RE: Response to the National Telecommunications and Information Administration Regarding Privacy, Equity, and Civil Rights Request for Comment (Docket No. 230103-0001)

To Whom It May Concern:

Palantir Technologies ("Palantir") is a US-based software company that builds platforms to enable public, private, and non-governmental organizations to integrate, analyze, and collaborate on their data in a secure and privacy-protective way. We are proud to make software that enables the institutions that serve our societies to use their data responsibly and effectively.

Palantir was founded in 2003 on the conviction that it is essential to preserve fundamental principles of privacy and civil liberties while using data. It is for this reason that Palantir established one of the world's first Privacy & Civil Liberties Engineering ("PCL") teams more than a decade ago, specifically to focus on the development of privacy-protective technologies and to foster a culture of responsibility around their development and use. Our PCL team serves as a specialist vanguard to advise and lead on these subjects, and we especially welcome efforts by the National Telecommunications and Information Administration ("NTIA") in highlighting several important issues at the intersection of privacy, equity, and civil rights.

Our response to this Request for Comment ("RFC") is based on insights gathered over nearly 20 years of experience building technology to uphold and enforce privacy principles across industries and jurisdictions. We have specifically focused our response on our areas of operational expertise, both answering select questions from the RFC and contributing our perspective on how organizations can uphold privacy and equity in practice. The Executive Summary provides a brief overview of our response and the specific approaches we recommend.

We are thankful to the NTIA for the opportunity to contribute to this valuable public discourse on privacy, equity, and civil rights. We would welcome any request for clarification, and we look forward to the NTIA's final report on these critical issues.

Sincerely,

Courtney Bowman
Global Director of Privacy and Civil Liberties Engineering, Palantir Technologies

Arnav Jagasia
Privacy and Civil Liberties Engineering Lead, Palantir Technologies

Helena Vrabec
Data Protection and Privacy Lead, Palantir Technologies

Executive Summary

Our response to this Privacy, Equity, and Civil Rights Request for Comment ("RFC") presents our perspective on how organizations can facilitate the responsible use of technology as well as several tactical technical and organizational recommendations for the National Telecommunications and Information Administration (NTIA) to consider.

To briefly summarize the response that follows, we encourage the NTIA:

- To recognize the prevailing (and sometimes, polarizing) positions present in the discourse around the appropriate use of technology, especially as it concerns marginalized communities. Pp. 5-10.
- To adopt a centrist position on how to make imperfect data work for imperfect environments by building resilience and fault tolerance into data use. Pp. 10-12.
- To consider two case studies about employee data and sensitive health information, in which we outline suggestions and best practices rooted in our experience with these domains. Pp. 12-17. Based on these case studies, we believe the following approaches are crucial to any strategy that encourages organizations to uphold privacy, equity, and civil rights:
    - Foundational data protection technologies
    - Purpose specification for sensitive actions
    - Scheduled deletion of sensitive data by default
    - Data minimization
    - Regular data quality assessment
    - Use limitation restrictions
    - Codes of conduct
    - Organizational data governance bodies

Our response below touches on several questions specifically asked in the RFC. We specify via footnote each response to a particular question from the RFC.

Table of Contents

Introduction

We recognize that lax data privacy regulations as well as the practices of certain commercial organizations can lead to significant social harms. The introduction to the RFC as well as the prior listening sessions highlighted several examples of commercial harms, primarily driven by the use of technology. These harms can come from a variety of technological solutions. For example, the use of AI/ML models that are poorly designed or deployed outside their area of intended application can lead to bias, equity, and fairness concerns. This can further lead to particularly adverse impacts for marginalized or underrepresented communities. Additionally, companies that engage in surreptitious collection and re-sale of potentially sensitive consumer data can violate individuals' safety and privacy. As more organizations use technology to conduct their business practices, these harms and abuses will continue to grow in scale and severity.

As a software technology provider, we recognize that technology can be misused or abused. It is for this reason that we have focused on privacy and security as foundational requirements for our software offerings. By building capabilities for granular access controls, lineage-aware deletion, robust AI/ML model governance, to name a few, into our software platforms, we have earned the trust of organizations to provide software that can be used for their most critical challenges.

In our experience building software through this privacy- and individual rights-protective lens, we have seen two common positions that frame the discourse around data harms and their potential mitigations.

1. Advocates of technology and data systems posit that using data-driven solutions is the most effective way to solve an organization's challenges by virtue of its data-based foundation. We agree that using data to solve operational challenges can be effective to deliver solutions grounded in an empirical reality, but this position has two shortcomings. First, proponents of this viewpoint can see technology as the solution to every problem. In our experience, however, it is imperative to consider whether technology – especially, the use of automated-decision making systems – should even be used in the first place to solve a given problem. As we have written in our comments to the National Security Commission for AI, technology should be used to solve a specific, well-defined objective.[1] Second, a data-driven solution will necessarily reflect the reality of the data it is provided and can therefore exacerbate existing biases. Such systems should not be used without first seriously considering such risks and potential mitigations.
2. Critics of the use of technology, especially for challenges in domains like law enforcement and healthcare, argue that data-driven solutions should largely be avoided due to concerns of bias in data and algorithms. We also agree that the crux of this argument is incontestably true: algorithms, data, and information more broadly will always reflect some inherently limited perspective of the world. What differentiates our

[1] *See* Eric Heller & Akash Jain, *Our Recommendations to the NSCAI*, PALANTIR BLOG (2021), https://blog.palantir.com/palantirs-recommendations-to-the-nscai-f5d7d5dad344 ("AI works best when it is considered holistically and is grounded in advancing specific mission objectives, rather than being introduced as techno-solutionism in search of a problem."); *See also Final Report*, NAT'L. SEC. COMM'N. A.I. (2021), https://www.nscai.gov/2021-final-report/.

position, however, is that we contend that some of those limitations will be acceptable in certain contexts. As we have written in our comments to NIST's Proposal for Identifying and Managing Bias in Artificial Intelligence, bias must be evaluated in the context of a particular deployment of technology.[2] Insights about bias must consider the context of application to determine whether the forms of bias are *in situ* desirable or undesirable. For example, applications of artificial intelligence in healthcare for legitimate diagnostic purposes might require access to sensitive data that could encode systemic biases. In some constrained contexts, this may be deemed acceptable, but its application to other use cases might not. We contend that any use of technology with sensitive data should be rigorously and thoroughly scrutinized, but we should not foreclose data-driven technologies solely because of the limitations inherent in all collected datasets or algorithms.

While certain aspects of both arguments are true, we find that this dichotomy tends to channel the discourse around data-driven techniques applied to some of society's most pressing challenges in an unproductive direction.[3]

Instead, we have long advocated for a more centrist, nuanced position. We contend that technologists, social scientists, and policymakers can best advance the responsible use of technology, while safeguarding against harms especially to marginalized populations, by focusing on building fault tolerance and resilience into technology systems. This framing refocuses the discussion on how to make imperfect data work for imperfect environments, rather than either exalting or completely foreclosing the use of technology.[4]

Echoing our statements in our response to the Federal Trade Commission's Advanced Notice of Proposed Rulemaking on Commercial Surveillance and Data Security, we believe that the thoughtful establishment of rules and standards that guide commercial organizations to optimize their technology and data practices for *both* business outcomes *and* privacy can free us from the zero-sum notion that the use of technology must necessarily come at the expense of privacy, civil rights, and equity. [5]

I. A Tale of Two Cultures

Let us define two ideological camps whose perspectives tend to dominate the discourse around data and technology applications in consequential consumer-facing settings.[6] On one side, "data

[2] Anthony Bak, Courtney Bowman, & Megha Arora, *Palantir Comments on NIST SP 1270*, (2021), available at https://www.nist.gov/system/files/documents/2021/09/15/20210910_Palantir%20Reponse%20to%20SP%201270.pdf

[3] *See infra* Part I. A Tale of Two Cultures, for a fuller exposition of these arguments and the dichotomy they present.

[4] *See infra* Part II. Reframing the Problem: A Pragmatic Synthesis of Data Evangelism and Data Detraction.

[5] Courtney Bowman, Arnav Jagasia, & Helena Vrabec, *Comments to the Federal Trade Commission Regarding Advanced Notice of Proposed Rulemaking on Commercial Surveillance and Data Security*, 1-2 (2020), available at https://www.regulations.gov/comment/FTC-2022-0053-0702 [hereinafter Palantir Response to FTC ANPR].

[6] In describing these two camps, we acknowledge the risk of being somewhat reductive. But, much like C. P. Snow's seminal lecture, *The Two Cultures,* we find that there is a "gulf of mutual incomprehension" between these two ideologies prevalent in many discussions around the use of technology. *See* C. P. Snow, *The Two Cultures*, 4

evangelists" are typically technologists and entrepreneurs who are motivated by an optimism that data and data-driven technologies tend to offer the best (in some cases, the only) solution to problems that consumers and consumer-facing organizations face. "Data evangelists" see traditional human-driven approaches as mired in personality flaws, subjectivity, and all the messiness of contingent personal lives. On the other side, "data detractors" are typically (but not exclusively) social scientists, advocates, and other non-technologists who view "data" and "data-driven technologies" as code for objective frameworks for viewing the world. These frameworks, they assert, have been regularly discredited as naive, simplistic, or pernicious in their tendency to uphold cultural and social worldviews that are unfair and favor the privileged over historically disadvantaged groups. They see data and technologies not as solutions to human problems, but as just the latest way that human frailties like bias become further enshrined in cultural practices.

Both camps represent views with legitimate concerns about the world, and in the best circumstances, laudable motivations to address those concerns. Both camps also represent forms of myopic thinking that lead to miscommunication, stalemate, and often direct conflict. Both sides would benefit by bracketing their respective ideologies and examining how they each have failed in the past to carry out their objectives.

A. "Data Evangelism"

"Data evangelists" need to reckon with the reality that data-driven techniques are not unequivocally salutary. The consumer internet and the proliferation of smart devices that serve as its end-user interfaces introduce an entire world of ready-to-hand information, communications, entertainment, consumption, and much more. For almost every consumer itch, there's an app that "solves" it. The frictional points of the analog world (passengers seeking drivers, tenants seeking renters, bachelors/bachelorettes seeking companions, etc.) are viewed as points of virtual intervention, interactions to be simplified through technological mediation. The evangelists speak in terms of "data democratization," bringing information and capabilities to the masses through the easy interaction with devices we all have come to treat as inseparable digital companions.

However, the more these problems are "solved," the more they can generate unexpected externalities and more problems. These problems accumulate along a spectrum of harms, including unreasonable flows of commercial data to government agencies, commercial surveillance practices that monetize and manipulate behavior, automated decision-making systems that unwittingly introduce or exacerbate unlawful discrimination. Take, for example, communications services and technologies that increasingly help to mediate the most intimate aspects of our personal and social lives. These tools, in the course of carrying out their explicitly intended functions, also often produce massive troves of digital exhaust that, placed in the wrong hands, could be used to reconstruct in vivid detail vast swaths of the private sphere. The same information that most citizens would vehemently guard against government intrusion, is

---

(Canto ed. 1993). By understanding the merits present in both arguments, we advocate for a more centrist position on the responsible use of technology. *See infra* Part II. Reframing the Problem: A Pragmatic Synthesis of Data Evangelism and Data Detraction. This, in turn, informs our perspective and recommendations on how to use technology without sacrificing privacy, equity, or civil rights.

willingly handed over to private sector service providers every hour of every day. Consumers – the extent to which they are truly aware of the staggering breadth of such data accumulation – may be willing to accept this as a necessary tax on the use of digital technologies. But most people almost certainly would not willingly consent to the unfettered sharing of their email, text messages, and other personal correspondences with any government agency, including law enforcement. Indeed, many of our most sacred constitutional protections – civil liberties enshrined in the Bill of Rights – exist to guard against unlawful state intrusion carried out absent legal authority and due process constraints. And yet, commercial service providers are less constrained, and entire marketplaces have emerged for the accumulation, packaging, and reselling of digital communications information to private and public sector buyers alike.

There is good reason to be alarmed by the risk of government agencies circumventing due process restrictions on access to citizens' data by using private sector data brokers and resellers as an end-run.[7] Privacy interests of citizens in the consumer-facing communications industry, in online ads, and other industries that trade in personal information should not end at the division of the public and private sectors. Government agencies, we believe, should not be permitted to circumvent due process considerations in seeking data from consumer facing applications and communications services that would otherwise require a warrant, subpoena, or other legal process to acquire directly. We believe it is in the interests of US citizens that due process considerations are factored into all requests for use of these classes of personal information. Where law enforcement, intelligence, and defense agencies have legitimate need to use sensitive and personal communications data to advance their critical missions, we believe that legal standards directed at upholding Fourth Amendment and other constitutional guarantees of US persons should be consistently applied by government agencies in their methods of access to this data.

There are other, more subtle, harms to consumers arising from the "solutionism" advanced by data evangelists.[8] Commercial "surveillance capitalism" does not simply seek methods of benignly monetizing our interactions with digital applications; it actually shapes our behaviors, beliefs, and interactions with the world.[9] From social media services that are optimized to increase and extend user engagements (and thereby serve up more ads and monetizable content) to online shopping portals that algorithmically propose to us what our consumer needs are via recommendations, digital technologies marketed as tools for making our lives easier may in fact be manipulating our very sense of what kind of lives we should be leading. Discussions around the impact of social media platform content promotion and moderation on political polarization is just the latest, most salient instantiation of the ways that these technologies – once lauded as epochal advances in democratizing information sharing and community engagement – have since

---

[7] *See generally,* proposed Fourth Amendment Is Not For Sale Act, S. 1265, 117th Cong. (2021), https://www.congress.gov/bill/117th-congress/senate-bill/1265; Press Release, *Wyden, Paul and Bipartisan Members of Congress Introduce The Fourth Amendment Is Not For Sale Act* (Apr. 21, 2021) https://www.wyden.senate.gov/news/press-releases/wyden-paul-and-bipartisan-members-of-congress-introduce-the-fourth-amendment-is-not-for-sale-act- (highlighting the gaps in Constitutional protections against unreasonable government searches, especially when government agencies use private data).

[8] *See* Evgeny Morozov, *To Save Everything, Click Here: The Folly of Technological Solutionism* (2013).

[9] Shoshana Zuboff, *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power* (2019).

demonstrated that the echo chamber effects they enable may have led to more social division than unity or connectedness.[10] Since social media is particularly popular with younger generations including teenagers and even preschoolers, these constant waves of polarization and manipulation are becoming dangerously ingrained as part of their upbringing.[11]

Of late, one of the loudest criticisms leveled against the dogma of "data evangelism" is focused on the role of algorithms and automated decision-making systems as tools for embedding discrimination in the digital world. This argument takes two, often conflated, forms. In one form, it devolves to an earlier critique of data bias: data is biased, and therefore models and systems trained and built upon that data will also reflect that bias. The other form is that independent of data deficiencies or biases, the models themselves are often constructed with biased features or optimization parameters.[12] In this latter example, the remedy to addressing bias is not just a matter of gaining access to better, more representative data. Rather, it becomes a question of exercising myriad other best practices of sound data science, such as: selecting appropriate model types; interrogating the legitimacy and causal relevance of data features the model might rely upon; and determining the necessary and contextually appropriate metrics and tuning parameters that might be used, for example, to minimize disparities associated with sensitive group categories represented in the data in favor of some quantified notion of "fairness."

Suffice to say, there are several and notably severe ways in which algorithms built to purpose can have consequential impacts on peoples' lives: credit and lending determinations by financial institutions, candidate selection by recruiting agencies and employers, screening services by health providers, to name a few. Data evangelists who fail to acknowledge these demonstrable risks and the harms they have already effectuated on the lives of real people are living in a world of fantasy. They perpetuate a doctrine of technology salvation at the expense of not just their own credibility, but with real consequences to individuals and potentially growing backlash against even the most responsible and defensible variations of their current and future wares. Data evangelists would do well to temper their technological religiosity and spend a bit more time interacting with the people and communities their innovations ultimately exist to serve.

B. "Data Detraction"

On the other side of this digital divide, "data detractors" need to come to terms with the fact that data and data-driven technologies are not unequivocally harmful in all the ways that have become knee-jerk responses to every new technology. Data, information science, information technologies have engendered benefits ranging from improved healthcare, access to once-privileged information sources, timely and reliable communications with loved ones, efficiency and safety improvements in navigating through the physical world, and much more.

---

[10] *See generally* Jenny Darroch, THE HUFFINGTON POST (December 1, 2009, 5:12 AM EST), https://www.huffpost.com/entry/what-do-social-media-and_b_305583 (for an early example of the perceived salutary impacts of social media).

[11] Privacy, Equity, and Civil Rights Request for Comment, Docket No. 230103-0001, Question 2(a). Younger generations are particularly vulnerable to the harms of social media due to their regular exposure to social media.

[12] In fact, there are other forms of potential bias in automated decision-making systems, including but not limited to, biases in user interfaces, biases in programs of application, etc. We focus here on only two classes that we find are most commonly discussed by the critiques of "data evangelism."

To the point that "data detractors" often commence with, it should be recognized as a banality that data is biased. Data is simply information encoded digitally. But all information captured, whether acquired by means of human observation or machine interaction, is a reflection of a specific vantage point or set of vantage points on the world. The perspective is not the world itself, nor is there a single framing that is godlike and omniscient. In both digital and manual terms, we are always dealing with limits — limited perspectives, limited grounding assumptions, limited abilities to interrogate questions of meaning and intent. This is data and information bias in the most basic sense, and it should be acknowledged as a persistent, unavoidable reality of the world of information.

But this does not imply that all data and data-driven technologies should be discarded out of hand any more than individual testimonials should be dismissed as just the opinion of one or several person(s). Data may be flawed as much as human attestations may be flawed.[13] We assess the validity and veracity of information according to qualities like sourcing, method, recency, completeness, etc. These evaluations help us to determine how much trust we should place in our information. They also help us to understand the degree of bias inherent in our data — whether human or machine sourced — and then make appropriate decisions about whether to build on, modify, augment, or discard that information. Good data and data-driven technologies are subject to many of the same assessment principles as we would apply to the trustworthiness of people responsible for executing consequential tasks that impact our daily lives. In all cases, we ultimately should be aiming to appropriately contextualize data and its contingencies, understand the limitations and trade-offs implicit in those contingencies, and focus on well-bounded data applications that, to the greatest extent possible, provide meaningful corrections to known limitations and sensible off-ramps for potential or likely failings.

As much as "data detractors" may be correct in alerting us to the risks of over-reliance on data and technology solutionism, their critiques must also provide an honest reckoning with the practical realities of the world we live in. It is a world of increasing digitalization, interconnectedness, specialization, and information dependencies. Putting aside the moral valence of whether or not those trends are good, they are impossible to ignore or discard. Dealing with the world as it is (even to the end of changing it to be otherwise) necessitates an appreciation — or at minimum, a recognition — of the ways that data and data technologies can serve as powerful tools for helping us address the compounding complexities of modern life.

"Data detractors" must open their eyes to the reality that our data laden world must be reasonably managed and taken stock of, if it is to be brought in fuller alignment with certain values orientations. For example, institutions with a documented history of cultural or racial bias, such as redlining in mortgage lending, do not become reformed or improved by categorically rejecting data-driven or data-dependent approaches to lending as necessarily biased. Neither the problem nor the proposed solutions are that simple. Algorithmic approaches may indeed further entrench

---

[13] *See generally* Jake Silberg and James Manyika, McKinsey Global Institute, *Tackling bias in artificial intelligence (and in humans)*, (2019), available at https://www.mckinsey.com/featured-insights/artificial-intelligence/tackling-bias-in-artificial-intelligence-and-in-humans (for a discussion on the relationship between human bias and technology bias).

the historical bias. But they may also do otherwise: they could be wielded to correct or to serve as calibrated procedural refinements by, for example, helping to flag subtle but important checks on the mortgage assessment process.

Our key point is that "data detractors" (similar to their ideological opposites) need to take a step back from their criticisms and look more closely at what data and data-driven technologies are actually doing, how they can be tuned and contextually situated to serve as tools of responsible application or even as accountability mechanisms, and engage with their interlocutors with less reflexive skepticism and more open willingness to find a convergent path to dealing with a world awash in data and information technologies.

II. Reframing the Problem: A Pragmatic Synthesis of Data Evangelism and Data Detraction

We readily acknowledge that the diametric framing laid out above is something of a caricature sketch made for illustrative purposes. But, unfortunately, in our experience it is not that far removed from how many such discussions commence (and quickly end).

Our argument is that these opposing dogmas must be called out for what they are (even in their more subtle, cloaked forms) in order to provide a pathway for navigating the trap of cross-discipline partisan dispute, which too often leads to paralyzed inaction.

In lieu of a "data evangelist" versus "data detractor" fight to the death, *we propose an alternative framing that refocuses the discussion on how to make imperfect data work for imperfect environments, building fault tolerance and resilience into data use, rather than treating data bias (a reality of the world) as a damning critique on the one hand or holding data-driven technologies as objective, clever "hacks" to otherwise intractable life problems.*

Identifying and understanding data quality limitations are a necessary starting point of any data-dependent or data-driven enterprise. The necessary process of working with data in the digital age should always begin with an effort to characterize data and data system limits (including relevant bias considerations, but also efficacy, privacy, security, and a host of other considerations and constraints) as much as any anticipated problem-solving virtues.[14] This process is not unidimensional; it is fundamentally political and social, as much or more than it is technologically dependent. It involves an intersection of equities and interests that must be jointly addressed through deliberative and discursive means. The goal is not to arrive at a uniformly perfect state of affairs (an impossibility!), but to achieve outcomes that all parties should be willing to accept as imperfect relative to their respective starting positions. This goal is as much a process as it is an outcome. It ultimately reflects a shared (even if not fully agreed upon) recognition of conscious choices, necessary trade-offs, and a common vision of the practical realities of operating in a messy, imperfect world.

In the following Part, we articulate some specific examples of how this process can be carried

[14] Privacy is only one component of a comprehensive framework necessary to address system limits. Privacy, Equity, and Civil Rights Request for Comment, Docket No. 230103-0001, Question 1(a).

out to achieve balanced, sustainable, responsible, and defensible outcomes in the use of data and data-driven technologies to address important societal challenges. Before we describe the specific case studies, however, we also wish to acknowledge the need for and role of legislation to help facilitate our proposed re-framing.[15]

Appropriate legislation can curb the most harmful parts of these data-driven techniques without unduly restricting innovation. There are existing regulatory and legislative standards to draw upon in pointing the way to remedying data harms experienced by underserved or marginalized groups. One such example to consider is the European GDPR's approach to establishing evaluation frameworks that are risk-based and adaptable, rather than universally prescriptive. GDPR's heightened requirements for entities that engage in high-risk processing of personal data include prior consultations with authorities, data protection impact assessments, privacy by design, and information security measures.[16] An EU framework that is perhaps even more relevant for our discussion is the proposed AI Act which differentiates AI systems based on the context in which they are used and attaches to them a label of a (non-)risky system which then triggers more or less rigorous compliance measures.[17] Among others, the proposed act envisages conformity assessments and the adoption of harmonized standards for the high-risk systems.[18] While neither the GDPR nor the proposed AI Act are without flaws, approaching the AI-driven workflows through the prism of risk is helpful.

Legislative language that allows for flexible application of organizational, contractual, and technical measures for companies to adopt to ensure sufficient data privacy and protection against harms to marginalized groups is essential for dealing with the contextually dependent complexities of data bias and discrimination. For example, protected categories like ethnicity and race, which should be categorically excluded as decision criteria in the context of consumer lending are instead essential considerations to be explicitly factored into pharmaceutical evaluations of congenital health risks. Similarly, biometric identification technologies such as facial recognition may be so thoroughly opposed by local communities as to be banned entirely from common public safety uses but may still be wholly justifiable in limited and controlled settings such as high security facilities.

As we have advised in previous public contributions, the EU sets a useful legislative precedent in focusing on set minimum standards that can stand the test of time, ensuring capacity for

---

[15] The subsequent three paragraphs respond to Privacy, Equity, and Civil Rights Request for Comment, Docket No. 230103-0001, Question 4(f).

[16] *See* Regulation (EU) 2016/679, of the European Parliament and the Council of 27 April 2016 on the Protection of Natural Persons with Regard to the Processing of Personal Data and on the Free Movement of Such Data, and Repealing Directive 95/46/EC (General Data Protection Regulation), arts. 25, 34, 36, 2016 O.J. (L 119) 1 [hereinafter GDPR]; *see also Id.*, rec. 94.

[17] Proposal for a Regulation of the European Parliament and of the Council Laying Down Harmonised Rules On Artificial Intelligence (Artificial Intelligence Act) and Amending Certain Union Legislative Acts, COM(2021) 206 final.

[18] *See* Hadrien Pouget, *The EU's AI Act Is Barreling Toward AI Standards That Do Not Exist*, LAWFARE (January 12, 2023 8:16AM), https://www.lawfareblog.com/eus-ai-act-barreling-toward-ai-standards-do-not-exist ("It is up to harmonized standards to fill in the blanks left by the act, and the standard setters therefore bear the brunt of the responsibility for this compromise").

exceptions for granting exceptions for specific situations, and leaving space for organizations and standards bodies who are closer to the context of technology application to provide more specific implementation guidance.

III. Bridging the Two Cultures: Case Studies from Practice

This Part details two cases studies of how Palantir has approached privacy and equity considerations in building software that handles sensitive data. At Palantir, our customers use our software platforms to process their most sensitive information. Palantir does not collect, store, or broker customer data. Instead, we provide customers with software platforms to integrate, analyze, and operationalize the data they already lawfully control.

In working with customers across a variety of industries, from defense to retail, healthcare to telecommunications, we have gained insights into technologies and strategies to best enable customers to responsibly use their most sensitive data. We have selected two case studies regarding the processing of sensitive data from marginalized or disadvantaged groups. By sharing these case studies, we hope to better inform policy considerations that promote audits and oversight;[19] software design paradigms;[20] and industry and company codes;[21] among other mitigations for data harms. We also believe these case studies serve as instructive examples of how the alternative framing proposed in Part II can help facilitate practicable approaches to complex institutional challenges.

A. Workers as a marginalized group: leveraging technology to ensure employee protection

The use of technology has become deeply integrated in today's workplace. Workers have praised the increased flexibility that technology allows and have recognized its positive effects on their efficiency. At the same time, the automation and digitalization of the workplace have widely exposed employees' data, for instance, on their performance, the number of breaks they take, and the ways in which they cooperate with others. With the recent shifts toward flexible, non-permanent, and less certain employment practices, this means that modern workers are, just like other marginalized groups mentioned in the RFC, increasingly pushed into a disadvantaged position in the workplace.[22]

In using our software for their most critical operational challenges, some of our customers process personally identifiable information (PII), which may include sensitive information about their personnel. For example, this could encompass understanding the risk behaviors of military personnel or the shift schedules of employees in a business.[23]

---

[19] Privacy, Equity, and Civil Rights Request for Comment, Docket No. 230103-0001, Question 6(c).
[20] *Id.,* Question 6(d).
[21] *Id.,* Question 6(e).
[22] *See* Bart Custers and Helena U. Vrabec, *Worker Privacy in a Digitalized World Under European Law*, 39 COMP. LAB. L. & POL'Y J. 323 (2018).
[23] *See generally* Press Release, *Army Vantage's CRRT Completes Rollout to all Army Components* (Apr. 5, 2021) https://www.eis.army.mil/newsroom/news/data/army-vantages-crrt-completes-rollout-all-army-components ("The Commander's Risk Reduction Toolkit (CRRT)… [enables] command teams across all Army components to readily access and analyze aggregated risk data on Soldiers to help mitigate and avoid Soldiers' risk behaviors."); Dynamic

In the modern workplace, the tensions between the two poles we described in Part I are particularly strong. On the one hand, there is a push for accelerated adoption of powerful technologies, often predicated on the belief that data can solve operational challenges and that data technologies are infallible. It therefore comes as no surprise that many of the technologies used in today's workplaces rely heavily on the collection and use of personal information. On the other hand, there are calls to move away from using data-driven technologies due to bias inherent in all data and algorithms. Data – including incorrect data – can be used to infer characteristics and expectations that can be ascribed to workers, dangerously affecting the ways in which they are viewed by employers.

We have built several capabilities into our software platforms to help encourage responsible use for critical workflows that may include personnel data. These tools are designed to operationalize regulatory and organizational requirements regarding the processing of sensitive data. Moreover, they allow our customers to uphold "privacy by design" and similar principles that encourage the use of defaults to nudge users to more privacy-protective behaviors and contribute to reduced risk of disparate impact on workers.

Based on this experience, we detail below four recommendations on how organizations can use similar privacy, data governance, and data protection approaches to mitigate data harms to employees.[24] These recommendations do not comprise a holistic approach to handling employee information, but rather specifically demonstrate the utility of certain data protection techniques in mitigating the potential misuse or abuse of such data.

- **Foundational data protection technologies.** Access controls, versioning, and data lineage tools are the bedrock of any robust data protection infrastructure. First, access controls are necessary to ensure that data is not misused or repurposed beyond its intended use. Second, versioning is critical for allowing collaborative, operational use of data, while still maintaining transparency, accountability, and safety. Versioning tools are common in software engineering, and in our experience, bringing these capabilities to enterprise software has allowed more collaborative and more responsible use of data. Third, understanding the full provenance of a dataset and any derived data is an essential aspect of characterizing the limits of the data. Without understanding where a dataset comes from, the further necessary steps to understand concerns of bias or privacy become futile. For instance, understanding the data provenance has been recommended as a safeguard to support the validity and interpretability of the decision-making in the hiring process.[25] While nothing can guarantee perfect data, these infrastructure components can

---

Scheduling Primitives, https://www.palantir.com/platforms/foundry/scheduling-primitives/ (last visited Feb. 24, 2022).

[24] Privacy, data protection, and data governance are related frameworks for conceptualizing shared values and potential harms. In our experience, technical solutions for data governance and data protection may differ from those strictly focused on privacy, and a comprehensive approach that includes all of these terms may lend itself to a broader set of solutions and mitigations to potential harms. *See* Privacy, Equity, and Civil Rights Request for Comment, Docket No. 230103-0001, Question 1(a) ("Are there more comprehensive terms or conceptual frameworks [than 'privacy'] to consider?"); *see* supra note 14 and accompanying text.

[25] *The Future of Work: Protecting Workers' Civil Rights in the Digital Age, Hearing Before the Subcomm. on Civil Rights and Human Services*, 166th Cong. 13 (2020) (statement of Jenny R. Yang, Esq.) ("When the source code,

contribute to a more secure data foundation, necessary for both data analysis and the use of data-driven techniques like artificial intelligence or machine learning. We have long encouraged organizations to first consider these fundamentals of robust data protection before pursuing more specialized privacy-enhancing technologies.[26]

- **Purpose specification for sensitive actions.** All software platforms that process employee data should be able to request antecedent purpose justification. Requiring a justification before taking a sensitive action, like accessing or exporting employee records, can mitigate data harms in two different ways. First, asking a user for a justification *before* they take a sensitive action provides an opportunity for that user to consider the impact of their action and be more intentional about potentially sensitive operations. For example, a justification prompt can inform users about the heightened privacy implications of employee data and ask for an acknowledgement before allowing the action. Second, reviewing actions in a software platform along with user-submitted justifications provides a clearer understanding of why individuals have taken certain actions such as reviewing an employee's usage data. User-specified purposes can be captured in a software system's audit logs, and this audit trail can be reviewed in real-time or retroactively to assess whether data is being used for appropriate purposes. For example, we have built our software platforms with frameworks for configurable purpose specification, which enables compliance, governance, and privacy teams to request a justification before a user accesses or performs a certain operation on data.[27]
- **Scheduled deletion of sensitive data by default.** Technologists tend to like the idea of keeping data forever in case it might prove useful at some later point. Data that is no longer necessary for a specific business purpose, however, should be deleted, especially if such data is sensitive. Storage limitation and the regular deletion of collected data can reduce the risk of data misuse or repurposing. Sensitive data such as employee information, in particular, should be stored with scheduled deletion dates by default. If scheduled deletion is not appropriate, then the employer should have the burden to explain why data should be held indefinitely. Moreover, data deletion should encompass not only the originally collected data, but also any derived data or models trained on the data. In our experience, building software with comprehensive deletion tooling can encourage privacy by default and more accountable use of sensitive data.[28]

---

training data, and outputs are made available in a format that is understandable to an external party, bias in the data can be identified, flagged, and corrected.")

[26] *Privacy-Enhancing Technologies (PETs): An adoption guide (Palantir RFx Blog Series, #6)*, PALANTIR BLOG (2023), https://blog.palantir.com/privacy-enhancing-technologies-pets-an-adoption-guide-palantir-rfx-blog-series-6-b02dad56e9da ("Any given PET [Privacy-Enhancing Technology] is […] only effective if the underlying data foundation it is constructed upon is sound. To this end, organizations need strong controls over their data processing operations, including the ability to check their data for quality, accuracy, and representativeness.").

[27] *See* Future of Privacy Forum, *PEPR 2021: Session 8.2 - Lightweight Purpose Justification Service for Embedded Accountability,* YOUTUBE (Jun. 16, 2021), https://www.youtube.com/watch?v=T3aRNTa2Bwg (in which we discuss such approaches to purpose specification).

[28] *See generally* Paula Cipierre & Annabelle Larose, *Designing for Deletion (Palantir Explained, #6)*, PALANTIR BLOG (2022), https://blog.palantir.com/designing-for-deletion-palantir-explained-6-adfe25fda810 (describing how to design systems to meet complex deletion requirements); Palantir Technologies, *Data Lineage & Deletion*, (2022), available at https://www.palantir.com/assets/xrfr7uokpv1b/7ruwRAh1hvQiOCFdHGrjPt/5f9f71bc63229939d3d7bd051a316ca7/PCL-data-deletion_whitepaper_2022.pdf (detailing the approach we took to implement lineage-aware deletion in our Foundry software platform).

- **Data minimization:** Sensitive personnel data should always be minimized by default, and a variety of encryption, tokenization, or other pseudonymization techniques can be used to achieve this result.[29] If sensitive data needs to be preserved in its raw form, the onus should again be placed on the employer to explain why the intended purposes of use preclude data minimization at the outset. In our experience, data minimization techniques should be made available to users of all technical skill levels as implementing such privacy and compliance controls should be an interdisciplinary process with lawyers, engineers, data scientists, and domain experts. We have built tools for low-code and no-code data minimization that uses encryption to provide a voluntary, second layer of obfuscation on top of encryption-at-rest and encryption-at-transit measures.

These approaches can help organizations better balance the legitimate use of employee data with the privacy interests that employees have in their information. None of these techniques render the employee data unusable. Rather, they are designed to facilitate the responsible use of such data for legitimate workflows. By using these and other technical approaches, however, organizations can mitigate the risk of misuse or abuse of this sensitive information.

B. Patients as a vulnerable population: the importance of trust and transparency

The Covid-19 pandemic surfaced another vulnerable population – individuals with pre-existing and/or emergent health conditions. Across many nations, the response to Covid-19 was driven by data-driven technologies. There was a concern, however, that inconsistent and flawed data could undercut efforts to collect timely, actionable information to improve access to vaccines, testing, and life-saving health care services for those most in need.

Once again, "data evangelists" and "data detractors" dominated much of the discourse about how to respond to the pandemic. On the one hand, health data has been increasingly used to draw predictions regarding medical care, drive diagnostic decisions, and support operational workflows. The pandemic increased the need for not only high-quality and efficient medical care, but also tools for operationalizing large-scale public testing and vaccination campaigns. These opportunities are well-suited for the efficiency that technological, data-driven solutions can provide. On the other hand, using technology in a public health emergency brings unprecedented risk. Biases in the data foundation that a government agency collates to respond to a pandemic can have systemic and long-lasting repercussions. Moreover, the use of artificial intelligence or machine learning models, which have their own biases, could perpetuate stigma, lead to poor cohort representation, or suggest ineffective treatment modalities that may further discriminate against those with more limited access to care, such as minorities and economically disadvantaged individuals.

Palantir was at the forefront of the international Covid-19 response, providing organizations with the technical infrastructure needed to respond to supply chain, public health, and reporting challenges. We partnered with several public health agencies including the Department of Health and Human Services (HHS) in the United States and the National Health Service (NHS) in the

[29] *See* GDPR, art. 5(1)(c); *Id.*, recital 28.

United Kingdom, among other international counterparts.[30] These organizations had to comply with strict data protection requirements for processing sensitive data in the context of a public health emergency. Moreover, Palantir has long worked with healthcare and life sciences organizations in the private sector as well that have similar regulatory and organizational requirements around appropriately processing sensitive healthcare information.

By working with some of the world's leading healthcare providers and government agencies, we have come to understand transparency and trust in data as two critical considerations when working with sensitive information. First, transparency about the collected data and use of sensitive data not only promotes accountability, but it allows for a diverse breadth of stakeholders to collaborate on the problem at hand. Many mission-critical problems in healthcare require the domain expertise of scientists, medical professionals, public health experts and technologists as well as the partnership of the care recipients themselves. Upholding transparency as a key principle when handling sensitive data can facilitate visibility into the quality and completeness of the data an organization uses. This visibility is particularly helpful for marginalized groups, who may otherwise lack access or the means to determine how their data is used. Second, building trust in data is imperative for the use of any data-driven technology. If an organization uses data-driven technology, the data it decides to use will serve as the foundation for the organization's success or failures. Poor data quality can have compounding adverse effects, leading to misinformed options and decisions.[31]

To uphold these two principles when working with healthcare data, we detail below four recommendations that we have found to be critical in practice.

- **Regular data quality assessment:** Data health checks are one of the most effective ways to assess data quality, which is critical for ensuring both accurate and fair data outcomes. Automated checks can be designed to detect aberrations in the data — whether timeliness of data updates, completeness, consistency, or even identify missing contents — to ensure robust data quality at scale. Data pipelines can also listen for these checks and prevent propagation of the data downstream if any of the checks fail. Systematically assessing data quality is key to ensuring reliability of any flow of data, but it is even more important for data that knowingly includes sensitive data and data representative of vulnerable populations.
- **Use limitation restrictions:** Use limitation or purpose limitation restricts how a certain piece of data can be used, and it is a common principle in many data protection regulations.[32] Especially in the context of health information, it might be acceptable to use certain pieces of sensitive health data in one context, but not in others. In practice, it

[30] *See generally* Press Release, *Palantir and U.S. Government to Continue Work on COVID-19 Vaccine Distribution* (Jul. 6, 2021) https://www.palantir.com/newsroom/press-releases/palantir-and-us-government-to-continue-work-on-covid-19-vaccine/ (describing our work with the Department of Health and Human Services in their Covid-19 response); *Understanding our Work with the NHS*, PALANTIR BLOG (2022), https://blog.palantir.com/understanding-our-work-with-the-nhs-6d451beea022 (describing our work with the National Health Service in their Covid-19 response).

[31] Alice Yu, *Trust in Data (Palantir Explained #4)*, PALANTIR BLOG (2021), https://blog.palantir.com/trust-in-data-palantir-explained-4-c2adcdc31325.

[32] GDPR, art. 5(1)(b) ("Personal data shall be … collected for specified, explicit and legitimate purposes and not further processed in a manner that is incompatible with those purposes").

is access control mechanisms that allow software platforms to define which users can perform which operations on which pieces of data. Some more coarse configurations of access control systems can provide all-or-nothing access to data, but this could be too restrictive or too permissive for a given use case. Instead, more granular access control systems better support use limitation principles by guarding against the misuse and repurposing of data, while still allowing for legitimate uses to proceed. In our response to the Covid-19 pandemic, we implemented Purpose-Based Access Controls (PBAC) to help customers enforce use limitation.[33] This method of granular access controls enabled healthcare organizations to rapidly bring data together to respond to the pandemic, while also allowing governance teams to granularly administer and oversee the use of sensitive data.

- **Codes of conduct**: Beyond technical measures for enabling transparency and trust, codes of conduct are a complementary method for reducing potential data harms. A code of conduct allows an organization to put into writing the values, frameworks, and philosophy that guides its decisions.[34] While codes of conduct may not be comprehensive of all factors an organization considers when making a decision, it still provides valuable insight and transparency into how an organization that may control or process an individual's data weighs frameworks like privacy, equity, or civil rights.
- **Organizational data governance bodies:** Organizations can establish data governance bodies to help guide their decision-making in accordance with principles like privacy, equity, and civil rights. As we mentioned in our reflections on our work to support the international Covid-19 response, "we strongly encouraged and supported efforts to establish data governance bodies that could oversee the programs employing our software."[35] Technology, organizational policies, and even regional or national regulation can only adapt so quickly to the challenges of using data for mission-critical problems in practice. Data governance bodies can be another valuable source of input for understanding how a particular organizational decision would impact privacy and equity, especially in ever-changing business and regulatory environments.

IV. Conclusion

Palantir is not a data broker or a data collector, but since we have over a decade of experience providing data integration and analytics platforms to customers across various industries, we feel a responsibility to actively participate in the public discourse and share our insights about the modern digital economy. When it comes to addressing negative impacts of personal data processing, particularly on members of marginalized communities, we urge the NTIA to adopt a view that is neither too technology-centric nor technology-detracting. Instead, we believe the right way forward is to strive for balanced, sustainable, responsible, and defensible outcomes in

[33] Basil Jennings, *Purpose-based Access Controls at Palantir (Palantir Explained, #2)*, PALANTIR BLOG (2020), https://blog.palantir.com/purpose-based-access-controls-at-palantir-f419faa400b3.
[34] *See* Courtney Bowman, Reflections and Lessons from the COVID-19 Crisis, PALANTIR BLOG (2022), https://blog.palantir.com/reflections-and-lessons-from-the-covid-19-crisis-b406c03fbb4e ("In all of our COVID-19 response work across a multitude of public and private institutions, we have stayed true to our core values … outlined in our company Code of Conduct"); *See also,* Palantir Code of Conduct, (2020), available at https://s26.q4cdn.com/381064750/files/doc_downloads/governance/Code-of-Conduct.pdf.
[35] Bowman, *supra* note 34.

the use of data and data-driven technologies. Through our experience, we have had the opportunity to develop and test various analytical tools and processes that enhance protection of personal data, including data of some of the most vulnerable groups in our society. Our work has demonstrated compelling results that confirm technological development and data privacy are not mutually exclusive concepts. We hope that these case studies will help inform NTIA's policy considerations and ultimately lead to more effective standards and expectations for data processing that protects privacy, equity, and civil rights of individuals.