing the model(s) as well.¹⁶⁰ Not only do privacy implications arise from using models trained on data of questionable or unknown provenance but

also with the collection of data from consumers using a system operating on the underlying model. For example, will that data be ingested by the frontline system to be used for training or fine-tuning purposes of the primary model? If the answer is yes, how will consumers be properly informed in order to assert their due process rights?

These issues raise the need for a data governance framework that is aligned with data privacy concerns that goes beyond the individual rights provided by the FIPs. While individuals and their personal data are most certainly implicated in these governance questions, they also raise societal level concerns that aren’t captured by considering these activities solely from an individual rights perspective.

Why Training Data Matters

Training data is AI’s oxygen; an AI system cannot exist without it. Data quality and training set size are important inputs that impact the outputs of an AI system. Defining data quality can be a slippery concept, but it generally refers to the accuracy and relevance of the data used to train the system given the system’s goals.¹⁶¹ “Dirty” data can include mislabeled data, biased data, error-ridden data, and data that is inadequate for representing the problem in question, to name a few.¹⁶²

Estimating the desired size of a training dataset isn’t an exact science; while there is evidence that larger datasets appear to improve system capabilities,¹⁶³ larger doesn’t definitionally mean better. For example, an LLM such as ChatGPT has apparently benefitted from its massive dataset in terms of the system’s quality of response that mimics human reasoning, but certainly not in terms of the accuracy of its responses.¹⁶⁴ In contrast, a relatively small synthetic dataset such as TinyStories using words and structure appropriate for young children has been successfully

used to train small language models in coherent English.¹⁶⁵ The creators argue that there are distinct advantages (namely interpretability and coherence of output) of a simpler approach. Similarly, UC Berkeley’s Koala chatbot research prototype, built atop Meta’s LLaMA model, uses a curated dataset that the creators claim offers results that are competitive with both Stanford’s Alpaca and ChatGPT.¹⁶⁶ The authors argue that their findings “suggest [...] that models that are small enough to be run locally can capture much of the performance of their larger cousins if trained on carefully sourced data. This might imply, for example, that the community should put more effort into curating high-quality datasets, as this might do more to enable safer, more factual, and more capable models than simply increasing the size of existing systems.”

Quality and size do not necessarily have a direct relationship—meaning a large dataset does not guarantee either higher or lower output quality. However, given the time and expense required to create a high-quality dataset, it is likely that generative systems in particular that use large datasets have more quality issues.¹⁶⁷ For very large models such as LLMs, the size of the training dataset has been decisively traded off for quality. Broadly scraping the internet will yield content that is severely biased, toxic, inaccurate, spam-laden, or includes “by-catch” such as personal data.¹⁶⁸ Going big on data yields benefits in some contexts, but scraping the world’s data should not be thought of as the inevitable (and only) path for creating advanced AI systems.

Privacy Issues with the Dataset Pipeline

Khan and Hanna, in their paper decomposing the components of dataset development, helpfully identify two types of actors—data subjects and model subjects—in the dataset development pipeline.¹⁶⁹ *Data*

subjects are the individuals whose data was collected and included in the training data; *model subjects* are those subject to the decisions of the downstream model.¹⁷⁰ Thus far, regulators have been focused more on the privacy issues raised for model subjects than data subjects, though the release of generative AI systems to the public surfaced the issue of what data protections applied to data subjects. While the bulk of the attention placed on AI systems has been with system outputs, system inputs also matter. The inclusion of personal or identifiable data in training datasets not only makes it possible that a model may memorize and output that data, it also raises the issue of consent. Do data subjects know that their data was included in a system’s training data?¹⁷¹ Were they asked for consent before being included? What rights do they have to request exclusion or deletion from these datasets? Can individuals have their personal data deleted from a model? And how do individuals even proceed in discovering whether their data was included?

Training datasets present new challenges for regulators. The assumption in the FIPs-based rights framework is that individuals are engaged in a first-party data relationship where the individual knows who the data collector is and who may have consented (albeit perhaps under less than ideal circumstances, such as through notice and consent) to having their data collected. This assumption is one reason FIPs-based privacy and data protection regulations have been inadequate in fighting data collection by third parties such as advertising networks or data brokers. Data that is scraped from unknown online sources, purchased or traded from data brokers, or reused by a data processor that had consent for the original purpose for which it was collected but not the subsequent reuse, might all be components of training datasets. While the GDPR may provide EU data

protection regulators with the means to curb some of these uses of data, presently in the United States these uses are broadly unregulated, and it is unlikely that an ADPPA-esque law would prevent all of them. Even so, without more specific data minimization and purpose limitation rules, or limits targeted at data gathered or reused for training purposes, individuals will continue to bear the burden of identifying where their data may be ingested, submitting deletion requests, and otherwise shoulder the bulk of the labor required to maintain their privacy.

Arguably there are potential harms and risks for data subjects raised by training datasets. The paradigm of data collection by default means that today individuals are having their personal information pulled into model development without their explicit knowledge. There are also complex technical issues around the question of whether one can request the deletion of one's data from a trained model without requiring the destruction of the model itself. To date, the FTC has required both data and model deletion in a few cases where they argued that a company used customer data without consent.¹⁷² Research by our Stanford colleague Professor James Zou demonstrates that "approximate data deletion" is a potential method for deleting data from models without requiring retraining, though whether this approach has practical applications is an open question.¹⁷³ Companies will be incentivized to argue that data deletion from trained models cannot be accomplished to avoid the cost of complying. Further, without transparency requirements for documenting training data, including inputs for retraining, being able to determine whether one's personal data was used to train (or retrain) a model will be a challenge.

Transparency and Governance for the Data Development Pipeline

The gaps in FIPs-based frameworks with respect to the privacy issues stemming from the creation and use of training data must be filled by a strategy that looks beyond individual rights to address them. Policymakers must establish new regulations or standards across the data development supply chain that create and mandate dataset documentation and transparency processes. These should include documenting the provenance or source of any data used for training and, if the data is related to an individual, include whether it was obtained with consent.¹⁷⁴ This approach could incentivize the creation of new technical standards for data subjects to uniquely identify their own data and designate its appropriate uses, a strategy that could also have copyright management implications as well.

Focus on the data supply chain

The principles of data minimization and purpose limitation should be part of any strategy to address dataset development. However, data minimization in particular is a concept that may prove challenging for regulators to enforce beyond more egregious abuses unless it is given more specificity. Similarly, developers should be able to demonstrate that their uses of data

Policymakers must establish new regulations or standards across the data development supply chain that create and mandate dataset documentation and transparency processes.

for training purposes are consistent with what was disclosed and consented to at the point of collection, without the use of obfuscatory tactics such as broad, indeterminate disclosures (i.e., “We use your personal information to improve our products and services”).

But these two principles are not enough. A supply chain approach for AI data governance would bring greater transparency to the entire life cycle of dataset development and foreground the need for responsible data management from creation to deletion (including highlighting the need for deletion when appropriate).¹⁷⁵ To be clear, we are not arguing for data governance for manufacturing supply chains, but rather to treat the process for building datasets used to train AI as a supply chain. However, transparency is not the only goal; attention to quality will both improve the performance of AI systems and resolve some of the issues with bias and poor generalizability, which will in turn increase trustworthiness.¹⁷⁶

There is an emergent field of dataset documentation practices that can aid in determining data provenance, tracking consent, and providing greater transparency about the source of the data used to train AI models.¹⁷⁷ But it is important to stress that these processes are still immature and policymakers should not mandate specific approaches at this stage without more research and experimentation, such as through regulatory sandboxes, to understand which forms of documentation clearly address privacy risks and harms. Should algorithmic impact assessments gain steam as a governance tool, they could require a form of dataset assessment as well to fully understand how dataset creation and management practices influence algorithmic development. It is also possible that data protection impact assessments are broadened in scope to include the impacts of datasets themselves, and that particular categories or types of data could

trigger higher risk classifications and obligations. As one such example, California is considering requiring risk assessments for certain high-risk category automated systems when processing training data that includes consumers’ personal information.¹⁷⁸

However, critical questions remain: What forms of documentation will be the most effective in achieving the goals of protecting both individual and societal data privacy? Are there other portions of the dataset supply chain, such as data labeling, that deserve more scrutiny? And, importantly, under what circumstances will these forms of documentation identify systems that should not be built or deployed? Unfortunately, we are not yet at a stage where the answer to this last question is clearly apparent. It is certainly true that there is value in creating meaningful data compliance measures by requiring dataset creators to adopt controls and responsible practices for collecting and managing data to at a minimum provide greater transparency into the creation process. This is an area where regulatory sandboxing would be an appropriate approach to help determine the types of documentation and processes that would provide regulators and researchers with greater transparency without creating an onerous burden on businesses.

Incentivize responsible practices

Any approach to resolving the privacy issues raised by training data must address existing incentives in the data marketplace to obtain data cheaply or unethically. While mandating compliance requirements can shift both business practice and culture, doing so is unlikely to address the incentives that drive a race to the bottom with anti-privacy data practices. In short, it’s hard to compete with free and unregulated data, particularly when competition exists across legal jurisdictions.

Shifting the paradigm of data collection by default to one where data collection requires permission and potentially negotiation is the prohibitive strategy; the incentive-based strategy would, ideally, make ethically sourced data cheaper, less risky, and of higher quality than either scraped data or data obtained from the third party data collection ecosystem. That is why adopting a purely regulatory approach may have limited effectiveness. We need public and private investment in ethically sourced datasets, not only to disincentivize scraping but also to open the data ecosystem to a broader set of actors than the largest tech companies who have a near monopoly on consumer data. We also need investment in the technical underpinnings of the data infrastructure that could support ethical data sourcing. This means creating new open standards for personal data management, as well as new legal vehicles for managing, pooling, and licensing data. Synthetic data, too, is an option for some contexts, as is using technical privacy measures, such as differential privacy or homomorphic encryption, to share or access data in non-identifiable or secure ways.

Finally, we need to make public data more accessible for research and industry uses, as we and our Stanford colleagues argue would come from creating a National AI Research Resource.¹⁷⁹ Of the many ways this could be accomplished, one idea is to support the creation of publicly available datasets that contend with issues of not only privacy and consent but ideally intellectual property as well. Data is as valuable a resource for economic development and human flourishing as our natural resources; accordingly, public investment in data resources can bring social benefits as well as create the conditions for ethical data use.

We need public and private investment in ethically sourced datasets, not only to disincentivize scraping but also to open the data ecosystem to a broader set of actors.

Suggestion 3: Flip the script on the management of personal data

Summary: Support the development of data intermediaries as a way to both support and automate the exercise of individual data rights and preferences, as well as collective privacy.

Nearly 30 years have elapsed since the creation of the commercial internet, and yet the fundamentals of online data exchange remain largely unchanged, specifically the unbalanced flow of data from consumers to companies. New vectors for data collection (e.g., third party cookies, smartphone APIs) have emerged since the late 1990s, but the basic paradigm of paying for access to online products and services with our data has only become more ossified. Our rapid transition from AI winter to the current, urgent AI gold rush has largely been enabled by the massive quantity of data held by private companies and across the publicly accessible internet. Without the access to data that researchers—and more impactfully private companies—have had over the past decade, the trajectory of AI development that we are witnessing would arguably be far less rapid and pronounced.

And therein lies the conundrum in our current AI race: AI is progressing at an enormous clip, in large part thanks to the availability and quantity of the data we have all generated. But we, individually and collectively, were not asked if we wished to contribute our data to this experiment, one that has already caused and will likely continue to cause significant harm. And even if those concerns prove to be overblown, whether any of us, or the public at large, materially or personally benefits from AI—as compared to those who are developing it and stand to profit handsomely from its growth—is undetermined. Free AI services may be of little value if you’ve lost your livelihood to AI automation. As we have argued at Stanford HAI, the private, for-profit development of AI is far more likely to result in applications that benefit the developers rather than society at large.¹⁸⁰ Without sufficient public investment in basic AI research and development, we are far less likely to see public, nonprofit uses of AI that materially benefit society.

Over the past decade, the trading of one’s data for access to free online services has drawn increased scrutiny and critique from the public.¹⁸¹ The explosion of generative AI over the past year ratcheted up those concerns once it became clear that the massive datasets required by generative systems were built on data that in some cases was obtained from any source possible. The public is increasingly aware—and often displeased—with the hidden cost of free products: paying with your data, a trade-off that isn’t always in one’s best interest. In particular, there is growing concern that one’s data isn’t just used as payment but also to target and manipulate one’s behavior and choices.

Increasingly, there are calls for data equity in the form of compensation for the use of personal data. Former presidential candidate Andrew Yang even made data

equity a platform in his 2020 campaign.¹⁸² Startups are forming to capitalize on this trend by making the exchange of data for access to online resources both explicit and direct by attempting to provide users with direct compensation for their data use. However, this approach is typically focused on selling individuals’ data, which may yield a short-term gain but ultimately lead to alienation and further, irrevocable loss of control over data.¹⁸³ Instead of focusing on property rights, the focus of these efforts should be on developing a permissioning or licensing regime for individuals’ data through regulated data intermediaries, allowing them to reap any direct benefits while maintaining control over its use.¹⁸⁴

We argue above that data collection by default is one of the key issues we must address if we are to have a fighting chance at improving our data privacy as AI development fuels the demand for personal data. Below, we elaborate on the two components that can aid with that goal: facilitating the creation of data intermediaries and building the technical architecture to support personal data licensing. These solutions are not quick, short-term fixes to the problems we discuss above, but instead take a longer view of the structural changes we need to make to our digital ecosystem as we face a world with ever larger appetites for data.

Data Intermediaries & Data Permissioning Infrastructure

Data intermediaries are a core component of “flipping the script” on data collection to move away from a digital landscape where data is acquired by digital services to one where consumers decide the terms by which they will allow companies to use their data. The goal of a data intermediary is to facilitate and manage “data relations between data rights holders (such as people or businesses), depending on the parties’ relationships, intentions and resources. They

do so by encapsulating, communicating and enacting the shared interests of the relevant parties and safeguarding their interests. At their most basic level they facilitate the exchange of information; at their most sophisticated they can assume decisionmaking, including on behalf of people.¹⁸⁵ They can take many forms: data trusts, collaboratives, cooperatives, commons, stewards; for-profit, nonprofit, and publicly owned.

A key goal for the consumer data space is to create an intermediary that can be entrusted to share data according to the individual's preferences, potentially even negotiating the terms, and to do so by delegating one's preferences to a software system that handles transactions without constant micromanaging and intervention by the individual. Instead of one's data being collected by hundreds or even thousands of distinct entities, a data intermediary could manage these sharing relationships and ensure that one's data is being used according to negotiated terms. Instead of a data ecosystem where one's data is spread hither and yon, and one has zero control over where it goes or how it is used, the data intermediary model centers the individual and their preferences for data sharing, ensuring that they maintain control and benefits flow directly to them.

While we are focused primarily on the consumer data space here, data intermediaries could manage personal data across multiple contexts: healthcare, genetics, education, research, and so on. They can also offer the power and leverage that come from collective bargaining; individuals today are disempowered in their data relationships with companies. Data intermediaries can provide the strength in numbers for individuals to negotiate in their best interests. In order to do so, intermediaries would need to be regulated and hold fiduciary responsibilities for their users'

Data intermediaries can provide the strength in numbers for individuals to negotiate in their best interests.

data. Accordingly, they may be better positioned to act in consumers' best interests than placing similar obligations on technology platforms for whom the customers' best interests will always be in tension with the company's.

However, without a technical layer to facilitate intermediary relationships, this vision will, at best, grow very slowly, and certainly not fast enough to stem the ever-increasing tide of data collection. This is not an insignificant challenge—as evidenced by data portability mandates. While data portability requirements have been on the books since the adoption of the GDPR, the vision of being able to take one's social media data specifically and move it to a competing service remains largely unfulfilled because it is a complex technical problem: taking proprietary data that is potentially co-managed (e.g., threaded posts) and sorting out who is the primary “owner” and how to make such data interoperable between services. The intermediary concept takes the challenging goal of data portability and turns it on its head, focusing on the individual first as the atomic unit rather than with disparate platforms' proprietary systems.

The case of social media interoperability may in fact be a far more difficult scenario than attempting to create a standard for the exchange of a core set of personal data. Sir Tim Berners-Lee has been working on this

precise problem for several years; his company, Inrupt, is creating an open-source standard for information “pods” (called Solid¹⁸⁶) to create the infrastructure to overcome this challenge.¹⁸⁷ Presently there is more focus on building data intermediaries between businesses than for the consumer marketplace. But until we see policymakers advocate for or incentivize consumer-centered solutions, we are left with the more piecemeal approaches built on exercising data rights. Earlier in this paper, we discussed the Global Privacy Control proposed standard and the need for policymakers to mandate its adoption by browser creators. Another example of using existing data rights to push for collective action is by Consumer Reports, which created the app Permission Slip to use California’s CCPA Do Not Sell and Delete My Data rights, which includes a provision to allow individuals to designate an agent to exercise these rights. Rather than having to identify and contact companies one by one, the app facilitates making these requests for consumers from a single location. These examples show the initial promise of automating this type of data control, but the possible solutions could be much more robust and advanced than these initial strategies.

Shifting the data ecosystem away from data collection by default will require more than policy change. It will require new legal entities for data governance, such as data intermediaries, with clearly defined duties of care such that we do not inadvertently create a new class of data brokers (or we create them with tight regulations and high ethical standards). It will require technical infrastructure that can enable (and incentivize) ethical, human-centered data exchanges that respect user consent and usage preferences. It may also require reopening the digital rights management debates of the 2000s to allow individuals to protect both their data privacy as well as intellectual property rights with content they share online. But, as we discussed

earlier, it will also require an investment in public data resources as well, so the value of large datasets does not rest solely in the hands of private actors.

Chapter 5: Conclusion

As we rapidly transition from an internet era into an AI era, not surprisingly, regulation continues to lag behind advances in technology. The privacy and data protection laws already in place or in development today do and will continue to impact the growth and use of AI systems, through both implicit privacy protections and explicit measures addressing automated decision-making. But they are a reaction to the business models of the past, not the future. They were developed in response to the privacy threats that emerged over the last 20 years of the commercial internet, before the widespread deployment of powerful AI technologies.

Today, the data protection and due process rights that existing regulations offer are jeopardized by emergent technologies such as generative AI. Data collection practices, such as data scraping or third party data collection, conflict with many regulatory assumptions pertaining to how companies collect data from and about individuals. In addition to increasing the demand for data, AI systems also create new avenues for privacy harm through novel approaches to data collection and output. Existing regulations and frameworks also do not consider the ways in which privacy risks and harms from data can be relational and social in nature. Data can implicate individual privacy through inferences made through others' data, and through known or inferred relationships. And widespread data collection and surveillance can present societal level risks and harms that individually focused regulations simply cannot address.

One thing we do know with certainty is that AI needs data to advance. Without data, AI advancement grinds to a halt. As AI innovation continues apace, it is therefore crucial that we get the data piece of AI development “right” by centering individual and

collective privacy harms. Our focus throughout this paper has been to highlight the need to focus on the data life cycle that feeds AI development as a mechanism for transparency and accountability—while addressing the data privacy issues raised by AI that cannot be resolved through the exercise of individual data rights.

Our goal is to provide policymakers and other stakeholders with sufficient background information to understand why existing privacy regulations and frameworks do not fully address these problems, and to offer suggestions for short and longer term actions to take to protect and preserve privacy while also ensuring greater transparency and accountability in the AI data development life cycle. There are many other options available to address these issues; this work will certainly not be the final word on them. But we hope that we have provided sufficient rationale as to why they must be addressed if we are to have both privacy and AI moving forward.

We agree with scholars who have argued that nothing about AI development is inevitable.¹⁸⁸ We also don't think that privacy is dead or that a lack of protection of individual and collective data rights is a foregone conclusion.¹⁸⁹ While data can be stored indefinitely, it isn't necessarily permanent, and it certainly ages, often poorly, with time. The infrastructures that we build to support it aren't immutable, and aren't necessarily resilient in the face of change or catastrophe. Even the foundations upon which many are presently building foundation models are mutable. We are the creators of technology, and we can shape it to support and reflect human values, rather than to undo them. In sum: the future of AI isn't yet written; we can choose how we want it to unfold.

Endnotes

1 Rishi Bommasani, Christie M. Lawrence, Lindsey A. Gailmard, Caroline Meinhardt et al., “Decoding the White House AI Executive Order’s Achievements,” *Stanford Institute for Human-Centered Artificial Intelligence*, November 2, 2023, <https://hai.stanford.edu/news/decoding-white-house-ai-executive-orders-achievements>.

2 See Lina M. Khan, “Joint Statement on Enforcement Efforts Against Discrimination and Bias in Automated Systems,” *Federal Trade Commission*, April 25, 2023, <https://www.ftc.gov/legal-library/browse/cases-proceedings/public-statements/joint-statement-enforcement-efforts-against-discrimination-bias-automated-systems>; California Legislature, An act to amend, repeal, and add Section 35 of the Code of Civil Procedure, and to amend, add, and repeal Section 20010 of the Elections Code, relating to elections, California Assembly Bill No. 730, October 3, 2019, https://leginfo.ca.gov/faces/billTextClient.xhtml?bill_id=201920200AB730, Chapter 493; New York City Council, A Local Law to amend the administrative code of the city of New York, in relation to automated employment decision tools, New York City Council Int. No. 1894-A, December 11, 2021, <https://legistar.council.nyc.gov/LegislationDetail.aspx?ID=4344524&GUID=B051915D-A9AC-451E-81F8-6596032FA3F9&Options=Advanced&Search>, Subchapter 25.

3 When we use the term ‘provenance,’ we refer to the general sourcing of data. This is a different use than that of groups such as Content Credentials (<https://contentcredentials.org/>), which aim to ensure trustworthy data creation. It is possible that efforts like these could eventually encompass a broader concept of provenance to include metadata that will allow for tracking the source of created data for supply purposes.

4 Colleen McClain, Michelle Faverio, Monica Anderson, and Eugenie Park. “How Americans View Data Privacy,” *Pew Research Center*, October 18, 2023. <https://www.pewresearch.org/internet/2023/10/18/how-americans-view-data-privacy/>.

5 Jennifer King, “Privacy, Disclosure, and Social Exchange Theory,” *University of California Berkeley*, Spring 2018, <https://escholarship.org/uc/item/5hw5w5c1>.

6 National Institute of Standards and Technology, “Artificial Intelligence Risk Management Framework (AI RMF 1.0),” January 26, 2023, <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf>.

7 Information Commissioner’s Office, “What is automated individual decision-making and profiling?” <https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/individual-rights/automated-decision-making-and-profiling/what-is-automated-individual-decision-making-and-profiling/>.

8 Nestor Maslej, Loredana Fattorini, Erik Brynjolfsson, John Etchemendy et al., “The AI Index 2023 Annual Report,” *Stanford Institute for Human-Centered Artificial Intelligence*, April 2023, <https://aiindex.stanford.edu/report/>, Chapter 6.

9 Khan, “Joint Statement on Enforcement Efforts Against Discrimination and Bias in Automated Systems.”

10 Katrina Zhu, “The State of State AI Laws: 2023,” *Electronic Privacy Information Center*, August 3, 2023, <https://epic.org/the-state-of-state-ai-laws-2023/>.

11 Raphael Gellert and Serge Gutwirth, “The legal construction of privacy and data protection,” *Computer Law & Security Review* 29(5), October 2013, <https://doi.org/10.1016/j.clsr.2013.07.005>, 522-530.

12 Gellert et al., “The legal construction of privacy and data protection.”

13 For a general background, see Robert Gellman, “Fair Information Practices: A Basic History – Version 2.22,” *SSRN*, April 2022, <https://doi.org/10.2139/ssrn.2415020>.

14 Chris Jay Hoofnagle, “The Origin of Fair Information Practices: Archive of the Meetings of the Secretary’s Advisory Committee on Automated Personal Data Systems (SACAPDS),” *SSRN*, July 15, 2014, <http://dx.doi.org/10.2139/ssrn.2466418>.

15 U.S. Secretary’s Advisory Committee on Automated Personal Data Systems, “Records, Computers and the Rights of Citizens, Chapter IV: Recommended Safeguards for Administrative Personal Data Systems,” June 30, 1973, <https://aspe.hhs.gov/reports/records-computers-rights-citizens>; See also Hoofnagle, “The Origin of Fair Information Practices.” The original FIPs are: A) There must be no personal-data record-keeping systems whose very existence is secret; B) There must be a way for an individual to find out what information about him is in a record and how it is used; C) There must be a way for an individual to prevent information about him obtained for one purpose from being used or made available for other purposes without his consent; D) There must be a way for an individual to correct or amend a record of identifiable information about himself; E) Any organization creating, maintaining, using, or disseminating records of identifiable personal data must assure the reliability of the data for their intended use and must take reasonable precautions to prevent misuse of the data.

16 Mark MacCarthy, *Regulating Digital Industries: How Public Oversight Can Encourage Competition, Protect Privacy, and Ensure Free Speech*, Brookings Institution Press, 2023.

17 For an extended discussion on the origin of the FIPs, see Gellman, “Fair Information Practices.”

18 Private actors were largely ignored in the development of the FIPs, in part because in the early 1970s, it was presumed that no private actor could ever collect as much data as a government entity.

19 Organisation for Economic Co-operation and Development, “Recommendation of the Council concerning Guidelines Governing the Protection of Privacy and Transborder Flows of Personal Data,” September 1980, <https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0188>.

20 Gellman, “Fair Information Practices,” p.24.

21 The 2013 OECD Principles include both purpose specification and use limitation principles, while the GDPR uses the term purpose limitation. While we acknowledge that there are subtleties between the two terms, this paper will use the term purpose limitation as specified in the GDPR, see European Parliament and European Council, Regulation 2016/679 of the European Parliament and of the Council [General Data Protection Regulation, hereinafter GDPR], April 27, 2016, <https://eur-lex.europa.eu/eli/reg/2016/679/oj>, Article 5(1)(b).

22 Chris Jay Hoofnagle, Bart van der Sloot, and Frederik Zuiderveen Borgesius, “The European Union general data protection regulation: what it is and what it means,” *Information & Communications Technology Law* 28(1), February 2019, <https://www.tandfonline.com/doi/full/10.1080/13600834.2019.1573501>, 65-98.

- 23 European Parliament and European Council, GDPR, Article 4(1): “personal data’ means any information relating to an identified or identifiable natural person (‘data subject’); an identifiable natural person is one who can be identified, directly or indirectly, in particular by reference to an identifier such as a name, an identification number, location data, an online identifier or to one or more factors specific to the physical, physiological, genetic, mental, economic, cultural or social identity of that natural person.”
- 24 European Parliament and European Council, GDPR, Article 22.
- 25 Jennifer King, Caroline Meinhardt, Abel Ribbink, Pete Warden et al., “Response to the Request for Comments on Trade Regulation Rule on Commercial Surveillance and Data Security,” *Stanford Institute for Human-Centered Artificial Intelligence*, November 21, 2022, <https://hai.stanford.edu/sites/default/files/2022-12/HAI%20-%20FTC%20ANPR%20Comments.pdf>.
- 26 Sebastião Barros and Gabriela Zanfir-Fortuna, “FPF Report: Automated Decision-Making Under the GDPR - A Comprehensive Case-Law Analysis,” *Future of Privacy Forum*, May 17, 2022, <https://fpf.org/blog/fpf-report-automated-decision-making-under-the-gdpr-a-comprehensive-case-law-analysis/>.
- 27 Information Commissioner’s Office, “Data protection impact assessments,” May 19, 2023, <https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/accountability-and-governance/guide-to-accountability-and-governance/accountability-and-governance/data-protection-impact-assessments/>.
- 28 117th U.S. Congress, American Data Privacy and Protection Act, H.R.8152, June 21, 2022, <https://www.congress.gov/bill/117th-congress/house-bill/8152>.
- 29 Cameron F. Kerry, “How privacy legislation can help address AI,” *Brookings*, July 7, 2023, <https://www.brookings.edu/articles/how-privacy-legislation-can-help-address-ai/>; The White House, Blueprint for an AI Bill of Rights, October 2022, <https://www.whitehouse.gov/ostp/ai-bill-of-rights/>.
- 30 Andrew Folks, “US State Privacy Legislation Tracker,” *International Association of Privacy Professionals*, December 8, 2023, <https://iapp.org/resources/article/us-state-privacy-legislation-tracker/#enacted-laws>.
- 31 George P. Slefo, “Marketers and Tech Companies Confront California’s Version of GDPR,” *AdAge*, June 29, 2018, <https://adage.com/article/digital/california-passed-version-gdpr/314079>.
- 32 Anupam Chander, Margot E. Kaminski, and William McGeveran, “Catalyzing Privacy Law,” *Minnesota Law Review* 105, 2021, https://scholarship.law.umn.edu/mlr/3305_1733.
- 33 California Privacy Protection Agency, “A New Landmark for Consumer Control Over Their Personal Information: CPPA Proposes Regulatory Framework for Automated Decisionmaking Technology,” November 27, 2023, <https://cppa.ca.gov/announcements/2023/20231127.html>.
- 34 Rogier Creemers and Graham Webster, “Translation: Personal Information Protection Law of the People’s Republic of China – Effective Nov. 1, 2021,” *Stanford DigiChina Project*, August 20, 2021, <https://digichina.stanford.edu/work/translation-personal-information-protection-law-of-the-peoples-republic-of-china-effective-nov-1-2021/>.
- 35 Alexa Lee, Mingli Shi, Qiheng Chen, Jamie P. Horsley et al., “Seven Major Changes in China’s Finalized Personal Information Protection Law,” *Stanford DigiChina Project*, September 15, 2021, <https://digichina.stanford.edu/work/seven-major-changes-in-chinas-finalized-personal-information-protection-law/>.
- 36 In its August 2021 IPO filing, AI giant SenseTime listed “complex and evolving laws, regulations and governmental policies regarding privacy and data protection” as a risk factor, citing the PIPL and Data Security Law, which were about to take effect. See Jane Li, “China’s new data laws are a risk factor in a facial-recognition giant’s IPO filing,” *Quartz*, August 30, 2021, <https://qz.com/2053040/chinas-new-data-laws-are-a-risk-factor-in-sensetimes-ipo-filing>.
- 37 Graham Webster, “Translation: Chinese Authorities Announce \$1.2B Fine in DiDi Case, Describe ‘Despicable’ Data Abuses,” *Stanford DigiChina Project*, July 21, 2022, <https://digichina.stanford.edu/work/translation-chinese-authorities-announce-2b-fine-in-didi-case-describe-despicable-data-abuses/>.
- 38 Mingli Shi, Jamie P. Horsley, and Xiaomeng Lu, “Forum: Unpacking the DiDi Decision,” *Stanford DigiChina Project*, July 22, 2022, <https://digichina.stanford.edu/work/forum-unpacking-the-didi-decision/>.
- 39 Inioluwa Deborah Raji, Peggy Xu, Colleen Honigsberg, and Daniel Ho, “Outsider Oversight: Designing a Third Party Audit Ecosystem for AI Governance,” *Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society*, July 2022, <https://dl.acm.org/doi/abs/10.1145/3514094.3534181>, 557–571.
- 40 Emanuel Moss, Elizabeth Anne Watkins, Ranjit Singh, Madeleine Clare Elish et al., “Assembling Accountability: Algorithmic Impact Assessment for the Public Interest,” *Data & Society*, June 29, 2021, <https://datasociety.net/library/assembling-accountability-algorithmic-impact-assessment-for-the-public-interest/>.
- 41 Brent Mittelstadt, “Interpretability and Transparency in Artificial Intelligence,” In *The Oxford Handbook of Digital Ethics*, edited by Carissa Véliz, Oxford University Press, October 20, 2021, <https://doi.org/10.1093/oxfordhb/9780198857815.013.20>.
- 42 Samantha Lai and Brooke Tanner, “Examining the intersection of data privacy and civil rights,” *Brookings*, July 18, 2022, <https://www.brookings.edu/articles/examining-the-intersection-of-data-privacy-and-civil-rights/>.
- 43 118th U.S. Congress, Algorithmic Accountability Act of 2023, H.R. 5628, September 21, 2023, <https://www.congress.gov/bill/117th-congress/house-bill/6580/text>. See also 117th U.S. Congress, American Data Privacy and Protection Act, Section 207(c).
- 44 Gopal Ratnam, “Data privacy law seen as needed precursor to AI regulation,” *Roll Call*, September 26, 2023, <https://rollcall.com/2023/09/26/data-privacy-law-seen-as-needed-precursor-to-ai-regulation/>.
- 45 Lauren Leffer, “Your Personal Information Is Probably Being Used to Train Generative AI Models,” *Scientific American*; Sara Morrison, “The Tricky Truth About How Generative AI Uses Your Data,” *Vox*, July 27, 2023, <https://www.vox.com/technology/2023/7/27/23808499/ai-openai-google-meta-data-privacy-nope>.
- 46 Kashmir Hill, *Your Face Belongs To Us: A Secretive Startup’s Quest To End Privacy As We Know It*, Random House, September 2023.
- 47 Natasha Lomas, “Clearview fined again in France for failing to comply with privacy orders,” *TechCrunch*, May 10, 2023, <https://techcrunch.com/2023/05/10/clearview-ai-another-cnll-gspr-fine/>.
- 48 Illinois General Assembly, Biometric Information Privacy Act, 740 ILCS 14/, October 3, 2008, <https://www.ilga.gov/legislation/ilcs/ilcs3.asp?ActID=3004&ChapterID=57>.

- 49 Gabriela Zanfir-Fortuna, “How Data Protection Authorities are De Facto Regulating Generative AI,” *Future of Privacy Forum*, September 12, 2023, <https://fpf.org/blog/how-data-protection-authorities-are-de-facto-regulating-generative-ai/>.
- 50 Office of the Privacy Commissioner of Canada, “Statement on Generative AI,” June 21, 2023, https://www.priv.gc.ca/en/opc-news/speeches/2023/s-d_20230621_g7/.
- 51 ACLU v. Clearview AI, May 11, 2022, <https://www.aclu.org/cases/aclu-v-clearview-ai>.
- 52 Pamela Samuelson, “Generative AI meets copyright,” *Science* 381, July 14, 2023, [10.1126/science.adi0656](https://doi.org/10.1126/science.adi0656), 158-161.
- 53 The FTC requested information on how training data was obtained, what categories of content were included in the corpus, who was responsible for reviewing the data, how reinforcement learning was used for training, the extent to which models produced information about individuals (including the accuracy of personal information), the extent to which fine-tuning was used to address these issues, and how personal information is ingested by the system. See Center for AI and Digital Policy, “In the Matter of Open AI (Federal Trade Commission 2023),” <https://www.caipd.org/cases/openai/>.
- 54 Tonya Riley, “The FTC’s biggest AI enforcement tool? Forcing companies to delete their algorithms,” *Cyberscoop*, July 5, 2023, <https://cyberscoop.com/ftc-algorithm-disgorgement-ai-regulation/>.
- 55 Lina M. Khan, “Lina Khan: We Must Regulate A.I. Here’s How,” *The New York Times*, May 3, 2023, <https://www.nytimes.com/2023/05/03/opinion/ai-lina-khan-ftc-technology.html>.
- 56 Luca Bertuzzi, “Italian data protection authority bans ChatGPT citing privacy violations,” *Euractiv*, March 31, 2023, <https://www.euractiv.com/section/artificial-intelligence/news/italian-data-protection-authority-bans-chatgpt-citing-privacy-violations/>.
- 57 Garante per la protezione dei dati personali [hereafter Garante], “Artificial intelligence: stop to ChatGPT by the Italian SA, Personal data is collected unlawfully, no age verification system is in place for children,” March 31, 2023, <https://www.garanteprivacy.it/web/guest/home/docweb/-/docweb-display/docweb/9870847#english>.
- 58 Garante, “Artificial intelligence: stop to ChatGPT by the Italian SA.”
- 59 Garante, “ChatGPT: Italian DPA notifies breaches of privacy law to OpenAI,” January 29, 2024, <https://garanteprivacy.it/home/docweb/-/docweb-display/docweb/9978020#english>.
- 60 Digital Curation Centre, Trilateral Research, and School of Informatics, The University of Edinburgh, “The Role of Data In AI: Report for the Data Governance Working Group of the Global Partnership of AI,” November 2020, <https://gpai.ai/projects/data-governance/role-of-data-in-ai.pdf>.
- 61 Mark MacCarthy, “In Defense of Big Data Analytics,” In *The Cambridge Handbook of Consumer Privacy*, edited by Evan Selinger, Jules Polonetsky, and Omer Tene, Cambridge University Press, 2018.
- 62 King et al., “Response to the Request for Comments on Trade Regulation Rule on Commercial Surveillance and Data Security.”
- 63 See Asia Biega and Michele Finck, “Reviving Purpose Limitation and Data Minimisation in Data-Driven Systems,” *Technology and Regulation*, 2021, <https://doi.org/10.26116/techreg.2021.004>, 44–61; Hongyi Wen, Longqi Yang, Michael Sobolev, and Deborah Estrin, “Exploring Recommendations Under User-Controlled Data Filtering,” *Proceedings of the 12th ACM Conference on Recommender Systems*, September 2018, <https://doi.org/10.1145/3240323.3240399>, 72-76.
- 64 According to Open AI’s technical report, GPT-4 was pre-trained “using both publicly available data (such as internet data) and data licensed from third-party providers.” According to Anthropic’s model card, Claude models “are trained on a proprietary mix of publicly available information from the Internet, datasets that we license from third party businesses, and data that our users affirmatively share or that crowd workers provide.” See OpenAI, “GPT-4 Technical Report,” *arXiv preprint*, March 15, 2023, <https://arxiv.org/pdf/2303.08774.pdf>; Anthropic, “Model Card and Evaluations for Claude Models,” July 8, 2023, <https://paperswithcode.com/paper/model-card-and-evaluations-for-claude-models>.
- 65 Kristen E. Busch, “Generative Artificial Intelligence and Data Privacy: A Primer,” *Congressional Research Service*, May 23, 2023, <https://crsreports.congress.gov/product/pdf/R/R47569>.
- 66 The draft EU AI Act includes in Annex IV (d): “where relevant, the data requirements in terms of datasheets describing the training methodologies and techniques and the training data sets used, including a general description of these data sets, information about their provenance, scope and main”
- 67 See Rishi Bommasani, Kevin Klyman, Shayne Longpre, Sayash Kapoor et al., “The Foundation Model Transparency Index,” *arXiv preprint*, October 19, 2023, <https://arxiv.org/abs/2310.12941>; Cat Zakrzewski, “FTC investigates OpenAI over data leak and ChatGPT’s inaccuracy,” *The Washington Post*, July 13, 2023, <https://www.washingtonpost.com/technology/2023/07/13/ftc-openai-chatgpt-sam-altman-lina-khan/>.
- 68 Danielle Keats Citron and Daniel J. Solove, “Privacy Harms,” *Boston University Law Review* 102, 2022, https://scholarship.law.gwu.edu/faculty_publications/1534/, 793.
- 69 Daniel J. Solove, “A Taxonomy of Privacy,” *University of Pennsylvania Law Review* 154, 2006, https://scholarship.law.upenn.edu/penn_law_review/vol154/iss3/1, 477.
- 70 Existing harms include: surveillance, secondary uses of data, loss of control over data, data security risks, increased identification, and accessibility that can cause intrusion. Hao-Ping Lee, Yu-Ju Yang, Thomas Serban von Davier, Jodi Forlizzi et al., “Deepfakes, Phrenology, Surveillance, and More! A Taxonomy of AI Privacy Risks,” *arXiv preprint*, October 11, 2023, <https://arxiv.org/abs/2310.07879>.
- 71 See Will Knight, “AI Chatbots Can Guess Your Personal Information From What You Type,” *Wired*, October 2023, <https://www.wired.com/story/ai-chatbots-can-guess-your-personal-information/>; Nicholas Carlini, Florian Tramèr, Eric Wallace, Matthew Jagielski et al., “Extracting Training Data from Large Language Models,” *30th USENIX Security Symposium*, June 15, 2021, <https://www.usenix.org/conference/usenixsecurity21/presentation/carlini-extracting>, 2633-2650; Lauren Leffer, “Your Personal Information Is Probably Being Used to Train Generative AI Models,” *Scientific American*, October 19, 2023, [https://www.scientificamerican.com.stanford.idm.oclc.org/article/your-personal-information-is-probably-being-used-to-train-generative-ai-models/](https://www.scientificamerican.com/stanford.idm.oclc.org/article/your-personal-information-is-probably-being-used-to-train-generative-ai-models/); Robin Staab, Mark Vero, Mislav Balunović, and Martin Vechev, “Beyond Memorization: Violating Privacy Via Inference with Large Language Models,” *arXiv preprint*, October 11, 2023, <https://arxiv.org/abs/2310.07298>.
- 72 Shoshana Zuboff, *The Age of Surveillance Capitalism: The Fight For A Human Future At The New Frontier of Power*, Profile Books, 2018.

- 73 Future of Privacy Forum, “Unfairness by Algorithm: Distilling the Harms of Automated Decision-Making,” December 11, 2017, <https://fpf.org/blog/unfairness-by-algorithm-distilling-the-harms-of-automated-decision-making/>.
- 74 See Cameron F. Kerry, “Protecting privacy in an AI-driven world,” *Brookings*, February 10, 2020, <https://www.brookings.edu/articles/protecting-privacy-in-an-ai-driven-world/>; Office of the Victoria Information Commissioner, “Artificial Intelligence and Privacy - Issues and Challenges,” August 2018, <https://ovic.vic.gov.au/privacy/resources-for-organisations/artificial-intelligence-and-privacy-issues-and-challenges/>.
- 75 M.R. Leiser and Cristiana Santos, “Dark Patterns, Enforcement, and the emerging Digital Design Acquis: Manipulation beneath the Interface,” *SSRN*, April 27, 2023, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4431048.
- 76 See Harry Surden, “Structural Rights in Privacy,” *SMU Law Review* 60, 2007, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1004675, 1605-1629; Woodrow Hartzog and Frederic D. Stutzman, “The Case for Online Obscurity,” *California Law Review* 101, 2013, <https://ssrn.com/abstract=1597745>, 1.
- 77 Kate Crawford and Jason Schultz, “Big Data and Due Process: Toward a Framework to Redress Predictive Privacy Harms,” *Boston College Law Review* 55, 2014, <https://ssrn.com/abstract=2325784>, 93.
- 78 Future of Privacy Forum, “Unfairness by Algorithm.”
- 79 Citron et al., “Privacy Harms.”
- 80 Nora A. Draper and Joseph Turow, “The corporate cultivation of digital resignation,” *New Media & Society* 21(8), March 8, 2019, <https://journals.sagepub.com/doi/full/10.1177/1461444819833331>, 1824-1839.
- 81 Bobbie Johnson, “Privacy no longer a social norm, says Facebook founder,” *The Guardian*, January 10, 2010, <https://www.theguardian.com/technology/2010/jan/11/facebook-privacy>.
- 82 Jeffrey Dastin, “Insight - Amazon scraps secret AI recruiting tool that showed bias against women,” *Reuters*, October 10, 2018, <https://www.reuters.com/article/us-amazon-com-jobs-automation-insight/amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK08G/>.
- 83 Andrew Lee Park, “Injustice Ex Machina: Predictive Algorithms in Criminal Sentencing,” *UCLA Law Review* 19, February 19, 2019, <https://www.uclalawreview.org/injustice-ex-machina-predictive-algorithms-in-criminal-sentencing>.
- 84 Will Douglas Heaven, “Predictive policing algorithms are racist. They need to be dismantled,” *MIT Technology Review*, July 17, 2020, <https://www.technologyreview.com/2020/07/17/1005396/predictive-policing-algorithms-racist-dismantled-machine-learning-bias-criminal-justice>.
- 85 Ratnam, “Data privacy law seen as needed precursor to AI regulation.”
- 86 Without clear operational criteria for purpose limitation and data minimization, interpreting these principles in practice will be dependent upon how a company’s practices are stated in their privacy policy. If a company’s statements are ambiguous (as many are today) regarding how data might be repurposed, companies will have too much latitude to interpret these principles as they wish.
- 87 See King et al., “Response to the Request for Comments on Trade Regulation Rule on Commercial Surveillance and Data Security”; Divya Shanmugam, Samira Shabaniyan, Fernando Diaz, Michèle Finck et al., “Learning to Limit Data Collection via Scaling Laws: A Computational Interpretation for the Legal Principle of Data Minimization,” *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, June 21, 2022, <https://doi.org/10.48550/arXiv.2107.08096>, 839-849.
- 88 Reuben Binns and Valeria Gallo, “Data Minimisation and Privacy-Preserving Techniques in AI Systems,” *Information Commissioner’s Office*, August 21, 2021, <https://ico.org.uk/about-the-ico/media-centre/ai-blog-data-minimisation-and-privacy-preserving-techniques-in-ai-systems/>.
- 89 Biega et al., “Reviving Purpose Limitation and Data Minimisation in Data-Driven Systems.”
- 90 Shanmugam et al., “Learning to Limit Data Collection via Scaling Laws.”
- 91 Woodrow Hartzog, “The Inadequate, Invaluable Fair Information Practices,” *Maryland Law Review* 76, 2017, <https://www.marylandlawreview.org/volume-76-issue-4-symposium/the-inadequate-invaluable-fair-information-practices>, 952.
- 92 King et al., “Response to the Request for Comments on Trade Regulation Rule on Commercial Surveillance and Data Security.”
- 93 King, “Privacy, Disclosure, and Social Exchange Theory.”
- 94 Daniel J. Solove, “Introduction: Privacy Self-Management and the Consent Dilemma,” *Harvard Law Review* 126, May 2013, <https://harvardlawreview.org/print/vol-126/introduction-privacy-self-management-and-the-consent-dilemma/>, 1880.
- 95 MacCarthy, *Regulating Digital Industries*.
- 96 European Parliament and European Council, GDPR, Article 6 (a-f).
- 97 European Commission, “What does ‘grounds of legitimate interest’ mean?” https://commission.europa.eu/law/law-topic/data-protection/reform/rules-business-and-organisations/legal-grounds-processing-data/grounds-processing/what-does-grounds-legitimate-interest-mean_en.
- 98 Since Brexit, the U.K. has observed its own (albeit nearly identical) version of the GDPR. See Information Commissioner’s Office, “The UK GDPR,” <https://ico.org.uk/for-organisations/data-protection-and-the-eu/data-protection-and-the-eu-in-detail/the-uk-gdpr/>.
- 99 Information Commissioner’s Office, “When can we rely on legitimate interests?” <https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/lawful-basis/legitimate-interests/when-can-we-rely-on-legitimate-interests/>.
- 100 Sam Schechner and Jeff Horwitz, “Meta to Let Users Opt Out of Some Targeted Ads, but Only in Europe,” *The Wall Street Journal*, March 30, 2023, <https://www.wsj.com/articles/meta-to-let-users-opt-out-of-some-targeted-ads-but-only-in-europe-44b20b6d>.
- 101 Noyb, “Meta (Facebook, Instagram) switching to ‘Legitimate Interest’ for Ads,” March 30, 2023, <https://noyb.eu/en/meta-facebook-instagram-switching-legitimate-interest-ads>.

102 Meta Newsroom, “Facebook and Instagram to Offer Subscription for No Ads in Europe,” October 30, 2023, <https://about.fb.com/news/2023/10/facebook-and-instagram-to-offer-subscription-for-no-ads-in-europe>.

103 The EU’s e-Privacy Directive, which is responsible for the deluge of cookie banners, is one exception. The enactment of the GDPR helped nudge the consent dialog for cookie banners toward a more simplified “reject all” choice, should individuals wish to use it, rather than just a notification that cookies were being collected, which does not provide the user with any actual choices. It remains to be seen whether the renegotiation of the e-Privacy Directive will incorporate a method like Global Privacy Control to make cookie acceptance or rejection more streamlined, and whether an infrastructure can be adopted to preserve one’s choices across multiple modalities, not just browsers (and potentially, not just cookies).

104 Colorado presently has the most detailed regulations on the books regarding automated decision-making. For a current listing of U.S. state privacy laws, see Folks, “US State Privacy Legislation Tracker.”

105 Organisation for Economic Cooperation and Development, “10 Policy Issues In Data Protection and Privacy: Concepts and Perspectives,” *Proceedings of the OECD Seminar 24th-26th June 1974*, June 1974, <https://glgonzalezfuster.files.wordpress.com/2022/06/policy-issues-in-data-protection-and-privacy-concepts-and-perspectives-proceeding-of-th-eoecd-seminar-24th-26th-june-1974.pdf>, 73.

106 Information Commissioner’s Office, “What is automated individual decision-making and profiling?” These documents in turn reference the Article 29 Working Party’s explainer on profiling and ADM, see European Commission, Guidelines on Automated individual decision-making and Profiling for the purposes of Regulation 2016/679 (wp251rev.01), Article 29 Data Protection Working Party, August 22, 2018, <https://ec.europa.eu/newsroom/article29/items/612053/en>.

107 European Parliament and European Council, GDPR, Article 22, Sec.1. Profiling and automated processing is defined in Recital 71: “The data subject should have the right not to be subject to a decision, which may include a measure, evaluating personal aspects relating to him or her which is based solely on automated processing and which produces legal effects concerning him or her or similarly significantly affects him or her, such as automatic refusal of an online credit application or e-recruiting practices without any human intervention. Such processing includes ‘profiling’ that consists of any form of automated processing of personal data evaluating the personal aspects relating to a natural person, in particular to analyse or predict aspects concerning the data subject’s performance at work, economic situation, health, personal preferences or interests, reliability or behaviour, location or movements, where it produces legal effects concerning him or her or similarly significantly affects him or her.”

108 European Parliament and European Council, GDPR, Article 4 (4).

109 Angelina Wang, Sayash Kapoor, Solon Barocas, and Arvind Narayanan, “Against Predictive Optimization: On the Legitimacy of Decision-Making Algorithms that Optimize Predictive Accuracy,” *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, June 2023, <https://doi.org/10.1145/3593013.3594030>, 626.

110 While many suggest George Orwell’s 1984 as an example of such a society, Terry Gilliam’s 1985 film “Brazil” is another excellent example: “Brazil (1985 film),” *Wikipedia*, February 1, 2024, [https://en.wikipedia.org/wiki/Brazil_\(1985_film\)](https://en.wikipedia.org/wiki/Brazil_(1985_film)).

111 The CCPA was initially introduced as a ballot initiative directly to the voters, not through the legislative process. See Nicholas Confessore, “The Unlikely Activists that Took On Silicon Valley—And Won,” *The New York Times Magazine*, August 14, 2018, <https://www.nytimes.com/2018/08/14/magazine/facebook-google-privacy-data.html>.

112 Natasha Singer, “Group Behind California Privacy Law Aims to Strengthen It,” *The New York Times*, September 24, 2019, <https://www.nytimes.com/2019/09/24/technology/group-behind-california-privacy-law-aims-to-strengthen-it.html>.

113 California Privacy Protection Agency, Addendum to Draft Automated Decisionmaking Technology Regulations: New Rules Subcommittee Format Proposal, December 2023, https://cppa.ca.gov/meetings/materials/20231208_item2_addendum.pdf.

114 California Privacy Protection Agency, “A New Landmark for Consumer Control Over Their Personal Information.”

115 Colorado Department of Law, Colorado Privacy Act Rules, 4 CCR 904-3, March 15, 2023, <https://coag.gov/app/uploads/2023/03/FINAL-CLEAN-2023.03.15-Official-CPA-Rules.pdf>.

116 Rabel J. Burdge, “A brief history and major trends in the field of impact assessment,” *Impact Assessment* 9(4), 1991, <https://www.tandfonline.com/doi/abs/10.1080/07349165.1991.9726070>, 93-104.

117 107th U.S. Congress, e-Government Act of 2002, H.R. 2458, December 17, 2002, <https://www.congress.gov/107/plaws/publ347/PLAW-107publ347.pdf>.

118 European Commission, “When is a Data Protection Impact Assessment (DPIA) required?” https://commission.europa.eu/law/law-topic/data-protection/reform/rules-business-and-organisations/obligations/when-data-protection-impact-assessment-dpia-required_en.

119 The California Privacy Protection Agency is in its rulemaking phase for these regulations as of early 2024. The regulations may not be finalized until July 2024.

120 See Theodore P. Augustinos, “Emerging Requirements for Data Protection Impact Assessments,” *Locke Lord*, Spring 2022, <https://www.lockelord.com/newsandevents/publications/2022/05/emerging-requirements>; California Privacy Protection Agency, New Rules Subcommittee Revised Draft Risk Assessment Regulations, December 2023, https://cppa.ca.gov/meetings/materials/20231208_item2_draft_clean.pdf.

121 For a general example, see 118th U.S. Congress, Algorithmic Accountability Act of 2023. For an overview of algorithmic impact assessments, see Moss et al., “Assembling Accountability.”

122 Anne Josephine Flanagan, Jen King, and Sheila Warren, “Redesigning Data Privacy: Reimagining Notice & Consent for human technology interaction,” *World Economic Forum*, July 30, 2020, <https://www.weforum.org/publications/redesigning-data-privacy-reimagining-notice-consent-for-humantechnology-interaction/>.

123 Marci Meingast, Jennifer King, and Deirdre K. Mulligan, “Embedded RFID and Everyday Things: A Case Study of the Security and Privacy Risks of the US e-Passport,” *2007 IEEE International Conference on RFID*, April 2007, <https://ieeexplore.ieee.org/abstract/document/4143504>, 7-14.

124 Karen Weise, Cade Metz, Nico Grant, and Mike Isaac, “Inside the A.I. Arms Race That Changed Silicon Valley Forever,” *The New York Times*, December 5, 2023, <https://www.nytimes.com/2023/12/05/technology/ai-chatgpt-google-meta.html>.

125 Mehtab Khan and Alex Hanna, “The Subjects and Stages of AI Dataset Development: A Framework for Dataset Accountability,” *Ohio State Technology Law Journal* 19, 2023, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4217148.

126 California Privacy Protection Agency, New Rules.

127 California Privacy Protection Agency, New Rules.

128 Salomé Viljoen, “A Relational Theory of Data Governance,” *Yale Law Journal* 131(2), November 2021, <https://www.yalelawjournal.org/feature/a-relational-theory-of-data-governance>.

129 See European Digital Rights, “Civil society calls for AI red lines in the European Union’s Artificial Intelligence proposal,” January 12, 2021, <https://edri.org/our-work/civil-society-call-for-ai-red-lines-in-the-european-unions-artificial-intelligence-proposal/>; Fanny Hidvegi, Daniel Leufer, and Estelle Massé, “The EU should regulate AI on the basis of rights, not risks,” *Access Now*, February 17, 2021, <https://www.accessnow.org/eu-regulation-ai-risk-based-approach/>.

130 Amba Kak and Sarah Myers West, “AI Now 2023 Landscape: Confronting Tech Power,” *AI Now*

131 Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell, “On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?,” *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, March 2021, <https://dl.acm.org/doi/abs/10.1145/3442188.3445922>, 610–623.

132 Dahlia Peterson and Samantha Hoffman, “Geopolitical implications of AI and digital surveillance adoption,” *Brookings*, June 2022, <https://www.brookings.edu/articles/geopolitical-implications-of-ai-and-digital-surveillance-adoption/>.

133 Karen Hao, “AI is sending people to jail—and getting it wrong,” *MIT Technology Review*, January 21, 2019, <https://www.technologyreview.com/2019/01/21/137783/algorithms-criminal-justice-ai/>.

134 Ivey Dyson, “How AI Threatens Civil Rights and Economic Opportunities,” *Brennan Center for Justice*, November 16, 2023, <https://www.brennancenter.org/our-work/analysis-opinion/how-ai-threatens-civil-rights-and-economic-opportunities>.

135 The 2023 Biden Executive Order on AI calls out this tension by mandating that privacy and security standards be incorporated into the software development life cycle, and that the Office of Management and Budget must identify the commercial data agencies may have procured that contains personal identifiable information and provide guidance as to how to mitigate privacy risks in its use. See White House, Executive Order on the Safe, Secure, and Trustworthy Development and use of Artificial Intelligence, October 30, 2023, <https://www.whitehouse.gov/briefing-room/presidential-actions/2023/10/30/executive-order-on-the-safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence/>. For example, §8(b)(i)(D) (“incorporation of safety, privacy, and security standards into the software-development lifecycle for protection of personally identifiable information”) and §9(a)(i) (“evaluate and take steps to identify commercially available information (CAI) procured by agencies, particularly CAI that contains personally identifiable information and including CAI procured from data brokers and CAI procured and processed indirectly through vendors, in appropriate agency inventory and reporting processes (other than when it is used for the purposes of national security); (ii) evaluate, in consultation with the Federal Privacy Council and the Interagency Council on Statistical Policy, agency standards and procedures associated with the collection, processing, maintenance, use, sharing, dissemination, and disposition of CAI that contains personally identifiable information (other than when it is used for the purposes of national security) to inform potential guidance to agencies on ways to mitigate privacy and confidentiality risks from agencies’ activities related to CAI”).

136 “It’s the economy, stupid,” *Wikipedia*, November 16, 2023, https://en.wikipedia.org/wiki/It%27s_the_economy_stupid.

137 European Parliament and European Council, Directive 2009/136/EC of the European Parliament and of the Council [ePrivacy Directive], November 25, 2009, https://edps.europa.eu/data-protection/our-work/subjects/eprivacy-directive_en.

138 Federal Trade Commission, “Privacy Online: Fair Information Practices in the Electronic Marketplace - A Report to Congress,” May 2000, <https://www.ftc.gov/sites/default/files/documents/reports/privacy-online-fair-information-practices-electronic-marketplace-federal-trade-commission-report/privacy2000text.pdf>.

139 King et al., “Online privacy notices don’t work.”

140 For example, see Federal Trade Commission, “Rite Aid Banned from Using AI Facial Recognition After FTC Says Retailer Deployed Technology without Reasonable Safeguards,” December 19, 2023, <https://www.ftc.gov/news-events/news/press-releases/2023/12/rite-aid-banned-using-ai-facial-recognition-after-ftc-says-retailer-deployed-technology-without>.

141 Federal Trade Commission, Trade Regulation Rule on Commercial Surveillance and Data Security, August 22, 2022, <https://www.federalregister.gov/documents/2022/08/22/2022-17752/trade-regulation-rule-on-commercial-surveillance-and-data-security>.

142 European Parliament and European Commission, Proposal for a Regulation of the European Parliament and of the Council concerning the respect for private life and the protection of personal data in electronic communications and repealing Directive 2002/58/EC [Regulation on Privacy and Electronic Communications], January 10, 2017, <https://www.europarl.europa.eu/legislative-train/theme-a-europe-fit-for-the-digital-age/file-jd-e-privacy-reform>.

143 European Commission, “Can data received from a third party be used for marketing?” https://commission.europa.eu/law/law-topic/data-protection/reform/rules-business-and-organisations/legal-grounds-processing-data/can-data-received-third-party-be-used-marketing_en.

144 Federal Trade Commission, “FTC Proposes Strengthening Children’s Privacy Rule to Further Limit Companies’ Ability to Monetize Children’s Data,” December 20, 2023, <https://www.ftc.gov/news-events/news/press-releases/2023/12/ftc-proposes-strengthening-childrens-privacy-rule-further-limit-companies-ability-monetize-childrens>.

145 Jennifer King, Daniel Ho, Arushi Gupta, Victor Wu et al., “The Privacy-Bias Tradeoff: Data Minimization and Racial Disparity Assessments in U.S. Government,” *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, June 2023, <https://dl.acm.org/doi/abs/10.1145/3593013.3594015>, 492–505.

146 Jennifer King, “A Bill Designed to Protect Kids Could Change the Internet for the Better,” *Tech Policy Press*, September 15, 2022, <https://www.techpolicy.press/a-bill-designed-to-protect-kids-could-change-the-internet-for-the-better/>.

147 Thomson Reuters Westlaw, “Barclays Official California Code of Regulations,” https://govt.westlaw.com/calregs/Document/IA39673F0D44D11ED8CB7A9089AB7E22E?h%2520Result&list=REGULATION_PUBLICVIEW&rank=8&t_T2=7004&t_S1=CA+ADC+s, 11 CA ADC § 7004(b-c).

148 Apple Support, “If an app asks to track your activity,” <https://support.apple.com/en-us/HT212025>.

149 Singular, “Limit Ad Tracking,” <https://www.singular.net/glossary/limit-ad-tracking/>.

- 150 Tiahn Wetzler, “ATT two years on: Opt-in rates keep climbing,” *Adjust*, July 5, 2023, <https://www.adjust.com/blog/app-tracking-transparency-opt-in-rates/>.
- 151 King, “Online privacy notices don’t work.”
- 152 Nick Doty, “It’s Time to Standardize the Global Privacy Control,” *Center for Democracy and Technology*, December 13, 2023, <https://cdt.org/insights/its-time-to-standardize-the-global-privacy-control/>.
- 153 See Colorado Department of Law, Colorado Privacy Act Rules, Rule 5.06(A)(1) (“The Universal Opt-Out Mechanism may communicate a Consumer’s opt-out choice by sending an opt-out signal.”); California Consumer Privacy Act, §1798.185(a)(19)(A) (“Issuing regulations to define the requirements and technical specifications for an opt-out preference signal sent by a platform, technology, or mechanism, to indicate a consumer’s intent to opt out of the sale or sharing of the consumer’s personal information and to limit the use or disclosure of the consumer’s sensitive personal information.”)
- 154 Global Privacy Control, <https://globalprivacycontrol.org/>.
- 155 Katherine Chaves and Jamie Nafziger, “Universal Opt-Out/Global Privacy Control: Preparing for the New Online World,” *JD Supra*, December 7, 2022, <https://www.jdsupra.com/legalnews/universal-opt-out-global-privacy-7739024/>.
- 156 This is ignoring the material side effects of data processing and storage on the environment, through energy use, materials for computer hardware, etc. Data collection and processing has material impacts, even if difficult to measure.
- 157 Article 10 of the AI Act (not final at publication time) is a data governance provision focused on training, validation, and testing data sets for high risk systems and may address some quality issues, but not specifically the privacy based concerns we raise here. See Council of the European Union, Artificial Intelligence Act.
- 158 Milad Nasr, Nicholas Carlini, Jon Hayase, Matthew Jagielski et al., “Extracting Training Data from ChatGPT,” November 28, 2023, <https://not-just-memorization.github.io/extracting-training-data-from-chatgpt.html>.
- 159 Drew Breunig, “Considering AI Privacy Scenarios,” *dbreunig.com*, May 15, 2023, <https://www.dbreunig.com/2023/05/15/ai-privacy-scenarios.html>.
- 160 Federal Trade Commission, “AI Companies: Uphold Your Privacy and Confidentiality Commitments,” January 9, 2024, <https://www.ftc.gov/policy/advocacy-research/tech-at-ftc/2024/01/ai-companies-uphold-your-privacy-confidentiality-commitments>.
- 161 Digital Curation Centre et al., “The Role of Data In AI.”
- 162 Nithya Sambasivan, Shivani Kapania, Hannah Highfill, Diana Akrong et al., “Everyone wants to do the model work, not the data work,” *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, May 2021, <https://dl.acm.org/doi/abs/10.1145/3411764.3445518>.
- 163 Alon Halevy, Peter Norvig, and Fernando Pereira, “The Unreasonable Effectiveness of Data,” *IEEE Intelligent Systems* 24(2), March 2009, <https://ieeexplore.ieee.org/document/4804817>, 8-12.
- 164 Craig S. Smith, “Hallucinations Could Blunt ChatGPT’s Success: OpenAI says the problem’s solvable, Yann LeCun says we’ll see,” *IEEE Spectrum*, March 13, 2023, <https://spectrum.ieee.org/ai-hallucination>.
- 165 Ronan Eldan and Yuanzhi Li, “TinyStories: How Small Can Language Models Be and Still Speak Coherent English?” *arXiv preprint*, May 12, 2023, <https://arxiv.org/abs/2305.07759>. The dataset is available on Hugging Face. See Hugging Face, TinyStories, <https://huggingface.co/datasets/ronaneldan/TinyStories>.
- 166 Xinyang Geng, Arnav Gudibande, Hao Liu, Eric Wallace, et al., “Koala: A Dialogue Model for Academic Research,” *Berkeley Artificial Intelligence Research*, April 3, 2023, <https://bair.berkeley.edu/blog/2023/04/03/koala/>.
- 167 Sambasivan et al., “Everyone wants to do the model work, not the data work.”
- 168 Abeba Birhane, Vinay Uday Prabhu, and Emmanuel Kahembwe, “Multimodal datasets: misogyny, pornography, and malignant stereotypes,” *arXiv preprint*, October 5, 2021, <https://arxiv.org/abs/2110.01963>.
- 169 Khan et al., “The Subjects and Stages of AI Dataset Development.”
- 170 Khan et al., “The Subjects and Stages of AI Dataset Development.”
- 171 Chloe Xiang, “AI Is Probably Using Your Images and It’s Not Easy to Opt Out,” *Vice*, September 26, 2022, <https://www.vice.com/en/article/3ad58k/ai-is-probably-using-your-images-and-its-not-easy-to-opt-out>.
- 172 See Federal Trade Commission, “Rite Aid Banned from Using AI Facial Recognition After FTC Says Retailer Deployed Technology without Reasonable Safeguards,” December 19, 2023, <https://www.ftc.gov/news-events/news/press-releases/2023/12/rite-aid-banned-using-ai-facial-recognition-after-ftc-says-retailer-deployed-technology-without>; Federal Trade Commission, “FTC and DOJ Charge Amazon with Violating Children’s Privacy Law by Keeping Kids’ Alexa Voice Recordings Forever and Undermining Parents’ Deletion Requests,” May 31, 2023, <https://www.ftc.gov/news-events/news/press-releases/2023/05/ftc-doj-charge-amazon-violating-childrens-privacy-law-keeping-kids-alexa-voice-recordings-forever>; Federal Trade Commission, “FTC Takes Action Against Company Formerly Known as Weight Watchers for Illegally Collecting Kids’ Sensitive Health Data,” March 4, 2022, <https://www.ftc.gov/news-events/news/press-releases/2022/03/ftc-takes-action-against-company-formerly-known-weight-watchers-illegally-collecting-kids-sensitive>; Federal Trade Commission, “FTC Finalizes Settlement with Photo App Developer Related to Misuse of Facial Recognition Technology,” May 7, 2021, <https://www.ftc.gov/news-events/news/press-releases/2021/05/ftc-finalizes-settlement-photo-app-developer-related-misuse-facial-recognition-technology>.
- 173 Zachary Izzo, Mary Anne Smart, Kamalika Chaudhuri, and James Zou, “Approximate Data Deletion from Machine Learning Models,” *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics*, March 2021, <https://proceedings.mlr.press/v130/izzo21a.html>, 2008-2016.
- 174 Ben Hutchinson, Andrew Smart, Alex Hanna, Emily Denton et al., “Towards Accountability for Machine Learning Datasets: Practices from Software Engineering and Infrastructure,” *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, March 2021, <https://doi.org/10.1145/3442188.3445918>, 560–575.
- 175 See generally Kenny Peng, Arunesh Mathur, and Arvind Narayanan, “Mitigating Dataset Harms Requires Stewardship: Lessons from 1000 Papers,” *arXiv preprint*, August 6, 2021, <https://arxiv.org/abs/2108.02922>; Amandalynne Paullada, Inioluwa Deborah Raji, Emily M. Bender, Emily Denton, et al., “Data and its (dis) contents: A survey of dataset development and use in machine learning research,” *Patterns* 2(11), November 12, 2021, <https://doi.org/10.1016/j.patter.2021.100336>.

- 176 Weixin Liang, Girmaw Abebe Tadesse, Daniel Ho, L. Fei-Fei et al., “Advances, challenges and opportunities in creating data for trustworthy AI,” *Nature Machine Intelligence* 4, August 17, 2022, <https://doi.org/10.1038/s42256-022-00516-1>, 669–677.
- 177 See Alex LaCasse, “Proposed data provenance standards aim to enhance trustworthiness of AI training data,” *International Association of Privacy Professionals*, January 17, 2024, <https://iapp.org/news/a/leading-corporations-proposed-data-provenance-standards-aims-to-enhance-quality-of-ai-training-data/>; Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan et al., “Datasheets for datasets,” *Communications of the ACM* 64(12), <https://dl.acm.org/doi/fullHtml/10.1145/3458723>, 86–92; Emily M. Bender and Batya Friedman, “Data Statements for Natural Language Processing: Toward Mitigating System Bias and Enabling Better Science,” *Transactions of the Association for Computational Linguistics* 6, December 1, 2018, https://direct.mit.edu/tacl/article/doi/10.1162/tacl_a_00041/43452/Data-Statements-for-Natural-Language-Processing, 587–604; The Data Nutrition Project, <https://datanutrition.org/>; Data Provenance Explorer, <https://www.dataprovenance.org/>.
- 178 California Privacy Protection Agency, New Rules.
- 179 Daniel E. Ho, Jennifer King, Russell C. Wald, and Christopher Wan, “Building a National AI Research Resource,” *Stanford Institute for Human-Centered Artificial Intelligence*, October 2021, <https://hai.stanford.edu/white-paper-building-national-ai-research-resource>.
- 180 Ho et al., “Building a National AI Research Resource.”
- 181 Zuboff, *The Age of Surveillance Capitalism*.
- 182 Hannah Klein, “Andrew Yang Wants You to Own and Sell Your Data,” *Slate*, June 23, 2020, <https://slate.com/technology/2020/06/yang-launches-data-dividend-project.html>.
- 183 Cameron F. Kerry and John B. Morris, “Why data ownership is the wrong approach to protecting privacy,” *Brookings*, June 26, 2019, <https://www.brookings.edu/articles/why-data-ownership-is-the-wrong-approach-to-protecting-privacy/>.
- 184 Pamela Samuelson, “Privacy As Intellectual Property?” *Stanford Law Review* 52, 1999, <https://heinonline.org/HOL/P?h=hein.journals/stflr52&i=1145>, 1125.
- 185 World Economic Forum, “Advancing Digital Agency: The Power of Data Intermediaries,” February 15, 2022, <https://www.weforum.org/publications/advancing-digital-agency-the-power-of-data-intermediaries/>. See also M. Micheli, E. Farrell, B. Carballa Smichowski, M. Posada Sanchez et al., “Mapping the landscape of data intermediaries: Emerging models for more inclusive data governance,” Publications Office of the European Union, August 24, 2023, <https://publications.jrc.ec.europa.eu/repository/handle/JRC133988>.
- 186 Solid, <https://solidproject.org/>.
- 187 Steve Lohr, “He Created the Web. Now He’s Out to Remake the Digital World,” *The New York Times*, January 10, 2021, <https://www.nytimes.com/2021/01/10/technology/tim-berners-lee-privacy-internet.html>.
- 188 Kak et al., “AI Now 2023 Landscape: Confronting Tech Power.”
- 189 Evan Selinger and Woodrow Hartzog, “Stop Saying Privacy is Dead,” *Medium*, October 18, 2018, <https://medium.com/@evanselinger/stop-saying-privacy-is-dead-513dda573071>.