

Public Utility for What? Governing AI Datastructures

Julie E. Cohen*

Both in the U.S. and in Europe, initiatives for AI governance have focused principally on identifying and mitigating the risks created by AI models and their downstream uses rather than on those created by the datasets on which the models are trained. However, some of the most intractable dysfunctions of generative AI systems involve datasets. In particular, the very large datasets amassed by dominant providers of generative AI and related services are rapidly taking on infrastructural characteristics and importance. Effective AI governance therefore requires an infrastructural turn in thinking about data.

First, the Article explains the significance of the infrastructure lens and sketches some of the distinctive implications of data infrastructures, in particular, for governance of networked digital processes and the social and economic activities that they facilitate. Next, it explores two interrelated problems manifesting within generative AI systems—simulation and sociopathy—that illustrate the extent to which the project of AI governance is, unavoidably, a data governance project. In brief, generative AI models trained on mass content from the open internet are also trained on data infrastructures that have been developed for behaviorist, extractive purposes and that encourage the production and spread of particular kinds of content and particular styles of communication. Last, the article considers whether the concept of public utility, now the subject of growing interest among legal

* Mark Cluster Mamolen Professor of Law and Technology, Georgetown Law. Thanks to Hilary Allen, Chinmayi Arun, Dan Bateyko, William Boyd, Gordon Hull, Jake Goldenfein, James Grimmelman, Gabriel Nicholas, Paul Ohm, Blake Reid, Andrew Selbst, Alicia Solow-Niederman, Salome Viljoen, Rebecca Wexler, and participants in the Governing Data Symposium at Yale Law School, the NPU Spring 2025 Workshop Series, the AI and Human Values Manuscript Workshop at the Cornell Bowers College of Computing and Information Science, and a workshop session at the 2024 Privacy Law Scholars Conference for their thoughtful comments and to James Burns, Sabrina Deng, and Gent Salihu for research assistance.

scholars who study regulated industries, might supply a possible foundation for tackling the data governance problems associated with generative AI systems. The public utility model, however, addresses only some of the considerations that the infrastructure lens highlights. It is highly attuned to questions about access to infrastructures and their outputs but relatively insensitive to questions about infrastructure configuration and input sourcing. The problems of simulation and sociopathy belong in the latter category.

Article Contents

Introduction	138
I. Infrastructures and Datastructures: A Primer	139
II. AI Datastructures and Their Dysfunctions.....	145
A. “It’s the [Data] Economy, Stupid”: What’s Out, What’s In, and What’s Salient.....	146
B. “Ten Thousand Spoons When All You Need Is A Knife”: Flavors of AI Solutionism.....	154
1. Open source	155
2. Alignment.....	157
3. Risk assessment	160
4. Acceleration	162
III. The Promise and Limits of “Public Utility” Regulation 164	
A. “You Keep Using that Word”: What’s [a] Public Utility, Anyway?.....	165
B. “You’re Gonna Need a Bigger Boat”: Elements of a Framework for AI Datastructure Governance.....	170
1. Open source revisited: Disclosing datastructuring	172
2. Alignment revisited: Preventing distortive datastructuring.....	173
3. Risk mitigation revisited: Deterring wasteful datastructuring.....	175
4. Acceleration revisited: Catalyzing democratic datastructuring	177
Conclusion.....	179

Introduction

Both in the U.S. and in Europe, initiatives for AI governance have focused principally on identifying and mitigating the risks created by AI models and their uses rather than on those created by the datasets on which the models are trained. However, some of the most intractable dysfunctions of generative AI models involve datasets and data infrastructures. Effective AI governance therefore requires an infrastructural turn in thinking about data.

The very large datasets amassed by dominant providers of AI-related services are rapidly taking on infrastructural characteristics and importance. Part I explains the significance of the infrastructure lens and sketches some of the distinctive implications of data infrastructures for governance of networked digital processes and the social and economic activities that they facilitate. Part II explores two interrelated problems manifesting within generative AI systems—simulation and sociopathy—that illustrate the extent to which the project of AI governance is, unavoidably, a data governance project. In brief, generative AI models trained on content from the public internet are also trained on data infrastructures that have been developed for behaviorist, extractive purposes and that encourage the production and spread of particular kinds of content and particular styles of communication. Most leading approaches to AI governance now emerging around the world, however, generally ignore or down-weight the problems created by data infrastructures.

Part III considers whether the concept of public utility, now the subject of growing interest among legal scholars who study regulated industries, might supply a possible foundation for tackling the data governance problems associated with generative AI systems. Within the contemporary scholarly movement to reinvigorate the concept of public utility, there is definitional ambiguity about whether “public utility” is a type of resource (that ought to be dedicated or constrained to serving public values), a type of regulatory regime (that works or might work to impose public values), or both. Definitional ambiguity translates into prescriptive ambiguity about what

exactly a regime of public utility regulation ought to require. More fundamentally, however, the public utility toolkit focuses principally on ensuring access to infrastructures and their outputs. As the evolving history of energy and water regulation illustrates, the public utility model has proved inadequate to deal with new kinds of challenges relating to infrastructure configuration and input sourcing. The problems of simulation and sociopathy belong in the latter category. A regulatory framework for AI governance must contend more directly with the complex path dependencies embedded in platformized ecosystems that mobilize vast, dispersed data supply chains, and it must address values that tend to sit outside the frame of access regulation.

I. Infrastructures and Datastructures: A Primer

As the outputs of generative AI models increasingly come to permeate economic, social, and governmental domains, both the models and their training datasets will come to function as infrastructures supporting a wide variety of downstream activities. Accumulated learning about infrastructures—what they are, how they work, and the kinds of governance challenges they create—therefore holds important lessons for the coming AI era. And, notably, for the project of understanding emerging AI infrastructures to be most illuminating and useful, it is important to disentangle the models and the datasets and consider each separately. Here, I explain some relevant characteristics of resources that function as infrastructure and some additional characteristics of data infrastructures. As the discussion will suggest, the very large datasets that perform infrastructural roles within the platformized information economy and that play essential roles in training generative AI models merit special attention.

Infrastructures are constructed formations designed to facilitate human and social activities across spaces and scales.¹ They are built around shared standards and conventions that

¹ See generally Susan Leigh Star & Karen Ruhleder, *Steps Toward an Ecology of Infrastructure: Design and Access for Large Information Spaces*, 7 INFO. SYS. RES. 111 (1996).

structure both the inputs and the manner of those activities. Therefore, and inevitably, infrastructures also embody and constitute a kind of sociotechnical governance authority that inheres in the details of their design and operation. They afford but also configure and constrain downstream uses.² Decisions about infrastructure design and deployment also reflect prevailing wisdom about the soundest and best forms of economic and socio-economic development.³ Both the construction and the ongoing maintenance of infrastructures therefore implicate fundamentally political and political-economic questions.

Importantly, the particular way that infrastructures govern involves habituation to the patterns of affordance and constraint that they encode. As those patterns become more widely and deeply embedded in the ordinary operation of social and economic life, they become *transparent to use*.⁴ Notably, transparency to use has a meaning that is almost exactly the opposite of what “transparency” signifies in legal and technical discourses. Features that are transparent to use have become more or less invisible even as they exert powerful behavioral effects. The behaviors they afford and constrain have become habituated, and the extent of those behaviors’ dependence on particular constructed formations becomes visible to users only when breakdown occurs. Correlatively, then, the ongoing maintenance that infrastructures require is not only physical but also normative and behavioral. In an influential article, Susan Leigh Star and Karen Ruhleder explored the relationality and social embeddedness of such work.⁵ Transparency to use and visibility upon breakdown may or may not correspond to technical transparency. However,

² Blake Hallinan & James Gilmore, *Infrastructural Politics Amidst the Coils of Control*, 35 CULT. STUD. 617 (2021).

³ Nikhil Anand, Akhil Gupta, & Hannah Appel, “Introduction: Temporality, Politics, and the Promise of Infrastructure,” in NIKHIL ANAND, AKHIL GUPTA, & HANNAH APPEL, *THE PROMISE OF INFRASTRUCTURE 1* (2018); William Rankin, *Infrastructure and the International Governance of Economic Development, 1950-1965*, in INTERNATIONALIZATION OF INFRASTRUCTURES 61 (2009).

⁴ Star & Ruhleder, *supra* note 1, at 113-14; *see also* Hallinan & Gilmore, *supra* note 2, at 635-36.

⁵ Star & Ruhleder, *supra* note 1.

technically black-boxed features of infrastructure can become transparent to use in the behavioral sense, and when breakdown occurs, lack of technical transparency may become very important.

Intangible resources, including datasets, can and do serve as infrastructures.⁶ Datasets have structures of their own, which result from many types of decisions about inclusion, categorization, indexing, meta tagging, and cross-linking.⁷ At scale, such decisions can become powerful vehicles for structuring and organizing social and economic activities.⁸ Put another way, under the right circumstances, datasets can come to function as infrastructures in their own right. The construction and maintenance of data infrastructures also is ongoing; Flyverbom and Murray refer to the sum total of this work as datastructuring.⁹ As we will see, data infrastructures (or, as I will call them, *datastructures*) introduce new difficulties into the study of infrastructures and infrastructuring work.¹⁰

⁶ See generally BRETT M. FRISCHMANN, *INFRASTRUCTURE: THE SOCIAL VALUE OF SHARED RESOURCES* (2012); Star & Ruhleder, *supra* note 1; Rankin, *supra* note 3, at 66-67, 72-73; Hallinan & Gilmore, *supra* note 2.

⁷ See generally GEOFFREY BOWKER & SUSAN LEIGH STAR, *SORTING THINGS OUT: CLASSIFICATION AND ITS CONSEQUENCES* (2000); LISA GITELMAN, ED., *RAW DATA IS AN OXYMORON* (2013).

⁸ Mikkel Flyverbom & John Murray, *Datastructuring—Organizing and Curating Digital Traces into Action*, *BIG DATA & SOC'Y* (2018), <https://doi.org/10.1177/2053951718799114>; Alex Gekker & Sam Hinds, *Infrastructural Surveillance*, *22 NEW MEDIA & SOC'Y* 1414 (2020); ROB KITCHIN, *THE DATA REVOLUTION: BIG DATA, OPEN DATA, DATA INFRASTRUCTURES AND THEIR CONSEQUENCES* (2014); see also Will Orr & Kate Crawford, *The Social Construction of Datasets: On the Practices, Processes and Challenges of Creating Datasets for Machine Learning*, *26(9) NEW MEDIA & SOC'Y* 4955 (2024), <https://doi.org/10.31235/osf.io/8c9uh>.

⁹ Flyverbom & Murray, *supra* note 8.

¹⁰ This usage is analogous to but distinct from the meaning of “data structure” in computer science. In computer science, “a *data structure* is a collection of data values, the relationships among them, and the functions or operations that can be applied to the data”—or, put differently, it is a technical arrangement for storing data values that makes particular manipulations possible. Peter Wegner & Edwin D. Reilly, *Data Structures*, in *ENCYCLOPEDIA OF COMPUTER SCIENCE* 507, Jan. 2003, <https://dl.acm.org/doi/10.5555/1074100.1074312>. As used here interchangeably, *datastructures* or *data infrastructures* are vast, scaled up datasets, developed and maintained via correspondingly scaled up

The contemporary, platformized information economy relies on a particular kind of datastructuring work.¹¹ The giant platforms that now dominate the emergent global political economy of informational capitalism have attained their dominance by pursuing so-called “ecosystem” strategies for layering proprietary data collection and exchange protocols systematically across the underlying infrastructure of the open internet. Those strategies rely on software developer kits (SDKs), which are off-the-shelf utilities prepared and distributed by platform proprietors to website and mobile app developers. Developers use the SDKs to enable a wide range of platformized functions—for example, location mapping, ad placement, and universal login. Platforms use SDKs to harvest data from website and mobile app users, transmit it to their data centers to augment already-existing user profiles, and push profile-driven functionality and content back to users via websites and mobile apps.¹²

The massive datasets derived from these activities—comprised of continuous flows of data derived from behavioral surveillance, continuously structured to conform to the protocols used by competing dominant platform ecosystems—perform a variety of infrastructural functions in their own right. They drive personalization of digital content and digital services and undergird complex targeted advertising strategies. For both individual and business users of platformized services, these functions are habituated and largely transparent to use; they are the way the digital economy “works.” Many—though not all—breakdowns also are labeled as ordinary events, quickly resolved via “patches” and “bug fixes.” Many other

arrangements for data collection, storage, and exchange, that facilitate, structure, and constrain social and economic activities.

¹¹ Some of the remaining material in this section is adapted from Julie E. Cohen, *Infrastructuring the Digital Public Sphere*, 25 YALE J.L. & TECH. 1 (2023) [hereinafter Cohen, *Infrastructuring*]; and Julie E. Cohen, *Data, Infrastructures, and Infrastructure Stacks*, in GLOBAL GOVERNANCE BY DATA: INFRASTRUCTURES OF ALGORITHMIC RULE (Fleur Johns, Gavin Sullivan, & Dimitri Van Den Meerssche, eds., forthcoming 2025 or 2026) [hereinafter Cohen, *Infrastructure Stacks*].

¹² See generally Ahmad Ghazawneh & Ola Henfridsson, *Balancing Platform Control and External Contribution in Third-Party Development: the Boundary Resources Model*, 23 INFO. SYS. J. 173 (2013).

third-party business models now coevolve with platformized datastructures. Third parties with data-driven business models, such as data brokers and data analytics consultancies, routinely seek and receive permission to embed their own SDKs in popular websites and apps and structure the services they offer around the types of data that such arrangements allow them to collect. Others scrape publicly available content and data from the ecosystems that platforms have created.

Platformized datastructures differ from more traditional physical infrastructures in some important ways. On one hand, they depend more heavily on other infrastructures. They are layered atop both platform protocols and underlying elements of the internet stack, and the very large datasets that comprise them are stored in large data centers that require enormous reserves of energy to operate. As a practical matter, the one who controls the platform protocols and data centers also has de facto control of the datastructures. On the other hand, because platformized datastructures—by design—scale both vertically and laterally, they are nimble, flexible, and adaptable to new uses that the platform or a third-party developer (or a government actor) might envision.¹³ Those characteristics, in turn, give datastructures (and their controllers) two kinds of power that physical infrastructures and technical protocols alone did not possess.

First, the malleability of platformized datastructures can render a wide range of activities more amenable to control and cooptation, with stark implications for both the nature and the distribution of downstream uses. For example, whoever controls a road can control the flow of traffic along that road, but those who control mapping and mobility datasets have more comprehensive and flexible power to reshape flows of people and vehicles through physical spaces—and also to encourage or discourage certain destinations or restrict certain people’s mobility.¹⁴ Datastructures that attain sufficient scale and power can begin to drive wholesale reorganization of

¹³ Cohen, *Infrastructure Stacks*, *supra* note 11.

¹⁴ *Id.*; Gekker & Hinds, *supra* note 8; Xiaohan Zhang, *Decoding China’s COVID-19 Health Code Apps: The Legal Challenges*, 10 HEALTHCARE (BASEL) 1479 (2022), <https://doi.org/10.3390/healthcare10081479>.

patterns of economic and social activity—and sometimes even reorganization of the underlying physical infrastructures on which those activities rely. The massive commercial and logistical datastructures controlled by Amazon and Alibaba Group have changed the ways that manufactured goods flow around the world, and, in the process, have produced significant reconfiguration of physical distribution networks around newly developed, proprietary hubs and transit fleets.¹⁵

Second, the malleability of platformized datastructures also facilitates other strategies for attaining dominance that operate via manipulation rather than by control. As data-driven, algorithmic processes structure and personalize *individual* processes of information retrieval and public discourse, they also shape *societal* processes of knowledge production—and they do so in ways both enabled and constrained by the datasets on which they have been trained and over which they operate.¹⁶

The emergent generative AI era will accelerate both kinds of effects and introduce new ones. More than half of U.S. adults self-report using generative AI for a wide variety of everyday activities. Increasingly, people use generative AI to search for information about everything from products and services to health and current events. They also use generative AI to compose, edit, and post text, images, and videos online.¹⁷ As the economist Cecilia Rikap has written, the market-leading generative AI systems are coming to operate as “intellectual monopolies,” shaping collective understanding, social interaction, and cultural production over millions of recursive

¹⁵ Cohen, *Infrastructure Stacks*, *supra* note 11; Jean-Paul Rodrigue, *The Distribution Network of Amazon and the Footprint of Freight Digitalization*, 88 J. TRANSPORT GEOG. 102825 (2020); Ellie Falcone, John Kent, & Brian Fugate, *Supply Chain Technologies, Interorganizational Network and Firm Performance: A Case Study of Alibaba Group and Cainiao*, 50 INT’L J. PHYS. DISTRIBUT. & LOGISTICS MGMT. 333 (2020).

¹⁶ Cohen, *Infrastructuring*, *supra* note 11, at 15; Salome Viljoen, Jake Goldenfein, & Lee McGuigan, *Design Choices: Mechanism Design and Platform Capitalism*, 8 BIG DATA & SOC’Y (2021), doi:10.1177/205395172111034312.

¹⁷ Lee Rainie, *Close Encounters of the AI Kind: The Increasingly Human-Like Way People Are Engaging with Language Models*, Imagining the Digital Future Center, Mar. 2025, <https://perma.cc/PY2B-L9V5>.

interactions.¹⁸ The very large datasets used to train such systems—which I will call AI datastructures—are therefore essential objects of study. As Part II explains, AI datastructures are inextricably linked to certain kinds of output dysfunctions.

II. AI Datastructures and Their Dysfunctions

Consider two different kinds of reliability problems with the outputs of generative AI models. One is that they sometimes produce fictional information while representing that information as factual. It is conventional to refer to this phenomenon as hallucination or confabulation, but that terminology conceals a pitfall that will become important later. The idea that a model might “hallucinate” or “confabulate” rests on an anthropomorphic analogy (i.e., models can “see” and “understand”), and so it may engender misconceptions about what AI models do. Therefore, I will call this problem the problem of *simulation*. A second problem is that generative AI models sometimes produce content that is amoral or antisocial when judged against prevailing cultural and ethical values (although, to the extent those values are contested, such judgments will also be contested). It is conventional to refer to this phenomenon as misalignment. As I will explain, though, that terminology reads society too far out of the equation, reframing a problem that originates in a behaviorist, extractive approach to human and social data as a predominantly technical challenge while taking background datastructuring practices largely for granted. To underscore the sociotechnical origins of this problem, I will call it the problem of *sociopathy*.

The infrastructure and datastructure lenses enable productive reframing of both problems because they direct attention to the ways that both are baked deeply into the design of platformized digital systems and the kinds of interactions those systems are designed to reproduce. They also help to explain the less-than-stellar track record of various interventions that have been touted as “solutions” to AI’s dysfunctions.

¹⁸ Cecilia Rikap, *Dynamics of Corporate Governance Beyond Ownership in AI*, COMMON WEALTH, May 15, 2024, <https://perma.cc/MHE5-MNPY>.

A. “*It’s the [Data] Economy, Stupid*”: *What’s Out, What’s In, and What’s Salient*

Data-driven practices of knowledge production do not simply reproduce pre-existing realities; they create new ones. In the book *An Engine, Not a Camera*, the sociologist Donald MacKenzie analyzed the operation of complex systems for collecting, communicating, and acting on financial information.¹⁹ It is conventional to think that such systems simply work like cameras, capturing snapshots of behaviors in markets that already existed. MacKenzie showed, however, that financial systems are deeply generative, inventing the relationships they purport merely to describe. They act, in other words, as engines, bringing socially constructed worlds into being. Choices about which types of data matter and how to juxtapose them have enormously important consequences. So too for AI datastructures. As this Section explains, datastructures derived from mass content on the open internet bring the socially constructed world into being in a very particular way.

Although the market leading providers of generative AI models tend to guard the specific details about their training datasets closely, it is clear that very large training datasets originating from the open internet function as AI datastructures for such models. For purposes of this paper, I take the “open internet” to include platformized ecosystems—e.g., video content from YouTube, social feeds from larger platforms such as Instagram, TikTok, or Twitter/X, discussion threads on smaller platforms such as Reddit and Yelp, content from gaming platforms such as Twitch and Discord, and self-publication platforms such as Medium and Substack—alongside the millions of other pages of openly available content made available by companies, organizations, and individuals via their personal domains. Some organizations that traditionally have relied on advertiser support—most notably media companies—also adopt a (partially) platformized structure for the content they provide, interpolating third-party adtech as well as specialized

¹⁹ DONALD MACKENZIE, *AN ENGINE, NOT A CAMERA: HOW FINANCIAL MODELS SHAPE MARKETS* (2008)

technologies for tracking user engagement internally.²⁰ I will refer to the content that circulates within platformized ecosystems (including advertiser-provided content) as “mass content.” Therefore, some of the analysis below also applies to the platformized content made available by media companies. The specific questions that will occupy the remainder of this Section concern the types of patterning that AI datastructures derived from mass content encode and their relation to problems of simulation and sociopathy.

Most of the scholarly discussion has focused on the kinds of content that AI datastructures include and exclude, so it is useful to begin—though not to end—there. AI datastructures predictably exclude or minimize the impact certain kinds of content have on model training. An influential paper by a group of now-former Google data ethicists explains, for example, that large language models predictably down-weight minority languages and cultures.²¹ (We have been here before, pondering the causes and effects of language failures in content moderation.²²) Such omissions are important because different languages encode different meaning structures and leave unstated different kinds of ambiguity. Grammar and usage, moreover, are not the same thing; the latter is informed by complex contextual judgments that humans learn to make automatically but that may pose challenges for non-humans. The linguist Noam Chomsky illustrated this point with the syntactically correct but substantively meaningless phrase “colorless green ideas sleep furiously.”²³

²⁰ See generally FRANKLIN FOER, *WORLD WITHOUT MIND: THE EXISTENTIAL THREAT OF BIG TECH* (2017); Erin C. Carroll, *Platforms and the Fall of the Fourth Estate: Looking Beyond the First Amendment to Protect Watchdog Journalism*, 79 MD. L. REV. 529 (2020).

²¹ Emily M. Bender, et al., *On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?*, FACCT '21 (2021), <https://doi.org/10.1145/3442188.3445922>; see also Paula Helm, et al., *Diversity and Language Technology: How Language Modeling Bias Causes Epistemic Injustice*, 26 ETHICS & INFO. TECH. 8 (2024), <https://doi.org/10.1007/s10676-023-09742-6>.

²² Gabriel Nicholas & Aliya Bhatia, *The Dire Defect of ‘Multilingual’ AI Content Moderation*, WIRED, May 23, 2023, <https://perma.cc/YN6T-SKPW>.

²³ NOAM CHOMSKY, *SYNTACTIC STRUCTURES* 15 (1957).

A related kind of exclusion problem relates to the fact that in human and social interactions, many kinds of cultural and even technical knowledge are contextual and tacit. Without an enormous amount of human guidance—which inevitably comes with built-in cultural filters of its own—generative AI systems cannot recognize, describe, or convey the tacit knowledge that may be required to interpret content elements contained within AI datastructures.²⁴ For example, ChatGPT can be trained to write funny headlines that emulate those from particular sources, but it has a harder time writing comedy routines.²⁵ AI chatbots carefully trained to provide certain types of mental health advice can supplement an overtaxed mental health system, but generative AI “companions” designed for more open-ended interactions have a harder time recognizing and responding appropriately to conversational danger signals that signal suicide risk.²⁶ Generative AI systems for image classification can identify patterns and use them to make predictive classifications, but they have a harder time decoding the multimodal and contextual elements that would enable a human observer to make sense of what the image depicts—for example, whether “two people walking together” are friends confiding in one another, coworkers arriving at the office simultaneously, or a suspect being led away by a police officer.²⁷ While it is a mistake to assume that generative AI

²⁴ For good summaries of the difficulties and challenges, see K. D. Fenstermacher, *The Tyranny of Tacit Knowledge: What Artificial Intelligence Tells us About Knowledge Representation*, PROC. 38TH ANN. HAWAII INT’L CONF. ON SYS. SCI. 243a (2005), doi:10.1109/HICSS.2005.620; Lene Pettersen, *Why Artificial Intelligence Will Not Outsmart Complex Knowledge Work*, 33 WORK, EMPL. & SOC’Y 1058 (2018), <https://doi.org/10.1177/0950017018817489>; see also Harry Collins, *Why Artificial Intelligence Needs Sociology of Knowledge: Parts I and II*, 40(3) AI & SOC’Y 1249 (2024), doi:10.1007/s00146-024-01954-8.

²⁵ Jack Hessel, et al., *Do Androids Laugh at Electric Sheep? Humor “Understanding” Benchmarks from the New Yorker Caption Contest* (2023), arXiv:2209.06293 [cs.CL].

²⁶ Kevin Roose, *Can A.I. Be Blamed for a Teen’s Suicide?*, N.Y. TIMES, Oct. 23, 2024, <https://perma.cc/P8G2-9CUR>; Julian De Freitas & I. Glenn Cohen, *The Health Risks of Generative AI-Based Wellness Apps*, 30 NAT. MED. 1269 (2024).

²⁷ Donald Geman et al., *Visual Turing Test for Computer Vision Systems*, 112 PNAS 3618 (2015); see Lisa Feldman Barrett, et al., *Emotional Expressions Reconsidered: Challenges to Inferring Emotion from Human*

systems could not “learn” over time to perform these and other open-ended tasks better or even well, it is also a mistake to assume that they will learn to do them perfectly or even well enough.

These points about exclusion and tacit knowledge might be understood as nesting comfortably within a large and growing corpus of scholarship on the ineradicability of bias in data-driven, predictive models. As is now well established, large scale ventures in data-driven, predictive patterning tend to produce results that reflect preexisting patterns of cultural and structural bias and disadvantage.²⁸ Making private or public policy based on the results of such patterning reifies accepted “truths” about the appropriate allocation of economic and social resources that are anything but neutral.²⁹ In legal scholarship, discussions of the pitfalls of data-driven predictive patterning have tended to revolve around the problem of bias and steps for addressing it (i.e., identifying bias, training it out of machine learning systems, evaluating the results of those efforts, and so on). Those are worthy projects, but the project of this paper is different.

In what follows, I reframe observations about what generative AI excludes or fails to notice as pointing to a larger

Facial Movements, 20 PSYCH. SCI. IN PUB. INT. 1 (2019), <https://doi.org/10.1177/1529100619832930>; Sangmin Lee et al., *Towards Social AI: A Survey on Understanding Social Interactions*, Sept. 5, 2024, arXiv:2409.15316 [cs.HC]; Maarten Sap et al., *Neural Theory-of-Mind? On the Limits of Social Intelligence in Large LMs*, Oct. 24, 2022, arXiv:2210.13312 [cs.CL]; Terrance DeVries, et al., *Does Object Recognition Work for Everyone?*, Jun. 18, 2019, arXiv:1906.02659v2 [cs.CV].

²⁸ Solon Barocas & Andrew Selbst, *Big Data's Disparate Impact*, 104 CAL. L. REV. 671 (2016); Bender, et al., *supra* note 21; Abeba Birhane, et al., *Multimodal Datasets: Misogyny, Pornography, and Malignant Stereotypes*, Oct. 5, 2021, arXiv:2110.01963v1 [cs.CY]; Jesse Dodge, et al., *Documenting Large Webtext Corpora: A Case Study on the Colossal Clean Crawled Corpus*, Sept. 30, 2021, arXiv:2104.08758v2 [cs.CL].

²⁹ Barocas & Selbst, *supra* note 28; Mehtab Khan & Alex Hanna, *The Subjects and Stages of AI Dataset Development: A Framework for Dataset Accountability*, 19 OHIO ST. TECH. L.J. 171 (2023); Yeshimabeit Milner & Amy Traub, *Data Capitalism + Algorithmic Racism*, DEMOS (2021), <https://perma.cc/NU6G-H96F>.

question about the kinds of epistemic force fields that exist in AI datastructures and argue that the answer depends in part on factors that human authors of AI-critical scholarship often fail to consider.

Begin by reconsidering the kinds of content that mass content datasets include. The political economy of informational capitalism depends heavily on free content, and free content has three particular characteristics. First, much of it is user generated.³⁰ Although it is clear that market-leading generative AI models are trained on a variety of content, including some that reflects more traditional and exacting editorial and curation standards, they are also trained on user generated material from sites such as Facebook and Instagram, X, YouTube, Reddit, Medium, and platformized mass media content from sites such as Huffington Post and Breitbart.³¹ AI systems trained on this content do not just simulate human social and cultural production; they simulate the particular types of social and cultural production that those constructed digital environments are designed to elicit. Second, free content from the open internet includes a lot of paid promotional material. Such material appears in many different places—including not only ad banners, popups, and video inserts but also as advertorials, sponsored features, and reviews by social media influencers.³² It also serves a range of purposes

³⁰ See generally Guy Pessach, *Beyond IP—The Cost of Free*, 54.1 OSGOODE HALL L.J. 255 (2016).

³¹ See Cade Metz, et al., *How Tech Giants Cut Corners to Harvest Data for A.I.*, N.Y. TIMES, Apr. 8, 2024, <https://perma.cc/ZLJ5-5PTN>; Alan D. Thompson, *What's In My AI?*, LifeArchitect.ai, Mar. 2022, <https://perma.cc/XNM8-66SL>; Farhad Manjoo, *Reddit and the Robots*, SLATE, June 21, 2024, <https://perma.cc/RB4K-39RR>; Annie Gilbertson, *Apple, Nvidia, Anthropic Used Thousands of Swiped YouTube Videos to Train AI*, WIRED, July 16, 2024, <https://perma.cc/HN2G-3F35>; Jess Weatherbed, *Anthropic's Crawler Is Ignoring Websites' Anti-AI Scraping Policies*, THE VERGE, July 25, 2024, <https://perma.cc/3HNZ-LATP>; Jess Weatherbed, *Meta Fed Its AI on Almost Everything You've Posted Publicly Since 2007*, THE VERGE, Sept. 12, 2024, <https://perma.cc/HB4P-QYHD>.

³² See Corinna Lauerer & Johannes and Beckert, *Pushing Boundaries—Hybrid Advertising in Digital News Media: A Content Analysis of Media Kits*, DIGIT. JOURNALISM 1(2024); Samantha Scott, *The Most Important Influencer Marketing Statistics for 2025*, MELTWATER, Feb. 2, 2024, <https://perma.cc/J6AY-UYLZ>.

that include not only selling goods and services but also furthering cultural and political influence campaigns.³³ AI systems deliberately or inadvertently trained on this content may be expected both to overweight the importance of persuasion and to internalize particular types of persuasion strategies. Third, an increasing amount of free content from the open internet is itself produced using generative AI, and such content can be untrustworthy for a variety of reasons.³⁴ Over time, generative AI models trained on the outputs of generative AI models may be expected to overweight the messages that such outputs convey.

Focusing exclusively on the kinds of *content* that are included in or excluded from AI training datasets, however, also reflects human epistemic biases. Human intelligence is trained on content and context in the first instance. To an extent, humans also engage in what we might optimistically call data-driven patterning, but the human capacity for patterning is also limited and can fall into well-recognized errors.³⁵ When a generative AI model is trained, the content-metadata polarity is reversed: there is no content and no context but only tokens and styles that inform continual, stochastic processes of pattern identification.³⁶ For models, datasets are always already datastructures. And what is most salient for models that train

³³ Anthony Nadler, Matthew Crain, & Joan Donovan, *Weaponizing the Digital Influence Machine: The Political Perils of Online Ad Tech*, DATA & SOC'Y (2018), <https://perma.cc/6QEC-XP78>.

³⁴ Samuel C. Woolley, *Bots and Computational Propaganda: Automation for Communication and Control*, in SOCIAL MEDIA AND DEMOCRACY: THE STATE OF THE FIELD, PROSPECTS FOR REFORM 89 (Nathaniel Persily & Joshua A. Tucker, eds. 2020); Josh Goldstein & Renée DiResta, *Propagandists Are Using AI Too—and Companies Need to Be Open about It*, MIT TECH. REV., Jun 8, 2024, <https://perma.cc/3MKD-55E7>; Creston Brooks, Samuel Eggert, & Denis Peskoff, *The Rise of AI-Generated Content in Wikipedia*, Oct. 10, 2024, arXiv:2410.08044v1 [cs.CL].

³⁵ The literature here is vast. For the foundational work, see Daniel Kahneman & Amos Tversky, *Judgment under Uncertainty: Heuristics and Biases*, 185 SCI. 1124 (1974).

³⁶ See Julia Witte Zimmerman et al., *Tokens, the Oft-Overlooked Appetizer: Large Language Models, the Distributional Hypothesis, and Meaning*, Dec. 14, 2024, arXiv:2412.10924 [cs.CL]; Bender, et al., *supra* note 21.

and retrain themselves on datastructures may be very different than what is salient to human observers.

Recall now that the disaggregated digital architectures that comprise platform ecosystems do not simply elicit content; they are designed to generate data, and not just any data. Rather, they have been systematically developed to produce data about the kinds of content that generate engagement and spread socially. Those engagement patterns, in turn, elicit additional, similar content and ensure that some kinds and some specific pieces of content are accessed and viewed far more frequently.³⁷ This, in turn, means that the tokens and styles of significance to generative AI models are produced and reproduced in particular patterns and not others. Data about the content and uptake of targeted promotional content also matter in the relevant sample sets, and the fact that such material will have been targeted to audiences, to content, or to both deepens and reinforces the relevant patterning. These points about the *nature* and *orientation* of datastructuring, moreover, should be understood as distinct from claims about the urgent necessity of *scale* in training datasets for successive generations of generative AI models (to which I return below).³⁸ Data do not exist in a vacuum; they derive from distributed sociotechnical systems that have been constructed for particular purposes. If those systems were constructed differently, the universe of “all the data” derived from them would look and behave differently.

All of this datastructuring shapes and should be expected to shape AI outputs. We should not be surprised, therefore, to see generative AI performing particular kinds of social and cultural reproduction. Begin with the content that amazes us. A colleague recently reported to me that he became convinced the most recent version of ChatGPT had made a great leap forward in functionality when he could instruct it to produce a really great role-playing game involving Supreme Court cases. From a different perspective, it is only logical that a system

³⁷ Nadler, Crain, & Donovan, *supra* note __.

³⁸ Jared Kaplan, et al., *Scaling Laws for Neural Language Models*, Jan. 23, 2020, arXiv:2001.08361v1 [cs.LG]; *see also* Birhane, Prabhu, & Kahembwe, *supra* note 28 (describing the prevailing thinking about generative AI training as “scale beats noise”).

trained at the intersection of gamer culture, gamified social media, gamified marketing, and gamified self-promotion would learn to gamify and would be able to apply that skill to a type of subject matter that is the repeated subject of public dissection and punditry. And it is only logical that a system trained on the voluminous outputs of policy think tanks, TedX, and Substack would be able to produce nonfiction essays whose flawless syntax and sophisticated vocabulary are matched only by their substantive banality.³⁹

Consider next the kinds of content that alarm or appall us. A decade of work on the pitfalls of data-driven patterning suggests that generative AI models will always present residual risks of simulation, but we should expect those risks to be far more significant for models trained on datastructures engineered to validate and boost conspiracy theories, alternative facts, “gotcha” moments, and herd-mentality cancellation campaigns. Some ongoing degree of dispute about the values that are or should be encoded in generative AI outputs is inevitable; this is what Grimmelmann, Reid, and Rozenshtein call the “baseline hell” problem.⁴⁰ But when generative AI models are trained on datastructures that include Logan Paul and Jackass videos, Tide Pod challenges, and Grand Theft Auto-style mayhem, we also should not be surprised when stochastic patterning processes produce, to take one noteworthy recent example, a video of SpongeBob SquarePants flying a plane into the twin towers in a demented parody of the apocryphal Hollywood pitch strategy (“It’s *Wreck-It Ralph* meets *United 93*”).⁴¹ When they are trained on datastructures that include nativist screeds and hate propaganda, we should not be surprised to see them generate flawless syllogisms presenting sociopathic logics about the

³⁹ For a thought-provoking exploration of the intersections between automated cultural production and aesthetics, see LEV MANOVICH & EMANUELE ARIELLI, *AUTOMATED AESTHETICS* (2024).

⁴⁰ James Grimmelmann, Blake Reid, & Alan Rozenshtein, *Generative Baseline Hell and the Regulation of Machine Learning Foundation Models*, LAWFARE, May 8, 2024, <https://perma.cc/DT4B-JXWK>.

⁴¹ Emilia David, *Bing’s AI Image Generator Tries to Block ‘Twin Towers’ Prompts, But It’s Not Working*, THE VERGE, Oct. 5, 2023, <https://perma.cc/9DED-ULWK>.

imminent danger of racial replacement.⁴² When they are trained to produce confident-sounding answers to questions about complicated social, economic, or scientific phenomena, we should not be surprised to discover that both the answers and the supporting citations are sometimes invented based on probabilistic predictions about the expected forms of such answers—to discover, in other words, that the outputs are bullshit.⁴³

Simulation and sociopathy are, so to speak, bred in the bone. They are as-if-innate, infrastructurally overdetermined attributes of a political economy organized around attention and extraction. None of this should come as a surprise because it is how infrastructures work. Even as the role of behavioral data within the contemporary, platformized political economy has come under increasingly intense criticism, the behavioral datastructures resulting from the attention and extraction imperatives have become increasingly transparent to use. It is urgently important to understand the repeated recurrence of simulation and sociopathy as markers of encoded dysfunction—as a kind of visibility upon breakdown that calls for a collective and considered response.

A. “*Ten Thousand Spoons When All You Need Is A Knife*”: *Flavors of AI Solutionism*

The major approaches that have been touted as ways to improve the performance of generative AI models are mostly nonresponsive to the dysfunctions lurking within AI datastructures. Here, I discuss four major clusters of approaches, which relate to model transparency, fine-tuning, risk assessment and mitigation, and acceleration. Each

⁴² Jonathan M. Katz, *Substack Has a Nazi Problem*, THE ATLANTIC, Nov. 28, 2023, <https://perma.cc/PL35-XYQB>; Kelsey Piper, *Grok’s MechaHitler Disaster Is a Preview of AI Disasters to Come*, VOX, Jul. 11, 2025, <https://perma.cc/83KG-9QML>.

⁴³ Klaudia Jazwinska & Aisvarya Chandrasekar, *AI Search Has a Citation Problem*, COLUM. JOURNALISM REV., Mar. 6, 2025, <https://perma.cc/ECW9-NAHV>; Kyle Orland, *Why Do LLMs Make Stuff Up? New Research Peers Under the Hood*, ARS TECHNICA, Mar. 28, 2025, <https://perma.cc/UR8L-2BH7>. See generally Michael Townsen Hicks, James Humpries & Joe Slater, *ChatGPT Is Bullshit*, 26 ETHICS & INFO. TECH. 38 (2024), <https://perma.cc/9VDS-ADF8>.

represents a very different approach to the accountability problem; none addresses datastructures. My very strong suspicion is that at least some market-leading AI developers and their investors understand this and view sound and fury about models and risks as a useful way of deflecting attention away from baseline data quality issues. In what follows, I will mostly bracket questions about motivation, which I address elsewhere.⁴⁴

1. Open source

One influential set of discussions about accountability in AI development draws on the traditions and practices of open source software development. Advocates for open source AI—including most prominently Meta—argue that model algorithms and weights should be publicly disclosed. Drawing on the open source community’s venerable maxim that “with enough eyeballs, all bugs are shallow,” they contend that widespread scrutiny of models and weights will turbo-charge the process of eliminating unreliable and/or undesirable outputs.⁴⁵ (Critics of this position argue that open access to models and weights by downstream actors will magnify the various social risks that AI creates.⁴⁶)

Open source rhetorics have powerful appeal in tech policy discussions for reasons that are both technical and ideological, but open source software development has notoriously failed to counteract one particular kind of technological closure

⁴⁴ Julie E. Cohen, *Oligarchy, State, and Cryptopia*, 94 FORDHAM L. REV. 563 (2025).

⁴⁵ See, most notably, Mark Zuckerberg, *Open Source AI Is the Path Forward*, Meta, Jul. 23, 2024, <https://perma.cc/9YVF-VRZS>.

⁴⁶ See, e.g., Anjali Gopal et al., *Will Releasing the Weights of Future Large Language Models Grant Widespread Access to Pandemic Agents?*, Oct. 25, 2023, [arXiv:2310.18233](https://arxiv.org/abs/2310.18233) [cs.AI]; Kelsey Piper, *Should We Make Our Most Powerful AI Models Open Source to All?*, VOX, Feb. 2, 2024, <https://perma.cc/7YF2-LSPW>; see generally Rebecca Ackermann, *The Future of Open Source Is Still Very Much in Flux*, MIT TECH. REV., Aug. 17, 2023, <https://perma.cc/8M9J-KQK7>; Bruce Schneier & Jim Waldo, *Big Tech Isn’t Prepared for A.I.’s Next Chapter*, SLATE, May 30, 2023, <https://perma.cc/M4XM-7LZC>.

relating to trade secrecy.⁴⁷ So, for example, the copyright industries have insisted that any media player licensed to render copyrighted content needs to incorporate proprietary technology that ensures an appropriately robust implementation of specific, secret access protocols. In litigation under provisions of the Copyright Act that prohibit trafficking in technologies for gaining unauthorized access to copyrighted works, open source advocates have failed to carve out shelter for the development of fully open source media players.⁴⁸ Licensed media players developed for open source environments exist, but they are not open in all respects.

Similarly, assessing whether any given AI system is “open” or “closed” requires a complex, multi-part inquiry that does not end with models and weights.⁴⁹ Debates about the benefits and risks of “open source AI” have largely ignored datasets and datastructures.⁵⁰ As we have just seen, however, models and weights are far from the only factors affecting the reliability of AI outputs and the potential for AI-caused mischief.

⁴⁷ On the role of ideology within open source software development, see GABRIELLA COLEMAN, *CODING FREEDOM: THE ETHICS AND AESTHETICS OF HACKING* (2012); CHRISTOPHER KELTY, *TWO BITS: THE CULTURAL SIGNIFICANCE OF FREE SOFTWARE* (2008).

⁴⁸ See *Universal City Studios, Inc. v. Reimerdes*, 111 F. Supp. 2d 294, 310-11 (S.D.N.Y. 2000), *aff'd*, *Universal City Studios, Inc. v. Corley*, 273 F.3d 429 (2d Cir. 2001); *RealNetworks, Inc. v. DVD Copy Control Ass’n, Inc.*, 641 F. Supp. 2d 913 (N.D. Cal. 2009).

⁴⁹ See generally Emily Black, Rakshit Naidu, Kit T. Rodolfa, Daniel E. Ho, & Hoda Heidari, *Toward Operationalizing Pipeline-Aware ML Fairness: A Research Agenda for Developing Practical Guidelines and Tools*, in *PROC. 3D ACM CONF. EQUITY & ACCESS IN ALGORITHMS, MECHANISMS, & OPTIMIZATION* (2023), <https://doi.org/10.1145/3617694.3623259>; David Gray Widder, Meredith Whittaker, & Sarah Myers West, *Why ‘Open’ AI Systems Are Actually Closed, and Why This Matters*, 635 *NATURE* 827 (2025), <https://doi.org/10.1038/s41586-024-08141-1>.

⁵⁰ See Black, et al., *supra* note 49; Rishi Bommasani, et al., *On the Opportunities and Risks of Foundation Models*, Jul. 12, 2022, arXiv:2108.07258 [cs.LG], at 8-9, 101-104; Francisco Eiras, et al., *Position: Near to Mid-Term Risks and Opportunities of Open-Source Generative AI*, May 24, 2024, arXiv:2404.17047 [cs.LG], at 4, 6-8; Widder, et al., *supra* note 49; *EleutherAI NTIA RFC Submission*, ELEUTHERAI, Mar. 27, 2024, <https://www.regulations.gov/comment/NTIA-2023-0009-0291>; *About ML Commons and Its Interest In This Request for Comment*, MLCOMMONS, Mar. 26, 2024, <https://www.regulations.gov/comment/NTIA-2023-0009-0166>.

Currently, market-leading developers of generative AI systems—including Meta—treat details about their datasets and data curation practices as proprietary.⁵¹ If openness is important for good AI governance—a question to which I return in Part III.B—it needs to extend to AI datastructures.

2. Alignment

A second widely endorsed approach to AI accountability involves fine-tuning AI models and/or adjusting model weights to produce better outputs.⁵² Although the regulatory discussions about AI in the U.S. and Europe differ in many ways, they have (until very recently) shared a common focus on fine-tuning to reflect human and more specifically Western values.⁵³ The U.S. is now clearly departing from that path as a formal matter, though private efforts may still continue.⁵⁴ But to the extent that alignment efforts focus on post hoc adjustments to models and/or weights rather than on systemic decisions about datastructuring, they are unlikely to get to the roots of generative AI's dysfunctions.

⁵¹ Domen Vake, et al., *Is Open Source the Future of AI? A Data-Driven Approach*, 15 APPLIED SCIS. 2790 (2025); NTIA Request for Comment Meta Response, META, Mar. 2024, <https://perma.cc/2HUR-DTXW> (offering justifications for dataset secrecy); see also Kate Knibbs, *Meta Secretly Trained Its AI on a Notorious Piracy Database, Newly Unredacted Court Docs Reveal*, WIRED, Jan. 9, 2025, <https://perma.cc/3QUW-D2F4>.

⁵² See generally Paul Ohm, *Focusing on Fine-Tuning: Understanding the Four Pathways for Shaping Generative AI*, 25 COLUM. SCI. TECH. L. REV. 214 (2024).

⁵³ See Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence and amending Regulations (Artificial Intelligence Act), OJ L 2024/1689, recital 110; Exec. Order No. 14,110, 88 FED. REG. 75191 (Oct. 30, 2023), rescinded by Exec. Order No. 14,179, 90 FED. REG. 8741 (Jan. 23, 2025); Vicki L. Birchfield, *From Roadmap to Regulation: Will There Be a Transatlantic Approach to Governing Artificial Intelligence?*, 46 J. EUR. INTEGRATION 1053 (2024); Manuel Wörsdörfer, *Biden's Executive Order on AI and the E.U.'s AI Act: A Comparative Computer-Ethical Analysis*, 37 PHIL. & TECH. 74 (2024).

⁵⁴ Exec. Order No. 14,179, 90 FED. REG. 8741 (Jan. 23, 2025); J.D. Vance, *Remarks by the Vice President at the Artificial Intelligence Action Summit in Paris*, The American Presidency Project, <https://perma.cc/SP4T-67VP>.

The regulatory discussions about AI training for alignment have important parallels in the debates about content moderation at scale. The “content moderation” label encompasses a wide range of activities, from identifying and removing harmful content (such as pseudoscience that threatens public health and calls for racially targeted, ethnonationalist, or insurrectionist violence) to detecting and countering so-called coordinated inauthentic activity (i.e., automated networks that post, like, and forward information in the service of covert influence campaigns). Generally speaking, all such activities are widely agreed to be very difficult to do perfectly (or even well) and to raise troubling risks of state and/or private censorship.⁵⁵ As I have explained elsewhere, however, the online content “climate” is not effectively regulable using post hoc interventions that focus on content but do not disturb underlying data harvesting and algorithmic optimization processes.⁵⁶ In important ways, the failures of content moderation at scale flow from widespread and systematic refusal to reckon more comprehensively with the deeper effects of datastructuring for content *immoderation* at scale.

An uncanny mirror image of content moderation’s failures has begun to manifest in the gradually coalescing Chinese approach to AI regulation, which relies in part on content moderation and in part on identifying and incentivizing the creation of training datasets that are deemed to be trustworthy because they align with “core socialist values.”⁵⁷ As a Chinese technology law expert explained to me, this approach reflects the perception that in Western discourses about AI regulation, “alignment” refers to efforts to institute cultural hegemony over AI outputs.⁵⁸ The focus on approved datasets is a strategy both for instituting an approach to alignment that advances the party-state’s worldview and for shifting the alignment project

⁵⁵ See, e.g., Eric Goldman & Jess Miers, *Why Can’t Internet Companies Stop Awful Content?*, ARS TECHNICA, Nov. 27, 2019, <https://perma.cc/2GVW-B6CS>; see generally Cohen, *Infrastructuring*, *supra* note 11, at 9.

⁵⁶ Cohen, *Infrastructuring*, *supra* note 11, at 28.

⁵⁷ Matt Sheehan, *China’s AI Regulations and How They Get Made*, CARNEGIE ENDOWMENT FOR INT’L PEACE, July 10, 2023, <https://perma.cc/Y975-9W9C>.

⁵⁸ I thank Haolin Li for educating me on this point.

to earlier stages of the AI development lifecycle. The party-state has also supported efforts to develop so-called “ambient content moderation,” with generative AI at its core, that does not simply flag forbidden content but rather uses the datastructures surrounding content to monitor and manage the affective “temperature” of social media ecosystems.⁵⁹ Such efforts, however, encounter what Eddie Yang and Margaret Roberts call the “authoritarian data problem”: models that have not been trained on the full range of expression available in more open information systems can fail to recognize prohibited content, making it necessary to import Western data for some purposes and uses.⁶⁰

As this description suggests, the ongoing saga of China’s AI regulations holds two important lessons for other countries engaged in the project of designing AI governance institutions. First, “alignment” is a deceptively innocuous way of referring to the seemingly ineradicable need to expend enormous amounts of effort on training the garbage out of the model (as discussed in Part III.B, below, at the further cost of enormous energy consumption) rather than on paying closer attention to how the garbage got there in the first place. Second, for that project to succeed as a non-authoritarian project, it will be important to disentangle the idea of content control from the idea of data quality and to focus on the ways that the behavioral datastructures we currently have overweight particular kinds

⁵⁹ Luzhou Li & Kui Zhou, *When Content Moderation is Not About Content: How Chinese Social Media Platforms Moderate Content and Why it Matters*, 27 NEW MEDIA & SOC’Y 6139 (2024), <https://doi.org/10.1177/14614448241263933>. China is, of course, not the only place where AI and content moderation are converging, and if recent controversies about matters ranging from Google’s relabeling of the Gulf of Mexico to misinformation-related training parameters for the Musk-developed GrokAI are any indication, the U.S. may not be as far behind in the march toward state-sanctioned, ambient censorship as we like to think. Dan Milmo, *Google Maps Will Rename Gulf of Mexico as Gulf of America in US*, THE GUARDIAN, Jan. 28, 2025, <https://perma.cc/NHB8-JLGF>; Effie Webb, *Grok Briefly Censored Criticism of Musk and Trump. It Was Blamed on a New Hire Who Hadn’t “Fully Absorbed” the Startup’s Culture*, BUS. INSIDER, Feb. 24, 2025, <https://perma.cc/4U8Q-ZZDJ>.

⁶⁰ Eddie Yang & Margaret E. Roberts, *The Authoritarian Data Problem*, 34 J. DEMOC. 141, 142 (2023).

of signals and underweight others. Put another way, disentangling the idea of content control from the idea of data quality is an important prerequisite for democratic datastructuring—a point to which I return in Part III, below. Some of the enthusiasm for alignment work is undoubtedly genuine—and, under any functioning regime for AI accountability, model training and fine-tuning will have important roles to play. But alignment should not be the only—or even the central—focus of institutional design efforts.

3. Risk assessment

Another influential family of approaches to accountability in AI development involves risk assessment and mitigation. There is considerable variation among risk-focused approaches. Some involve *ex ante* categorization, whereas others primarily involve *post hoc* testing and/or benchmarking.⁶¹ For generative AI models, *post hoc* testing may focus on the models or on narrower sets of use cases.⁶² As of this writing, there are no commonly agreed standards for any of these activities. Generally speaking, developing standards for risk assessment of generative AI systems is a very good idea. Risk assessment focused on models and use cases, however, is not the same as risk assessment focused on datastructures.

Risk assessment relating to human economic and technological development is a complex and hotly debated topic for reasons well beyond the scope of this article. In brief: affected industries and their investors both have strong incentives to ignore or minimize the longer-term risks their

⁶¹ See Artificial Intelligence Act, *supra* note 53, recitals, 53, 165, arts. 5, 6, Annex III, XI; Roberts, et al., *supra* note 60; *AI Risk Management Framework (AI RMF 1.0)*, Nat'l Inst. of Standards & Tech., Jan. 2023, <https://www.nist.gov/itl/ai-risk-management-framework/ai-risk-management-framework-engage>. *But see* Leila Abboud & Melissa Heikkila, *US and UK Refuse to Sign AI Safety Declaration at Summit*, ARS TECHNICA, Feb. 11, 2025, <https://arstechnica.com/ai/2025/02/us-and-uk-refuse-to-sign-ai-safety-declaration-at-summit/>.

⁶² *Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile*, Nat'l Inst. of Standards & Tech., July 2024, <https://perma.cc/3VBM-6MMF>; *see also* Lily Newman, *The US Government Wants You—Yes, You—to Hunt Down Generative AI Flaws*, WIRED, Aug. 21, 2024, <https://perma.cc/CVP6-3LFB>.

activities create. Devolving testing to third-party intermediaries or to wholly self-regulatory processes, as is routine within a wide variety of regulatory regimes in the U.S., creates additional moral hazard problems.⁶³ Last but not least, prevailing methods for predicting and quantifying risks tend to overweight the past and therefore to discount certain types of systemic risk that are increasingly important and severe. Recent examples include financial crashes precipitated in part by increasingly complex synthetic financial instruments and increasingly sudden and severe hurricanes precipitated by warming oceans.⁶⁴ These “black swan” risks now also include those created via the pervasive interventions of powerful generative AI systems.

Bracketing those important (and often interrelated) disputes about moral hazard and methodology, my point here is more basic: testing the outputs of AI development processes is, at best, an indirect way of addressing datastructuring issues. Sorting use cases into risk tiers and imposing direct risk mitigation duties, as the European AI Act does, still does not shed enough light on the distinctive upstream risks datastructuring introduces. The AI Act links generative AI models trained at scale to systemic risk and prescribes special safeguards for their training, assessment, and distribution. It further requires documentation of training activities, including “information on the data used for training, testing, and

⁶³ See Kenneth Bamberger, *Technologies of Compliance: Risk and Regulation in a Digital Age*, 88 TEX. L. REV. 669 (2010); Kenneth Bamberger, *Regulation as Delegation: Private Firms, Decisionmaking, and Accountability in the Administrative State*, 56 DUKE L.J. 377 (2006).

⁶⁴ See, e.g., Hilary J. Allen, *Regulatory Managerialism and Inaction: A Case Study of Bank Regulation and Climate Change*, 86 L. & CONTEMP. PROBS. 71 (2023); NASSIM NICHOLAS TALEB, *THE BLACK SWAN: THE IMPACT OF THE HIGHLY IMPROBABLE* (2010); Intergovernmental Panel on Climate Change, *Weather and Extreme Climate Events in a Changing Climate*, in IPCC SIXTH ASSESSMENT REPORT 1513 (2021), <https://perma.cc/HN2B-ZNCU>; See generally Jose Luis Bermudez & Michael Pardo, *Risk, Uncertainty, and “Super-Risk,”* 29 NOTRE DAME J. L. ETHICS & PUB. POL’Y 471 (2015).

validation” of all “general-purpose” AI models.⁶⁵ But generative AI’s effects on societal processes of knowledge production cannot be so easily firewalled. Without development of risk assessment methodologies addressed more specifically to the non-content elements of datastructures, policymakers can neither account for nor begin to address the effects of datastructuring for engagement and extraction on the capabilities and behaviors of AI systems.

4. Acceleration

A fourth approach to AI governance, increasingly ascendant in the US, is known by the moniker “accelerationism.” Accelerationist thinking about AI holds that radical breakthroughs in AI-driven capabilities—leading ultimately to so-called artificial general intelligence (AGI) that exceeds human capabilities—are hovering just beyond our reach.⁶⁶ Reaching them requires more inputs for AI development—more data, more compute, and more energy.⁶⁷ As to outputs, accelerationists argue that the dysfunctions of generative AI systems are mere hiccups, which increasing the power of generative AI systems will inevitably correct.⁶⁸

The infrastructure lens suggests, however, that both halves of the accelerationist narrative are wrong. Begin with inputs.

⁶⁵ Artificial Intelligence Act, *supra* note 53, arts. 51, 53, 55, Annex XI; *see also* Third Draft of the General-Purpose AI Code of Practice, Transparency Section, Mar. 11, 2025, <https://perma.cc/59AB-QFK2>.

⁶⁶ *See, e.g.*, RAY KURZWEIL, *THE SINGULARITY IS NEAR* (2005); RAY KURZWEIL, *THE SINGULARITY IS NEARER: WHEN WE MERGE WITH AI* (2024); Dario Amodei, *Machines of Loving Grace: How AI Could Transform the World for the Better*, Oct. 2024, <https://perma.cc/8KPO-KG6K>.

⁶⁷ Jeffrey Dastin, *OpenAI CEO Altman Says at Davos Future AI Depends on Energy Breakthrough*, REUTERS, Jan. 16, 2024, <https://perma.cc/L6T3-FN3X>; Samuel Hammond, *Why AGI Is Closer than You Think*, SECOND BEST, Sept. 22, 2023, <https://perma.cc/R8W9-46LA>; *see also* Veronika Samborska, *Scaling up: How Increasing Inputs Has Made Artificial Intelligence More Capable*, OUR WORLD IN DATA, Jan. 19, 2025, <https://perma.cc/C2TR-2R75>. I consider AI’s energy consumption briefly in Part III.B, below.

⁶⁸ Marc Andreessen, *The Techno-Optimist Manifesto*, Andreessen Horowitz, Oct. 16, 2023, <https://perma.cc/8C8H-ZQ5Z>; Kevin Roose, *This A.I. Subculture’s Motto: Go, Go, Go*, N.Y. TIMES, Dec. 10, 2023, <https://perma.cc/8D6D-ULNQ>.

According to accelerationists, we are living in a time of data scarcity. After training ever more powerful generative AI models has consumed all of the available data—an event that is projected to occur soon—the AI development race will hit a wall.⁶⁹ (Whether or not such predictions are right, they tend to confirm that the market-leading models are being trained on the behavioral datastructures that Part II.A described.) Accelerationists typically propose to solve the coming data deficit problem by generating synthetic data. That process, however, is likely to involve automation and patterning based on what AI datastructures already contain.⁷⁰ Even more troublingly, many researchers have concluded that reliance on synthetic data can eventually lead generative models to collapse, producing results that are increasingly incoherent and nonsensical.⁷¹ Although the extent of this risk is contested, it seems to be common ground that some amount of continued direct human involvement in behavioral data generation at scale is an important determinant of data quality.

Next, recall that the datastructures that supply AI's inputs both afford and constrain AI's outputs. As long as issues relating to AI datastructure quality remain largely overlooked, proposals to “solve” the problems of simulation and sociopathy by making the models more powerful and feeding them more

⁶⁹ Ross Andersen, *What Happens When AI Has Read Everything?*, THE ATLANTIC, Jan. 18, 2023, <https://perma.cc/629S-8EZB>; Pablo Villalobos, Anson Ho, Jamie Sevilla, Tamay Besiroglu, Lennart Heim, & Marius Hobbhahn, *Will We Run Out of Data? Limits of LLM Scaling Based on Human-Generated Data*, Oct. 26, 2022, arXiv:2211.04325 [cs.ML].

⁷⁰ Ruibo Liu, et al., *Best Practices and Lessons Learned on Synthetic Data for Language Models*, arXiv:2404.07503v1 [cs.CL], Apr. 11, 2024; see also Daniel Susser, et al., *Synthetic Health Data: Real Ethical Promise and Peril*, 54 HASTINGS CTR. REP. 8, (Sept.-Oct. 2024), doi: 10.1002/hast.4911.

⁷¹ Aatish Bhatia, *When A.I.'s Output Is a Threat to A.I. Itself*, N.Y. TIMES, Aug. 26, 2024, <https://perma.cc/6R3M-Q3C7>; Matteo Wong, *AI Is an Existential Threat to Itself*, THE ATLANTIC, Jun. 21, 2023, <https://perma.cc/6JKS-NNS5>; Ilya Shumailov et al., *AI Models Collapse When Trained on Recursively Generated Data*, 631 NATURE 755 (2024), <https://doi.org/10.1038/s41586-024-07566-y>. But see Rylan Schaeffer, Joshua Kazdan, Alvan Caleb Arulandu, & Sanmi Koyejo, *Position: Model Collapse Does Not Mean What You Think*, arXiv:2503.03150v2 [cs.LG], Mar. 18, 2025 (arguing that the “model collapse” label conflates different types of scenarios, some of which are “readily avoidable”).

data—whether extracted directly from human data subjects or synthesized based on the properties of existing datasets—are virtually certain to make those problems worse. That eventuality must be weighed against rosier speculations about what AGI will do. AGI developed within current platformized datastructures might cure disease and reverse climate change, but it might also throw support to know-nothing cults and offer up prescriptions for selective population culling. Arguably, the looming failure of what has amounted to an attempt to brute-force AGI using behavioral data extracted from humans points to a somewhat different lesson: Perhaps it is the nature of the human involvement that matters. Attaining new thresholds in AI capability may not require more data and less human involvement; it may require better datastructuring and different human involvement—possibilities to which I return in Part III.B.

III. The Promise and Limits of “Public Utility” Regulation

The dysfunctions of AI datastructures, which are infrastructural and manifest at scale, suggest the need for correspondingly large-scale projects of redesign and reconstruction. In some ways, the challenge evokes the mid-twentieth century “public works” projects that marshaled government and private inputs to build highways and electrical grids. Scholarship in law and political economy has identified the public utility model that emerged to govern the latter group of projects as centrally relevant to the project of governing both AI development and digital platform power more broadly.⁷² This Part considers that suggestion and concludes that, although the public utility model is relevant to some pieces of the AI governance puzzle, its lessons are more complex than its advocates acknowledge.

⁷² Tejas Narechania & Ganesh Sitaraman, *An Antimonopoly Approach to Governing Artificial Intelligence*, 43 YALE L. & POL’Y REV. 95 (2024); K. Sabeel Rahman, *The New Utilities: Private Power, Social Infrastructure, and the Revival of the Public Utility Concept*, 39 CARDOZO L. REV. 1621 (2018); MORGAN RICKS, GANESH SITARAMAN, SHELLEY WELTON, & LEV MENAND, NETWORKS, PLATFORMS, AND UTILITIES: LAW AND POLICY (2023).

In a nutshell, the problem is that the public utility model addresses only some of the considerations that the infrastructure lens highlights. It is highly attuned to questions about access to infrastructures and their outputs and to the division of labor between government and market actors in coordinating infrastructure provisioning. At the same time, it is relatively insensitive to questions about infrastructure configuration and input sourcing, and it has proved poorly designed for handling problems of infrastructure redesign and transition. It therefore represents a cramped frame for managing the full array of policy considerations that the emerging AI era now requires us to consider.

Section A reviews some of the gaps and ambiguities that exist in the public utility framework and considers their implications for the project of AI governance writ large, and Section B distills some important lessons for the more specific project of governing AI datastructures.

A. *“You Keep Using that Word”: What’s [a] Public Utility, Anyway?*

The growing chorus of support for public utility thinking about contemporary regulatory problems involving networked digital technologies often passes too quickly over questions about the extent to which preexisting models of public utility regulation can simply be transposed into new contexts.⁷³ Assumptions about regulatory portability, in turn, point to a more fundamental underlying definitional muddiness; people do not always seem to be on the same page about what the term public utility actually signifies. I begin with the definitional question and work forward to the ways in which its various possible answers throw the portability question into sharper relief.

One possible way to understand “public utility” is as a noun referring to a resource with particular characteristics that

⁷³ One notable exception is the Rahman & Teachout call to ban behavioral advertising on digital platforms, which I note in Part III.B, below. K. Sabeel Rahman & Zephyr Teachout, *From Private Bads to Public Goods: Adapting Public Utility Regulation for Informational Infrastructure*, KNIGHT FIRST AMEND. INST., Feb. 4, 2020, <https://perma.cc/JES4-E6SW>.

require it to be managed in a particular way. Some scholars and scholarly traditions designate natural monopoly resources as (potential) public utilities.⁷⁴ Another increasingly influential definition identifies “essential, general-purpose services” that are provided to users “more or less continuously.”⁷⁵ In practice, however, neither definition is sensitive to the infrastructural ramifications of different methods for generating and providing resource-related services at scale. The terminology of natural monopolies and essential facilities, which foregrounds consumable outputs while submerging the other kinds of choices entailed in infrastructuring work, seems at least partly to blame for this oversight.⁷⁶ So, for example, although an electrical grid serving a particular region is a natural monopoly in the formal sense, it might receive inputs from a variety of sources (e.g., coal, nuclear, hydropower, wind, solar). Similarly, a municipal water system might rely on a river, underground aquifers, desalinated water, or a watershed system hundreds of miles away. In both examples, choices about inputs and input sourcing have enormous public policy significance.

Another way to understand “public utility” is as an adjective referring to a particular type of regulatory regime created for certain types of infrastructural services (that may satisfy one or both of the criteria listed in the previous paragraph or, alternatively, may be defined more broadly).⁷⁷ Over time, however, fragmentation of the underlying resource categories has produced some predictable incoherence in the

⁷⁴ See, e.g., William Boyd, *Just Price, Public Utility, and the Long History of Economic Regulation in America*, 35 YALE J. REG. 721 (2018) [hereinafter Boyd, *Just Price*]; Herbert Hovenkamp, *Technology, Politics and Regulated Monopoly: An American Historical Perspective*, 62 TEX. L. REV. 1263 (1984).

⁷⁵ RICKS, ET AL., *supra* note 72, at 7; *But see* Josh Macey & Genevieve Lakier, *What Are Networks, Platforms, and Utilities, and What Should We Do with Them?*, Notice and Comment, YALE J. REG. ONLINE, Jan. 24, 2023, on this definition’s lack of analytical heft.

⁷⁶ *Cf.* Boyd, *Just Price*, *supra* note 74, at 770-771 (arguing that both the natural monopoly idea and the subsequent rise of neoliberal deregulatory ideology have worked to obscure an older conception of economic justice embedded in the public utility model).

⁷⁷ See generally K. Sabeel Rahman, *Infrastructural Regulation and the New Utilities*, 35 YALE J. REG. 911 (2018).

public utility regulatory toolkit. The twentieth-century public utility toolkit relied largely on rate regulation and access mandates to ensure both adequate returns to service providers and fair treatment of service users.⁷⁸ Certainly, both goals remain extremely important, but other considerations that are equally important—if not more so—have been added to the list. For example, in addition to the supply-side considerations described in the previous paragraph, an electrical grid’s outputs might be supplemented in various ways (e.g., installation of solar panels, purchases of alternative energy from private providers), each of which has different ecological and sustainability implications.⁷⁹ For certain types of essential services, such as electricity and water, regulation of demand also has become essential for meeting increasingly pressing sustainability goals, and such regulation also might be achieved in a number of different ways.⁸⁰

A third way to understand “public utility” is as shorthand for complex sets of decisions about infrastructure provisioning to supply “utility” for the public—and, more specifically, for a regulatory regime that attempts to manage the complexities by providing appropriate incentives to market actors. Processes for provisioning infrastructure, however, are more than just exercises in government-market coordination. Scholarship in energy law and policy has highlighted the unique arrangements for government-coordinated market provisioning that produced the first energy utilities, and it has also shown how more recent problems relating to electric grid deterioration and water system depletion now challenge conventional wisdom about older public provisioning strategies for reasons

⁷⁸ See generally William Boyd, *Public Utility and the Low Carbon Future*, 61 UCLA L. REV. 1614 (2014) [hereinafter Boyd, *Low Carbon Future*].

⁷⁹ See *id.* at 1669; William Boyd & Ann. E. Carlson, *Accidents of Federalism: Ratemaking and Policy Innovation in Public Utility Law*, 63 UCLA L. REV. 810, at 890 (2016).

⁸⁰ See generally Subhasis Panda, et al., *A Comprehensive Review on Demand Side Management and Market Design for Renewable Energy Support and Integration*, 10 ENERGY REP. 2228 (2023); A.H. Alias, C.A. Boyla, & S. Hassim, *Water Demand Management: A Review on the Mechanisms to Reduce Water Demand and Consumption*, 8 INT’L J. CIV. ENG. 554 (2017).

relating to both input sourcing and system configuration.⁸¹ As the country and the world confront an increasingly urgent climate crisis, it is insufficient—and in the case of water, soon to be impossible—to simply to rebuild old systems. Meanwhile, some of the loudest voices resisting redesign for decarbonization and conservation have come from incumbents protected by existing regulatory arrangements designed with the principal goal of ensuring a steady stream of cheap inputs.⁸² It has become necessary both to dismantle such arrangements and to reimagine how regimes for infrastructure provisioning and governance might become more systematically attentive to the sustainability implications of both inputs and outputs.

A different way of putting all of these points is that the portmanteau “public utility” elides some important questions about how “utility” ought to be understood once the costs of different choices about input sourcing and configuration are considered, and that within the public utility regulatory model more specifically, that is so by design. The public utility model is a political economy framework for infrastructure provisioning and operation that incorporates robust government intervention to solve problems relating to market entry, competition, and control over access. At the same time, at least as a baseline, it leaves provisioning problems related to input sourcing and infrastructure design for market actors themselves to solve. When incumbents create new bottlenecks or resist the kinds of innovation needed to address pressing social problems, scholars and policymakers wanting to supplement the public utility toolkit tend to lean into proposals for restructuring rates and/or markets to recalibrate incentives, which are forms of intervention that the public utility frame encourages.⁸³

⁸¹ See generally William Boyd, *Decommodifying Electricity*, 97 S. CAL. L. REV. 937 (2024); Alexandra Klass, Joshua Macey, Shelley Welton, & Hannah Wiseman, *The Key to Electric Grid Reliability: Modernizing Governance*, KLEINMAN CTR. FOR ENERGY POL’Y, Mar. 2024, <https://perma.cc/J2HS-8NDZ>; Newsha K. Ajami, Barton H. Thompson Jr., David G. Victor, *The Path to Water Innovation*, THE HAMILTON PROJECT, Oct. 14, 2014, <https://perma.cc/8W9E-ERXD>.

⁸² See sources cited in the previous footnote.

⁸³ See, e.g., Boyd, *Low Carbon Future*, *supra* note 78; Boyd, *Decommodifying Electricity*, *supra* note 81; Gabriel Chan & Alexandra B.

For exactly this reason, moreover, other regulatory regimes that sit outside the frame of the public utility model have long played vital roles in energy and water policy. Two in particular—environmental regulation and critical infrastructure regulation—are squarely relevant to the current debates about AI policy. Environmental regulation focuses on the ecosystem-level consequences of decisions about the production and consumption of both private and public goods. So, for example, both the Clean Air Act and the Clean Water Act impose strict limits on pollutant discharges from power plants.⁸⁴ Critical infrastructure regulation, meanwhile, is directly concerned with safeguarding the resilience of systems that support social and economic life at scale.⁸⁵ As researchers and policymakers already recognize, attempting to map the possible futures of AI regulation and governance without considering the lessons of these regimes would be a mistake.

Like electrical grids and water systems, AI systems present complex, multimodal provisioning and governance challenges. Proposals to use the emerging antimonopoly toolkit to ensure appropriately widespread access to upstream resource layers (such as data centers and compute), thereby fostering important types of downstream competition (regarding models

Klass, *Regulating for Energy Justice*, 97 N.Y.U. L. REV. 1426 (2022); Daniel E. Walters, *Lumpy Social Goods in Energy Decarbonization: Why We Need More than Just Markets for the Clean Energy Transition*, 93 COLO. L. REV. 541 (2022).

⁸⁴ See generally Richard K. Lattanzio, *Clean Air Act: A Summary of the Act and Its Major Requirements*, CONG. RES. SERV., Sept. 13, 2022, <https://perma.cc/3GWM-44V6>; Laura Gatz, *Clean Water Act: A Summary of the Law*, Cong. Res. Serv., Oct. 18, 2016, <https://perma.cc/4E9X-9SG4>; Statement of Laura Gatz, Hearing on The Clean Water Act at Fifty: Highlights and Lessons Learned from a Half Century of Transformative Legislation, Subcomm. On Water Resources & Env't., Comm. on Transp. & Infra., Sept. 20, 2022, <https://perma.cc/K9LJ-275A>.

⁸⁵ See generally *Critical Infrastructure Sectors*, Cybersecurity and Infrastructure Security Agency (last visited May 4, 2025), <https://perma.cc/4RWS-5UVR>; see also Rebecca Slayton & Aaron Clark-Ginsberg, *Beyond Regulatory Capture: Coproducing Expertise for Critical Infrastructure Protection*, 12 REG. & GOV. 115 (2018).

and uses), are surely good ideas.⁸⁶ The infrastructure lens also teaches, however, that it is important to pay attention to the ways policy choices at different layers intersect. So, for example, certain chip configurations both presume and require access to large arrays of compute and large amounts of energy, while others do not.⁸⁷ Increasing access to compute to foster competition at the model layer without carefully considering the systemic implications of growing dependency on large, centralized compute arrays would be unwise.

Data is a non-rival and non-excludable resource, so the problems of securing an adequate supply of training data and adequate access to training datasets might be framed as conventional public goods problems addressable via a combination of licensing requirements, access requirements, and structural separations.⁸⁸ As this article has shown, however, data—like electricity and water—can be generated, consumed, and replenished in ways that cause significant systemic degradation. And datasets and datastructuring exert powerful effects on the generative AI development process that are distinct from and antecedent to those associated with models. Subsuming questions about datasets and datastructuring into questions about model development therefore represents a significant category error.⁸⁹

A. *“You’re Gonna Need a Bigger Boat”: Elements of a Framework for AI Datastructure Governance*

The ongoing struggles over electric grid decarbonization and water conservation that have overtaxed the public utility model should inform the challenge of AI datastructure governance in three ways. First, they underscore the need for a regulatory framework that contends more directly and successfully with the complex path dependencies embedded in the design and ongoing maintenance of infrastructures that

⁸⁶ See, e.g., Narechania & Sitaraman, *supra* note 72; Jai Vipra & Sarah Myers West, *Computational Power and AI*, AI NOW INST. (2023), <https://perma.cc/GDR9-PQ6J>.

⁸⁷ Cade Metz, et al., *How A.I. Is Changing the Way the World Builds Computers*, N.Y. TIMES, Mar. 16, 2023, <https://perma.cc/JCE3-3PC7>.

⁸⁸ Narechania & Sitaraman, *supra* note 72.

⁸⁹ See *supra* Part II.B; see also Elettra Bietti, *Data Is Infrastructure*, 26 THEOR. INQ. L. 55 (2025); Narechania & Sitaraman, *supra* note 72.

supply essential public goods. Decarbonization and water conservation are difficult partly because current arrangements for providing electricity and water are embedded at scale within constructed arrangements that are both transparent to use and behaviorally habituated. So too with AI datastructures: Platformized ecosystems mobilize vast, dispersed data supply chains that cross conventional organizational boundaries—and that already have resisted many other kinds of regulatory oversight and accountability.⁹⁰

Second, the struggles over energy and water sustainability underscore the need to contend with the power of entrenched incumbents whose incentives and behaviors have produced systemic pathologies. Due in large part to the winner-take-all characteristics of the platformized information economy, the very large datasets and very large amounts of compute are concentrated in the hands of a few powerful commercial actors (and under the de facto control of their oligarchic leaders).⁹¹ A functioning regime for AI datastructure governance will need to contend with incumbent power.

Last but not least, the struggles over energy and water sustainability remind us to pay attention to values that tend to sit outside the frame of regulatory paradigms concerned principally with access and competition. As Part II explained, AI datastructures directly implicate core goals and challenges of democratic governance. Although the problem has not been a principal focus of this article, AI systems also directly implicate core goals and challenges of sustainable energy regulation. The development, training, fine-tuning, and use of generative AI currently consume mind-boggling and rapidly mushrooming amounts of energy, and continuing along the current trajectory threatens sustainable development goals on a planetary scale.⁹² Discussions about AI datastructure

⁹⁰ Jennifer Cobbe, Michael Veale, & Jatinder Singh, *Understanding Accountability in Algorithmic Supply Chains*, FACCT '23 (2023), <https://doi.org/10.1145/3593013.3594073>.

⁹¹ See generally Rikap, *supra* note 18; Vipra & West, *supra* note 86; Cohen, *Oligarchy*, *supra* note 44.

⁹² Bender, et al., *supra* note 21; KATE CRAWFORD, *ATLAS OF AI: POWER, POLITICS, AND THE PLANETARY COSTS OF ARTIFICIAL INTELLIGENCE* (2021); International Energy Agency, *Electricity 2025: Analysis and*

governance cannot focus narrowly on considerations of access and competition to the exclusion of questions about the kind of automated intelligence we want to build and the full tally of human and social costs.

Building from those lessons, this section identifies some options that policymakers who want to be serious about the potential costs and harms of AI-engendered simulation and sociopathy will need to explore. A warning to the reader: my goal is not to tie off loose ends but instead to open a Pandora's Box of disruptive and innovative regulatory options.⁹³

1. Open source revisited: Disclosing datastructuring

It is useful to begin with the idea of open access—and to recall that, as discussed above, openness means different things to the various players in generative AI development and in multisided platform ecosystems generally. Meaningful access to information about datasets and datastructures—for both market participants and regulators—is a necessary first step in structuring effective interventions.

Private initiatives have made significant progress toward understanding the various elements that meaningful disclosures would need to contain. In particular, the “Datasheets for Datasets” project undertaken by a group of highly regarded researchers has developed a detailed specification for dataset disclosure modeled on the datasheets that typically accompany electronic components, and the researchers report that some industry participants have begun using some elements of their proposed approach.⁹⁴ The project

Forecast to 2027, IEA.org, Feb. 2025, <https://perma.cc/96P9-MMQD>; see also Cristina Criddle & Stephanie Stacey, *Big Tech Data Center Buildouts Have Led to \$5.4 Billion in Public Health Costs*, ARS TECHNICA, Feb. 24, 2025, <https://perma.cc/X3Q4-5YVK>.

⁹³ On reclaiming disruption and innovation as legitimate and essential regulatory domains, see generally Jodi L. Short, *Regulatory Managerialism as Gaslighting Government*, 86 L. & CONTEMP. PROBS. 1 (2023); Julie E. Cohen, Nina-Simone Edwards, Meg Leta Jones, & Paul Ohm, *Designing Tools and Mechanisms for Regulatory Dynamism: Preliminary Concept Paper*, Redesigning the Governance Stack Project (Apr. 2025), <https://perma.cc/6NFR-PEXT/>.

⁹⁴ Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hal Daume III, & Kate Crawford, *Datasheets for Datasets*:

is oriented specifically toward bias detection, but it is easily adaptable to a broader range of uses.

As the history of financial market regulation demonstrates, however, relying on market participants or even third-party audit intermediaries to ensure adequate disclosure would be foolhardy.⁹⁵ For disclosures about datastructures to be trustworthy and adequate to the regulatory tasks at hand, they must be mandatory and enforced.⁹⁶ Note that although industry participants sometimes cite privacy considerations as militating against disclosure, a mandate to disclose data provenance and dataset structure need not entail full disclosure of dataset contents. There are many possible ways of structuring disclosure and provenance tracing regimes.⁹⁷

2. Alignment revisited: Preventing distortive datastructuring

Next, return to the problem of alignment—and recall that, as discussed above, it is both more accurate and more compatible with democratic values to reframe alignment projects around issues of datastructure design rather than

Documentation to Facilitate Communication between Dataset Creators and Consumers, 64 COMM'CS ACM 86 (2021), doi:10.1145/3458723; see also Shayne Longpre, Robert Mahari, Naana Obeng-Marnu, William Brannon, Tobin South, Jad Kabbara, & Sandy Pentland, *Data Authenticity, Consent, and Provenance for AI Are Broken: What Will It Take to Fix Them?*, AN MIT EXPLORATION OF GENERATIVE AI, Mar. 2024; <https://doi.org/10.21428/e4baedd9.a650f77d>; see generally Weixin Liang, Girmaw Abebe Tadesse, Daniel Ho, L. Fei-Fei, Matei Zaharia, Ce Zhang, & James Zou, *Advances, Challenges and Opportunities in Creating Data for Trustworthy AI*, 4 NATURE MACHINE INTELLIGENCE 669 (2022).

⁹⁵ See generally John C. Coffee, Jr., *Understanding Enron: "It's About the Gatekeepers, Stupid"*, 57 BUS. LAW. 1403 (2002); Thomas J. Fitzpatrick IV & Chris Sagers, *Faith-Based Financial Regulation: A Primer on Oversight of Credit Rating Agencies*, 61 ADMIN. L. REV. 557 (2009); Saule T. Omarova, *Bankers, Bureaucrats, and Guardians: Toward Tripartism in Financial Services Regulation*, 37 J. CORP. L. 621 (2012).

⁹⁶ See generally Julie E. Cohen, Brenda Dvoskin, Meg Leta Jones, Paul Ohm, & Smitha Krishna Prasad, *Regulatory Monitoring in the Information Economy: Preliminary Concept Paper*, Redesigning the Governance Stack Project (Jan. 2024), <https://perma.cc/6NFR-PEXT>.

⁹⁷ Cf. Christopher Morten, Gabriel Nicholas, & Salome Viljoen, *Researcher Access to Social Media Data: Lessons from Clinical Trial Data Sharing*, 38 BERKELEY TECH. L.J. 109 (2024).

issues of content control. Public utility regulation traditionally has approached what might be called alignment projects principally through the lenses of rate regulation and nondiscrimination, both of which are output focused. The types of interventions those lenses might suggest are unlikely to succeed. Consider, for example, the proposal by Sabeel Rahman and Zephyr Teachout to ban behavioral advertising, which can be understood as a nonprice intervention in the rate structure of the platformized information economy that alters the terms of arrangements between platforms and individual users of their services.⁹⁸ The proposal is a very good idea for reasons that relate to other goals a public-spirited regulator might identify. But simply banning the use of behavioral data in ad targeting would not address *collection* of behavioral data and would only indirectly address its use for AI training purposes (in the latter case, via possible trickle-down effects on the universe of tokens and styles derived from paid promotional content). And, as discussed in Part II.A, it is already well understood that post hoc nondiscrimination interventions can be stymied by bias embedded in datastructures.

A regulatory framework for AI governance needs to focus squarely on the path dependencies that AI datastructures introduce. That exercise can borrow from thinking in environmental and health law regarding chemicals that are deemed hazardous to human and/or ecosystem health in any amount. When such a finding is made, the chemicals need to be removed from supply and distribution chains—a process that begins where the chemicals are manufactured and introduced. The analysis in Part II.A suggests that supply chain restrictions on three aspects of AI datastructuring will be particularly important: behavioral data generated by and about patterns of engagement, data derived from paid promotional content, and data generated by and about coordinated inauthentic activity.⁹⁹

⁹⁸ Rahman & Teachout, *supra* note 73.

⁹⁹ The proposals described here would need to navigate the expansionist first amendment, which has come to function as a secular theology of sorts for many technology lawyers and their clients. Properly understood, the first amendment should not be a serious objection to economic regulation concerned with the way AI datastructures function to mediate flows of economic and social activity. See generally Dan L. Burk, *Asemic*

Additionally, the analysis in Part II.B suggests that deliberate generation and use of synthetic data to train generative AI should be carefully supervised and documented. Ensuring compliance with datastructuring mandates will require expansion and some reinvention of regulatory monitoring authority and capabilities, but that project is long overdue.¹⁰⁰

3. Risk mitigation revisited: Deterring wasteful datastructuring

Recall, next, that risk mitigation in AI governance will be effective only to the extent that it targets the datastructuring practices that create the largest systemic risks. And because so many of AI's dysfunctions relate to scale—scale in behavioral datastructuring, scale in power consumption, and scale as a matter of accelerationist ideology—thinking about regulatory carrots and sticks needs to scale up commensurately. The market-leading AI developers do not need incentives to grow larger and can well afford their outrageous power bills. Proposals to limit excessive data and energy use need to consider the enormous economic resources that the dominant tech platform firms and their oligarchic leaders already possess.

If taken seriously, the documentation and firewalling requirements described above would operate as taxes of sorts, but they need to be supplemented by more direct mandates addressing scale in datastructuring and underlying energy consumption. Regulatory approaches that are durably pro-social likely will need to involve advances in smaller-data learning (i.e. training on carefully curated datasets) coupled with innovation around techniques for minimizing compute. The DeepSeek AI system recently released by Chinese researchers appears to be an example of the latter, though perhaps not the former.¹⁰¹ Regulators working to craft policies

Defamation, or, the Death of the AI Speaker, 22 FIRST AMEND. L. REV. 189 (2024). But it is also necessary to bracket the question to prevent it from consuming the entire thought exercise I undertake here.

¹⁰⁰ See generally Cohen, et al., *supra* note 96.

¹⁰¹ Xiao Bi, et al., *DeepSeek LLM: Scaling Open-Source Language Models with Longtermism*, Jan. 5, 2024, <https://perma.cc/FR8X-4MEF>; Caiwen Chen, *How a Top Chinese AI Model Overcame US Sanctions*, MIT TECH. REV., Jan. 24, 2025, <https://perma.cc/WP3E-CVPF>; Gemma Conroy &

for democratic datastructuring should use federal grant-making power to encourage both kinds of innovation. They should also use federal procurement power to prohibit provisioning that creates or deepens government dependencies on AI resulting from distortive and/or wasteful datastructuring. Federal initiatives for greening energy production, now largely suspended by the current Administration, also are vital to the long-term prospects for deterring wasteful datastructuring.

Both kinds of efforts, moreover, will need to be crafted with careful attention to possible workarounds. The evolving stories of electricity and water regulation offer examples of the many different ways interventions intended to reduce consumption might be structured. Such examples include congestion pricing, congestion caps, fines for excess consumption, quality of service pricing, and information feedback strategies of various sorts, including some that attempt to force attention to sustainable input sourcing. To put it mildly, such efforts do not always succeed. In particular, when certain consumers and certain uses are relatively insensitive to price, strategies based on congestion pricing, fines, and quality of service pricing do not reliably meet their targets, and even when they do, they may produce distributional effects that do not (or should not) accord with public priorities.¹⁰² Efforts (or lack thereof) to control Bitcoin mining and its insatiable demands for energy and compute supply a related and perhaps even more instructive cautionary tale: Those determined to mine Bitcoin have found workarounds by crossing borders and/or using botnets to commandeer resources owned by others.¹⁰³ To counteract the seeming imperative to generate, process, and

Smriti Mallapaty, *How China Created AI Model DeepSeek and Shocked the World*, 638 NATURE 300 (2025).

¹⁰² See generally Thomas Wiedmann, Manfred Lenzen, Lorenz T. Keyser, & Julia K. Steinberger, *Scientists' Warning on Affluence*, 11 NATURE COMM'NS 3017 (2020); Elisa Savelli, et al., *Urban Water Crises Driven by Elites' Unsustainable Consumption*, 6 NATURE SUSTAINABILITY 929 (2023); Florian G. Kaiser & Jan Urban, *Wealth as an Obstacle and a Support for Environmental Protection*, 100 J. ENVTL. PSYCH. 102449 (2024).

¹⁰³ Peter Guest, *Bitcoin Mining Was Booming in Kazakhstan. Then It Was Gone.*, MIT TECH. REV., Jan. 12, 2023, <https://perma.cc/C5ES-ES35>; Opeyemi Sule, *US Bitcoin Miners Move Old Equipment Overseas - Here's Why*, BITCOINIST, Mar. 24, 2024, <https://perma.cc/Z864-YEGF>.

store ever-increasing amounts of data, it will become necessary to craft a different type of societal response that takes aim more directly at excessive resource consumption by the wealthiest.¹⁰⁴ Successful strategies likely will need to include oligarch-sized sticks as well as conventional regulatory carrots.

4. Acceleration revisited: Catalyzing democratic datastructuring

Finally, return to the accelerationist ethos that now underpins a global AI development race. Both policymakers wanting to take AI datastructures and their dysfunctions seriously and politicians wanting to “win” the fast-coalescing global race for supremacy in AI development should be asking questions about how to structure incentives for datastructure development within a nonauthoritarian political system. In part, that project requires foregrounding the questions now being asked in other parts of the world about whether AI should be used at all, and, if so, how it should be used. In part, it requires looking beyond the acceleration/rejection binary to ask different kinds of questions about what Salome Viljoen has called data as a democratic medium.¹⁰⁵

The first step toward democratic datastructuring must be conceptual. Legal thinking about data access and data use falls too readily into the familiar lines dictated by speech, property, and privacy frames, none of which is well suited to address the challenge that democratic datastructuring represents. The speech frame is profoundly mismatched to the functional roles that data and datastructuring perform within generative AI systems.¹⁰⁶ Taking the property frame more seriously—principally via requirements for copyright licensing and other kinds of access licensing—would not necessarily ensure more systematic reliance on higher quality data. Many of the kinds of content described in Part II.A, above, are also copyrighted.

¹⁰⁴ See generally Clinton G. Wallace & Shelley Welton, *Taxing Luxury Emissions*, 109 CORNELL L. REV. 1153 (2024).

¹⁰⁵ Salome Viljoen, *A Relational Theory of Data Governance*, 131 YALE L.J. 573, 634 (2021).

¹⁰⁶ See Burk, *supra* note 99; see also Cohen, *Infrastructuring*, *supra* note 11; LAURA DENARDIS, *THE INTERNET IN EVERYTHING: FREEDOM AND SECURITY IN A WORLD WITH NO OFF SWITCH* (2020).

As Alicia Solow-Niederman explains, requirements to license are far more likely to produce (and are already producing) workarounds that have the effect of facilitating access to mass content.¹⁰⁷ And finally, even if data minimization requirements in privacy laws were taken seriously, generative AI systems use more than just personal data; mass content datastructures that process personal data also include many other elements. It is therefore important to begin a broader discussion about datastructure redesign strategies that would target the “systemic and waveform attributes” of information flows within platformized online communication systems.¹⁰⁸

Development of practices and conventions for democratic datastructuring also will require significant experimentation. An extended discussion of that project is beyond the scope of this paper, so I will simply suggest three important directions for research. One, which builds on Viljoen’s discussion of public-data trusts, involves exploring what non-content specifications for publicly maintained datastructures should look like.¹⁰⁹ Such an exercise, in turn, might inform thinking about governance of private sector datastructuring in a way that moves beyond prohibitions on distortive datastructuring and envisions democratic datastructuring more systematically and affirmatively. Another research direction, which builds on work by Vicki Jackson on the essential role of knowledge institutions in democratic societies, involves consideration of what lessons such institutions might bring to the project of crafting more sustainably democratic datastructures.¹¹⁰

¹⁰⁷ Alicia Solow-Niederman, *AI and Doctrinal Collapse*, 78 STAN. L. REV. (forthcoming 2026), <https://perma.cc/KA5A-MJC4>; see also Blake Reid, *What Copyright Can’t Do*, 52 PEPP. L. REV. 519 (2025).

¹⁰⁸ Cohen, *Infrastructuring*, *supra* note 11, at 27; see also Ayelet Gordon-Tapiero, Paul Ohm, & Ashwin Ramaswami, *Fact and Friction: A Case Study in the Fight Against False News*, 57 U.C. DAVIS L. REV. 171 (2023).

¹⁰⁹ Viljoen, *supra* note 105, at 645-49.

¹¹⁰ Vicki C. Jackson, *Knowledge Institutions in Constitutional Democracy: Reflections on ‘the Press’*, 14 J. MEDIA L. 275 (2022); Vicki C. Jackson, *Knowledge Institutions in Constitutional Democracies: Of Objectivity and Decentralization*, HARV. L. REV. BLOG, Aug. 29, 2019, <https://perma.cc/6UK8-RMWW>; see also Erin C. Carroll, *A Free Press Without Democracy*, 56 U.C. DAVIS L. REV. 289 (2022); Erin C. Carroll, *Beyond Democracy: How a Free Press Supports the Rule of Law*, 47 CARDOZO L. REV. (forthcoming 2025).

(Projects one and two may overlap; for example, libraries are knowledge institutions that embody time-tested and squarely public-regarding datastructures.) A third important research direction involves developing more coherent thinking about how to harden democratic information systems against authoritarian attacks, whether foreign or domestic.¹¹¹ This issue is now widely recognized as important, but thinking about how to approach it has too often fallen into familiar but unproductive ruts, mistakenly conflating data controls (and their effects) with content controls (and their effects) and/or mistakenly conflating concerns about data-driven manipulation with foreign ownership of instrumentalities for data-driven communication.¹¹² It is important, instead, to focus carefully on the ways that current datastructuring practices, whether owned by foreign or domestic actors, facilitate weaponization in the service of authoritarian goals and on how democratic datastructuring might begin to alter those affordances.

Conclusion

The linked challenges of AI development and governance have been significantly misframed in three important ways. First, the very large datasets used to train generative AI systems are not simply inert inputs. They are essential infrastructures whose composition and configuration both enables and constrains AI development. Second, the composition and configuration of current platformized AI datastructures lie at the root of two very significant problems—simulation and sociopathy—that generative AI systems persistently manifest. Governance approaches that focus

¹¹¹ Cf. Henry Farrell & Bruce Schneier, *Common Knowledge Attacks on Democracy*, Publication No. 2018-7, Berkman Klein Center for Internet & Society (Oct. 2018), <https://perma.cc/Y7F8-QM3S>.

¹¹² For an example of the former, see Brief of First Amendment and Internet Law Professors as *Amici Curiae* in Support of Petitioners, *TikTok, Inc. v. Garland*, No. 24-656, <https://perma.cc/RQE9-LLNT>. For examples of the latter, see Ganesh Sitaraman, *The Regulation of Foreign Platforms*, 74 STAN. L. REV. 1073 (2022); Brief of *Amici Curiae* Former National Security Officials in Support of Respondent, *TikTok v. Garland*, Nos. 24-656, 24-657, <https://perma.cc/E9N9-HEGL>.

principally on models and their uses rather than on the composition and configuration of AI datastructures will not address those problems. Third and finally, addressing the governance challenges posed by AI datastructures requires both more careful thinking about the limits of preexisting regulatory toolkits and more deliberate and systematic investment in regulatory innovation. A regulatory model for the emerging AI era must contend directly with the design of essential infrastructures and with the upstream and downstream implications of infrastructure design choices.