

Machine Learning as Natural Monopoly

Tejas N. Narechania*

ABSTRACT: Machine learning is transforming the economy, reshaping operations in communications, law enforcement, and medicine, among other sectors. But all is not well: Many machine-learning-based applications harvest vast amounts of personal information and yield results that are systematically biased. In response, policy makers have begun to offer a range of incomplete solutions. In so doing, they have overlooked the possibility—suggested intuitively by scholars across disciplines—that these systems are natural monopolies and have thus neglected the long legal tradition of natural monopoly regulation.

Drawing on the computer science, economics, and legal literatures, I find that some machine-learning-based applications may be natural monopolies, particularly where the fixed costs of developing these applications and the computational costs of optimizing these systems are especially high, and where network effects are especially strong. This conclusion yields concrete policy implications: Where natural monopolies exist, public oversight and regulation are typically superior to market discipline through competition. Hence, where machine-learning-based applications are natural monopolies, this regulatory tradition offers one framework for confronting a range of issues—from privacy to accuracy and bias—that attend to such systems. Just as prior natural monopolies—the railways, electric grids, and telephone networks—faced rate and service regulation to protect against extractive, anticompetitive,

* Robert and Nanci Corson Assistant Professor of Law, University of California, Berkeley, School of Law. For helpful comments and suggestions, I owe many thanks to Michael Abramowicz, Abhay Aneja, Abbye Atkinson, Ken Bamberger, Pam Bookman, Maureen Carroll, Andrew Chin, Zach Clopton, Erin C. Delaney, Charles Duan, Robin Effron, Seth Endo, Mary Fan, Amit Gandhi, Rebecca Goldstein, Mark Gergen, David Grewal, Scott Hemphill, Chris Hoofnagle, Bert Huang, Ben Johnson, Sonia Katyal, Aniket Kesari, F. Scott Kieff, Mark Lemley, Joy Milligan, Aileen Nielson, David Noll, Michael Murray, Manisha Padi, Claudia Polsky, Asad Rahim, Steve Ross, David Schwartz, Paul Schwartz, David Simon, Aaron Simowitz, Erik Stallman, Maurice Stucke, Daniel Walters, Steve Weber, Rebecca Wexler, Tim Wu, Christopher Yoo, Adam Zimmerman, and audiences at the Center for Long Term Cybersecurity at the University of California, Berkeley, School of Information, the George Washington University Law School, the University of California, Berkeley, School of Law, the University of North Carolina School of Law, and the University of Pennsylvania Carey Law School. For outstanding research assistance, I thank the stellar law librarians at Berkeley Law, Tian Kisch, Delia Scoville, and Kaavya Shah. I also thank Adam Garcia and the editors of the *Iowa Law Review* for their helpful suggestions and careful edits.

and undemocratic behaviors, so too might machine-learning-based applications face similar public regulation to limit intrusive data collection and protect against algorithmic redlining, among other harms.

I. INTRODUCTION.....	1544
II. NATURAL MONOPOLIES AND REGULATION	1557
A. <i>UNDERSTANDING NATURAL MONOPOLIES</i>	1557
B. <i>NATURAL MONOPOLIES AND REGULATORY POLICY</i>	1559
1. Rate Regulation.....	1560
2. Service Specification.....	1563
3. Induced Competition	1566
III. MACHINE LEARNING AS NATURAL MONOPOLY.....	1569
A. <i>UNDERSTANDING MACHINE LEARNING</i>	1570
B. <i>MACHINE LEARNING AS NATURAL MONOPOLY</i>	1574
1. Application Development.....	1574
2. Training and Optimization	1578
3. Machine Learning's Virtuous Circle.....	1584
IV. MACHINE LEARNING AND REGULATORY POLICY	1588
A. <i>SOME PROBLEMS OF MACHINE-LEARNING-BASED</i> <i>APPLICATIONS</i>	1590
B. <i>MACHINE LEARNING AND MONOPOLY REGULATION</i>	1596
1. Informational Rate Setting for Privacy Harms.....	1596
2. Service Standards for Accuracy and Bias Harms.....	1603
3. Toward Competition.....	1608
V. CONCLUSION: PUBLIC REGULATION AND PRIVATE POWER.....	1613

I. INTRODUCTION

Our modern economy has been built atop industries regulated as natural monopolies. Railroads, for example, helped to transform the country's agrarian economy to an industrial one, opening distant markets to mass-produced goods.¹ Power utilities brought electricity to homes and businesses

1. See generally SARAH H. GORDON, *PASSAGE TO UNION: HOW THE RAILROADS TRANSFORMED AMERICAN LIFE, 1829-1929* (1996) (describing and examining the economic, legal, and social effects of the expansion of the railways nationwide).

nationwide.² And the telephone helped to knit together far-flung communities by way of a single unified communications network.³

Now a new class of natural monopoly—one that again seems critical to our evolving modern economy—may be emerging. Machine-learning-based applications (sometimes termed artificially intelligent systems) are transforming the ways we generate, store, use, and interpret information.⁴ Such applications are embedded in a vast range of systems—from medicine to law enforcement to internet platforms and beyond.⁵ In short, machine-learning-based applications are beginning to pervade the economy, transforming operations across sectors, much as the railways, electric grids, and communications platforms once did.⁶

Natural monopolies are exemplary candidates for regulation.⁷ This is for two reasons: One, because they are *monopolies*, consumers are likely to face high prices, poor quality, or both. Two, because these markets tend *naturally* to monopoly control, competition is unlikely to address these price- and quality-related problems—indeed, attempts to introduce competition may be more harmful than helpful. So how might we address the problems that attend to inefficiently high prices or substandard quality without turning to competitors (who, in other contexts, typically aim to offer better products at better prices)? Regulation. As Judge Posner once explained, regulating natural monopolies helps avoid “wasteful” competition by trading a promise

2. See, e.g., Eric Biber, Sarah E. Light, J.B. Ruhl & James Salzman, *Regulating Business Innovation as Policy Disruption: From the Model T to Airbnb*, 70 VAND. L. REV. 1561, 1589–90 (2017).

3. See, e.g., TIM WU, *THE MASTER SWITCH: THE RISE AND FALL OF INFORMATION EMPIRES* 4 (Vintage Books 2011). See generally JONATHAN E. NUECHTERLEIN & PHILIP J. WEISER, *DIGITAL CROSSROADS: TELECOMMUNICATIONS LAW AND POLICY IN THE INTERNET AGE* (2d ed. 2013) (describing the development of the nation’s telecommunications industry (and of its telecommunications policy)).

4. See, e.g., STUART RUSSELL, *HUMAN COMPATIBLE: ARTIFICIAL INTELLIGENCE AND THE PROBLEM OF CONTROL* 62–102 (2019) (cataloguing near-term and long-term changes based on artificial intelligence and suggesting that, cumulatively, these developments “could change the dynamic of history”).

5. See, e.g., Jason Furman & Robert Seamans, *AI and the Economy*, 19 INNOVATION POLY & ECON. 161, 161–62, 177, 181–82 (2019); see also James Bessen, Stephen Michael Impink, Lydia Reichensperger & Robert Seamans, *The Business of AI Startups* 2–3, 24–25 (BU Sch. of L., L. & Econ. Series, Working Paper No. 18-28, 2018) (describing the impact of artificial intelligence on labor markets across sectors).

6. STUART J. RUSSELL & PETER NORVIG, *ARTIFICIAL INTELLIGENCE: A MODERN APPROACH* 28 (3d ed. 2015) (explaining that “many thousands of AI applications are deeply embedded in the infrastructure of every industry” (quoting RAY KURZWEIL, *THE SINGULARITY IS NEAR: WHEN HUMANS TRANSCEND BIOLOGY* 204 (2005))); see also PEDRO DOMINGOS, *THE MASTER ALGORITHM: HOW THE QUEST FOR THE ULTIMATE LEARNING MACHINE WILL REMAKE OUR WORLD*, at xi (2015) (“You may not know it, but machine learning is all around you.”).

7. E.g., George L. Priest, *The Origins of Utility Regulation and the “Theories of Regulation” Debate*, 36 J.L. & ECON. 289, 295 (1993) (explaining that natural monopolies “have traditionally provided the strongest public interest justification for regulation”).

of monopoly status to one provider “in exchange [for] a commitment to provide reasonable service at reasonable rates.”⁸

Hence, to return to the telecommunications precedent, regulators would grant telephone providers a license to operate in exchange for public oversight.⁹ Such oversight has historically encompassed rate regulation (to ensure that consumers were not charged inefficiently high prices), service specification (to ensure that the monopoly service satisfied public standards), and build-out requirements (to ensure that the network served the entire community), among other rules.¹⁰ Of course, this process was not perfect. It was difficult, for example, for regulators to calculate the real costs of a telephone provider’s network, to predict the investments required for upkeep and expansion, and to assess a reasonable profit. And providers themselves had a stake in inflating their costs, exaggerating their difficulties, and currying favor with local authorities.¹¹ But policy makers agreed that public regulation and oversight of this sort—such defects notwithstanding—were preferable to any alternative (such as unregulated monopolist control of a given market).¹²

8. *Omega Satellite Prods. v. City of Indianapolis*, 694 F.2d 119, 126 (7th Cir. 1982); see also Richard A. Posner, *Natural Monopoly and Its Regulation*, 21 STAN. L. REV. 548, 548–50 (1969) [hereinafter Posner, *Natural Monopoly*] (describing and explaining the natural monopoly concept). Judge Posner was ultimately skeptical that this theory for regulation applied to cable television systems, as *Omega Satellite* (and his other writings) suggest. *Omega Satellite*, 694 F.2d at 126 (describing the finding of natural monopoly conditions “sketchy”); Richard A. Posner, *The Appropriate Scope of Regulation in the Cable Television Industry*, 3 BELL J. ECON. & MGMT. SCI. 98, 111 (1972) [hereinafter Posner, *Scope*]; see also Richard A. Posner, *Theories of Economic Regulation*, 5 BELL J. ECON. & MGMT. SCI. 335, 337–39 (1974) [hereinafter Posner, *Theories*] (finding a reformulation of the public interest theory to be unsatisfactory).

9. See, e.g., *Cnty. Commc’ns Co. v. City of Boulder*, 455 U.S. 40, 43–46 (1982).

10. NUCHESTERLEIN & WEISER, *supra* note 3, at 32–40; see also DELOS F. WILCOX, MUNICIPAL FRANCHISES: A DESCRIPTION OF THE TERMS AND CONDITIONS UPON WHICH PRIVATE CORPORATIONS ENJOY SPECIAL PRIVILEGES IN THE STREETS OF AMERICAN CITIES 252–396 (1910) (describing, in detail, provisions in some of the earliest municipal franchise contracts).

11. NUCHESTERLEIN & WEISER, *supra* note 3, at 34–35.

12. An alternative theory, advanced by George Stigler (among many others), is that regulation is itself a commodity that is best understood as demanded by some and supplied by others. See George J. Stigler, *The Theory of Economic Regulation*, 2 BELL J. ECON. & MGMT. SCI. 3, 11–16 (1971). Judge Posner compared Stigler’s theory to the “public interest” theory of regulation elaborated above to a somewhat more cynical “capture” theory of regulation and concluded that neither adequately explained all utility regulation (which includes, but is not necessarily limited to, natural monopoly regulation). See Posner, *Theories*, *supra* note 8, at 343–44. Others have since taken a more pluralistic approach. See Priest, *supra* note 7, at 293–94.

I do not mean to weigh in on this longstanding debate over the true origins of utility regulation, though I suspect that no single theory can adequately explain the true origins of utility regulation across all states and municipalities nationwide. It suffices for my purposes to note that certain forms of regulation can be—and have long been—justified by the (apparent) existence of natural monopoly conditions. See WILCOX, *supra* note 10, at 28–33 (explaining that municipal utility franchising is explained, in part, by “the natural tendency toward monopoly” in such markets).

Such regulation has since fallen out of favor in many of these contexts—not necessarily for any defect in this basic theory undergirding natural monopoly regulation generally, but rather because these industries have since been disrupted in ways that undermine their status as natural monopolies.¹³ For example, once new telecommunications competitors like MCI proved it possible to provide long-distance service wirelessly,¹⁴ in competition with AT&T's wired networks, policy makers no longer saw that market as naturally tending toward monopoly, thereby undermining the case for regulatory control and leading to that market's deregulation.¹⁵ The railways were likewise disrupted by technological advancements in automotives, trucking, and air travel, leading to deregulation across the transportation industries.¹⁶

These sectors have experienced waves of deregulation. But policy makers are only now beginning to consider the possibilities for addressing the problems—rooted, perhaps, in monopoly conditions—that attend to machine-learning-based applications. For example, the White House recently announced the creation of the National Artificial Intelligence Initiative Office to “set[] AI technical standards” and to develop a “risk assessment framework” for such applications, among other initiatives.¹⁷ But the possibility that machine-

13. See, e.g., Biber et al., *supra* note 2, at 1591–92; Michael E. Levine, *Revisionism Revised? Airline Deregulation and the Public Interest*, 44 L. & CONTEMP. PROBS. 179, 191–93 (1981); see also *infra* notes 14–15; cf. CLAYTON M. CHRISTIANSEN, *THE INNOVATOR'S DILEMMA: WHEN NEW TECHNOLOGIES CAUSE GREAT FIRMS TO FAIL* 9–19 (1997) (describing such disruption). But see Priest, *supra* note 7, at 290 (suggesting that it was skepticism for this “public interest” natural monopoly theory of regulation—skepticism driven by the *Journal of Law and Economics* itself—that “significantly contributed to the movement toward broad deregulation”). One other explanation for the fading popularity of such regimes lies in the problems of regulatory capture and the difficulties with rate setting described above. I address these problems in more detail *infra* notes 74–78 and accompanying text. And, as noted above, I do not mean to discount the possibility that changes in the regulatory apparatus were motivated by something other than—or something in addition to—these technological developments. Rather, as above, it suffices for my purposes to note that changes in the regulatory regime may be—and have been—justified by technological changes that affect market structure (no matter whether we can pinpoint any single or set of historically-accurate rationale(s) for such shifts in the prevailing regulatory paradigm). See *supra* note 12 and accompanying text.

14. MCI was initially known, more formally, as Microwave Communications, Inc., but was later renamed to MCI Communications Corp. In either event, the company was one of AT&T's primary competitors in the long-distance telephony market (among other markets). See generally PHILIP LOUIS CANTELON, *THE HISTORY OF MCI 1968–1988* (1993) (describing a history of MCI from 1968 to 1988).

15. See W. KIP VISCUSI, JOSEPH E. HARRINGTON, JR. & JOHN M. VERNON, *ECONOMICS OF REGULATION AND ANTITRUST* 534–44 (4th ed. 2005) (describing the economic significance of microwave transmission); see also NUCHECHTERLEIN & WEISER, *supra* note 3, at 40–51 (describing the developments leading to the enactment of the (largely deregulatory) Telecommunications Act of 1996).

16. See, e.g., ALFRED E. KAHN, *THE ECONOMICS OF REGULATION: PRINCIPLES AND INSTITUTIONS*, at xv–xxxvii (1988); see also STEPHEN BREYER, *REGULATION AND ITS REFORM* 237–39 (1982) (discussing deregulation in the trucking industry).

17. Press Release, White House Off. of Sci. & Tech. Pol'y, *The White House Launches the National Artificial Intelligence Initiative Office* (Jan. 12, 2021), <https://trumpwhitehouse.archiv>

learning-based applications—that artificially intelligent systems—are natural monopolies suggests an even more robust range of regulatory possibilities. Just as in the early days of the railways and the telephone networks, society may be better off with regulated applications, ones with legally specified rates and service conditions. In short, the tradition of natural monopoly regulation offers policy makers the opportunity to regulate machine-learning-based applications in line with shared public values regarding privacy, accuracy, and algorithmic bias, all while granting the public the opportunity to participate in those regulatory proceedings.

In this Article, I analyze the claim, occasionally made intuitively by scholars across disciplines, that machine-learning-based applications are natural monopolies.¹⁸ Some lessons from the computer science and economics

es.gov/briefings-statements/white-house-launches-national-artificial-intelligence-initiative-office [https://perma.cc/2EVA-5TT5]; see National Artificial Intelligence Initiative, 15 U.S.C. § 9411 (2021); Exec. Order No. 13,859, 84 Fed. Reg. 3967 (Feb. 11, 2019); Christopher S. Yoo & Alicia Lai, *Regulation of Algorithmic Tools in the United States*, 13 J.L. & ECON. REGUL. 7, 9–15 (2020) (providing overview of federal regulatory initiatives for various artificial intelligence applications).

The Risk Assessment Framework is meant to help “better manage risks to individuals, organizations, and society associated with artificial intelligence . . . [by] improv[ing] the ability to incorporate trustworthiness considerations into the design, development, use, and evaluation of AI products, services, and systems.” *AI Risk Management Framework*, NAT’L INST. OF STANDARDS & TECH., <https://www.nist.gov/itl/ai-risk-management-framework> [https://perma.cc/8YTT-ZWHV].

18. See, e.g., Oren Bracha & Frank Pasquale, *Federal Search Commission? Access, Fairness, and Accountability in the Law of Search*, 93 CORNELL L. REV. 1149, 1180 (2008) (“It is unclear whether search engines fall under the strict definition of a natural monopoly, but they exhibit very similar characteristics.” (footnote omitted)); Steven Weber & Gabriel Nicholas, *Data, Rivalry and Government Power: Machine Learning Is Changing Everything*, 14 GLOB. ASIA 23, 26 (2019) (“[T]here’s justification for concern about natural monopolies”); see also AD J.W. VAN DE GEVEL & CHARLES N. NOUSSAIR, *THE NEXUS BETWEEN ARTIFICIAL INTELLIGENCE AND ECONOMICS* 97 (2013) (“For the first in a new class of products, a natural monopoly may briefly exist.”); A. Michael Froomkin, Ian Kerr & Joelle Pineau, *When AIs Outperform Doctors: Confronting the Challenges of a Tort-Induced Over-Reliance on Machine Learning*, 61 ARIZ. L. REV. 33, 88 (2019) (suggesting that machine-learning-based applications face “high fixed costs and low marginal costs[,] resembl[ing] the economic profile of a so-called natural monopoly”); Benjamin L. Mazer, Nathan Paulson & John H. Sinar, *Protecting the Pathology Commons in the Digital Era*, ARCHIVES PATHOLOGY & LAB’Y MED. 1037, 1038 (2020) (“[O]paque [machine learning] algorithms can lead to a natural monopoly.”); AD J.W. van de Gevel & Charles N. Noussair, *Artificial Intelligence and Economics: From Homo Sapiens to Robo Sapiens* 5 (International Conference on Advances in Computing, Electronics and Communication, 2015), <https://www.seekdl.org/conferences/paper/details/4669.html> [https://perma.cc/5Q7Y-9YKK] (suggesting that “AI systems, including autonomous robots, can be expected to generally have th[e] type of cost structure” associated with natural monopolies); Romesh Vaitilingam, *How Leading Economists View Antitrust in the Digital Economy*, LONDON SCH. OF ECON. & POL. SCI. (Nov. 18, 2020), <https://blogs.lse.ac.uk/businessreview/2020/11/18/how-leading-economists-view-antitrust-in-the-digital-economy> [https://perma.cc/5PDQ-7DXR] (“The tech industry is rife with natural monopolies, which are routinely regulated in other sectors.”) (internal quotation marks omitted). But see Herbert Hovenkamp, *Antitrust and Platform Monopoly*, 130 YALE L.J. 1952, 1971 (2021) (“Few platforms are natural monopolies.”).

Many commentators have focused on search markets in particular (without regard to whether those search applications rely on machine learning technology). See FRANCESCO DUCCI, *NATURAL MONOPOLIES IN DIGITAL PLATFORM MARKETS* 5, 47–75 (2020) (concluding that “[t]he

literatures may seem to confirm this contested intuition in some cases, particularly where development and data costs are high and network effects are strong.¹⁹ And lessons from the legal literature chart a regulatory path forward.²⁰

The defining characteristic of a natural monopoly is that the benefits of scale increase over the entire relevant market, such that the average cost to serve each consumer approaches some lower limit.²¹ Stated similarly, it is, on average, more efficient to serve a given market with only one provider rather than with many (no matter the number of providers actually in the market at any time).²² In at least some cases, this definition of natural monopolies seems to apply to machine-learning-based applications. This is for at least three reasons.

One, for some applications, the fixed costs of developing a machine-learning-based application overwhelm the marginal costs: It is far more expensive to develop the application than it is to make it available to its consumers. Of course, this is true for most software development: It is comparatively costless to deploy a software program (like Microsoft Windows or Google Sheets) once it is through development.²³ But machine-learning-based applications have even higher fixed costs.²⁴ The costs of acquiring the specialized hardware and volumes of data necessary to “train” the application

predictive and universal nature of horizontal search are such that data can give rise to a natural monopoly,” but reaching different conclusions for different sorts of applications); Rufus Pollock, *Is Google the Next Microsoft: Competition, Welfare and Regulation in Online Search*, 9 REV. NETWORK ECON., Dec. 2010, at 1, 3 (noting that the “service” section of a search engine bears a “much greater resemblance to a natural monopoly” than does the “software” section); Ioannis Lianos & Evgenia Motchenkova, *Market Dominance and Search Quality in the Search Engine Market*, 9 J. COMP. L. & ECON. 419, 435 (2013) (“[T]he cost structure of Google’s universal search engine is close to that of natural monopolies . . .”); Vaitilingam, *supra* note 18 (interviewing two academic economists who conclude that search is a natural monopoly). *But cf.* Eric K. Clemons & Nehal Madhani, *Regulation of Digital Businesses with Natural Monopolies or Third-Party Payment Business Models: Antitrust Lessons from the Analysis of Google*, 27 J. MGMT. INFO. SYS. 43, 63 (2010) (contending that “Google is not a natural monopoly” but drawing from prior natural monopoly regulatory regimes to consider approaches in internet search).

19. See *infra* Section III.B.

20. See *id.*

21. See, e.g., VISCUSI ET AL., *supra* note 15, at 401–02; William J. Baumol, *On the Proper Cost Tests for Natural Monopoly in a Multiproduct Industry*, 67 AM. ECON. REV. 809, 819 (1977) [hereinafter Baumol, *Proper Cost*].

22. See Posner, *Natural Monopoly*, *supra* note 8, at 548 (“If the entire demand within a relevant market can be satisfied at lowest cost by one firm rather than by two or more, the market is a natural monopoly, whatever the actual number of firms in it.”). I explain this in greater detail, *infra* Section II.A; see also William J. Baumol, *Quasi-Permanence of Price Reductions: A Policy for Prevention of Predatory Pricing*, 89 YALE L.J. 1, 11–14 (1979) [hereinafter Baumol, *Quasi-Permanence*] (comparing and contrasting the concept of scale with the concept of subadditivity).

23. See, e.g., *Microsoft Corp. v. AT&T Corp.*, 550 U.S. 437, 439 (2007). This is especially so when considering the costs of setting up a distribution architecture as among the fixed costs of software development.

24. See *infra* Section III.B.

can yield fixed costs that can outpace marginal costs at an even greater pace.²⁵ Competition in such a market (among aspiring natural monopolists) may thus yield higher losses (through wasteful duplication) and fewer gains.

Two, the computational process of training a machine-learning-based application is itself another significant fixed cost.²⁶ Machine-learning-based applications can be understood as developed in two stages—training and then prediction or classification.²⁷ In training, such applications must surmount a steep learning curve until optimized—until the costs of further training (in terms of, say, additional data acquisition and computing power) outweigh the benefits of such training (in terms of accuracy).²⁸ Training is computationally costly: Though training complexity (computationally speaking) varies across applications and specific machine-learning implementations, it is generally far more costly to train an application than it is to ask a trained application to respond to a query—to, say, predict a successful treatment regimen, identify a criminal suspect, or respond to an online search query.²⁹ Hence, this vast difference between the fixed costs of training and the marginal costs of prediction suggests that the average marginal computational cost of each prediction is likely to decrease across all users of the application—in short, machine-learning-based applications appear computationally subadditive, a sufficient condition for natural monopoly. This is particularly true for deep learning applications that are based upon large text- or image-based datasets and that demand intense computational power, including, for example, GPT-3 (a powerful language prediction model).³⁰ A second such application is likely

25. I use the term “train” here in quotation marks in view of the controversy over whether it is accurate or appropriate to say that these algorithmic systems are capable of being trained or of learning. *See, e.g.*, RUSSELL & NORVIG, *supra* note 6, at 1041–48 (summarizing this debate). After this instance, I no longer place the term in quotation marks—but, by doing so, do not mean to diminish the significance of this philosophical, etymological, and, indeed, existential question.

26. *See infra* Section III.B.2.

27. *See, e.g.*, David Lehr & Paul Ohm, *Playing with the Data: What Legal Scholars Should Learn About Machine Learning*, 51 U.C. DAVIS L. REV. 653, 655 (2017) (suggesting “that legal scholars should think of machine learning as consisting of two distinct workflows,” first, the design, development and training of an application, and second, an application “deployed and making decisions in the real world”).

28. *See, e.g.*, IAN GOODFELLOW, YOSHUA BENGIO & AARON COURVILLE, DEEP LEARNING 114–17 (2016) (explaining that as “the training set increases” the “optimal capacity” of the machine learning system also increases until it “plateaus” in view of having “reach[ed] sufficient complexity to solve the task”).

29. *See infra* Section III.B.2.

30. *See, e.g.*, TOM B. BROWN, BENJAMIN MANN, NICK RYDER, MELANIE SUBBIAH, JARED KAPLAN, PRAFULLA DHARIWAL, ARVIND NEELAKANTAN, PRANAV SHYAM, GIRISH SASTRY, AMANDA ASKELL, SANDHINI AGARWAL, ARIEL HERBERT-VOSS, GRETCHEN KRUEGER, TOM HENIGHAN, REWON CHILD, ADITYA RAMESH, DANIEL M. ZIEGLER, JEFFREY WU, CLEMENS WINTER, CHRISTOPHER HESSE, MARK CHEN, ERIC SIGLER, MATEUSZ LITWIN, SCOTT GRAY, BENJAMIN CHESS, JACK CLARK, CHRISTOPHER BERNER, SAM MCCANDLISH, ALEC RADFORD, ILYA SUTSKEVER & DARIO AMODEI, CORNELL UNIV., LANGUAGE MODELS ARE FEW-SHOT LEARNERS 5 (2020), <https://arxiv.org/abs/2005.14165> [<https://perma.cc/8XRK-NLPV>].

to duplicate the costs and training efforts of the first application, while adding little new capacity and apparent comparative benefit.³¹ The second provider is largely redundant: It is akin to a series of parallel railroad tracks, or a second, overbuilt telephone network.

Three, some machine learning algorithms—deep learning or continual learning systems, for example—continue training even after development, using information drawn from their practical applications to further improve accuracy or performance.³² This process of retraining—a virtuous circle, to adopt the industry’s terminology—can give the first system deployed at scale an insurmountable lead.³³ The reason is straightforward: The first system to market gains access to new training data—data from market consumers—before any competitor. That new information makes the first system better still (compared to its competitors)—and so that system may earn even more subscribers. Those new subscribers give that first system even more information with which to further improve—and so on. In such cases, one provider is consistently better—or, at minimum, has access to better and more complete data—than the rest.³⁴ Hence, given these network effects, limiting the market to one provider helps ensure that that system “learns” from the entire corpus of information provided by the market it serves.³⁵

31. See *infra* Section III.B. As I explain there, my claim turns on some tricky questions of market definition. And, as I readily concede, some applications may operate in markets that can sustain competition because, say, different applications may serve different segments of the same market. See *infra* note 185–86 and accompanying text (noting possibilities for product differentiation). In short, I am careful not to suggest that all machine-learning-based applications necessarily operate in a natural monopoly market. Rather, my claim is, as I say above, more narrow: I find it likely that at least some machine-learning-based applications, when serving a single market, are natural monopolists. See also *infra* note 35–36 and accompanying text. (“I do not mean to contend that all algorithmic systems . . . fall into this category. But I do conclude that some machine-learning-based applications [in some scenarios] . . . seem to tend toward natural monopoly.” (footnote omitted)).

32. Importantly, not all machine learning systems learn continually. See, e.g., Arti K. Rai, Isha Sharma & Christina Silcox, *Accountability, Secrecy, and Innovation in AI-Enabled Clinical Decision Software*, 7 J.L. & BIOSCIENCES, Jan.–June 2020, at 1, 3, 5 (explaining that machine-learning-based applications “could either be locked, or [] could be allowed to learn in real-time from new data to which it was exposed,” but that the “FDA has thus far cleared only locked models” for use in clinical settings).

33. See ANDREW NG, *LANDING AI, AI TRANSFORMATION PLAYBOOK: HOW TO LEAD YOUR COMPANY INTO THE AI ERA* 4 (2019), https://landing.ai/wp-content/uploads/2020/05/Landing_AI_Transformation_Playbook_11-19.pdf [<https://perma.cc/Y5LW-A4BW>]; see also Jaehun Lee, Taewon Suh, Daniel Roy & Melissa Baucus, *Emerging Technology and Business Model Innovation: The Case of Artificial Intelligence*, 5 J. OPEN INNOVATION: TECH., MKT. & COMPLEXITY 44, 49–51 (2019) (similarly describing the “virtuous cycle of AI”).

34. See, e.g., Frank Pasquale, *Rankings, Reductionism, and Responsibility*, 54 CLEV. ST. L. REV. 115, 129–30 (2006) (“Growing numbers of searches on a given service give that service ever more data to refine and improve its index. The ‘rich get richer,’ making the search and rankings field a very difficult one to enter.” (footnote omitted)).

35. See *supra* note 25 and accompanying text (regarding my use of the term “train”).

I do not mean to contend that all algorithmic systems—a set that includes a vast range of technology, including other systems (e.g., expert systems) that are sometimes also called artificial intelligence—fall into this category.³⁶ But I do conclude that some machine-learning-based applications—especially where data are expensive, computational demands are high, and network effects are strong—seem to tend toward natural monopoly.

This conclusion carries some specific consequences. As noted above, natural monopolists may resist competition, charge supracompetitive prices, and skimp on service quality. In response, policy makers have long sought to regulate natural monopolies in order to retain the gains associated with consolidation, such as the effects of the virtuous circle, while minimizing these harms of monopoly power.³⁷ Drawing upon this tradition of regulation, I begin to consider a policy approach for machine-learning-based systems, focusing on three regulatory levers—rate regulation, service specification, and induced competition—each designed to address problems arising from the natural monopoly condition.

First, monopolists (including natural monopolists) typically charge inefficiently high prices, giving rise to problems that sound in both distributive and welfare concerns. In response, policy makers have traditionally regulated natural monopolists' rates to protect consumers from such charges. In the case of machine learning regulation, this can encompass direct regulation of the data collected and used by these systems and their developers. Drawing from the (perhaps tired) metaphor that consumers “pay” for various services with personal information,³⁸ we can understand providers'

36. See, e.g., Bruce G. Buchanan, *Can Machine Learning Offer Anything to Expert Systems?*, 4 MACH. LEARNING 251, 251 (1989) (explaining that, historically, most “expert systems do not employ a learning component to construct parts of their knowledge bases from libraries of previously solved cases”); cf. DIGIT. COMPETITION EXPERT PANEL, UNLOCKING DIGITAL COMPETITION: REPORT OF THE DIGITAL COMPETITION EXPERT PANEL 54–55 (2019), https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/785547/unlocking_digital_competition_furman_review_web.pdf [<https://perma.cc/4BXX-VT8S>] (analogizing digital markets generally to natural monopolies, but concluding that the analogy may not hold across the entire sector).

37. See, e.g., *Cantor v. Detroit Edison Co.*, 428 U.S. 579, 595–96 (1976); *Otter Tail Power Co. v. United States*, 410 U.S. 366, 386–89 (1973) (Stewart, J., concurring in part and dissenting in part); see also Priest, *supra* note 7, at 301 (explaining that municipalities regulated utilities, through franchise contracts that set rates and service quality, as early as “the early decades of the nineteenth century”).

38. See Ann Brody Guy, *Reinventing Cybersecurity*, BERKELEY ENG'G (Apr. 14, 2020), <https://engineering.berkeley.edu/news/2020/04/reinventing-cybersecurity> [<https://perma.cc/3GED-4C9S>] (noting the “conventional wisdom” that where “a service is free, you are the product”); see also TIM WU, THE ATTENTION MERCHANTS: THE EPIC SCRAMBLE TO GET INSIDE OUR HEADS 376–81 (2016) (explaining how human behavior is monetized by companies providing free services online); Chris Jay Hoofnagle & Jan Whittington, *Free: Accounting for the Costs of the Internet's Most Popular Price*, 61 UCLA L. REV. 606, 608–12 (2014) (describing how online businesses are offering services marketed as “free” but use such programs to obtain consumer's data in return); Shira Ovide, *Google Ends Its Free Digital Photo Storage*, N.Y. TIMES (Nov. 12, 2020), <https://>

privacy-invading data collection practices as an unduly extractive charge, as well as an inefficiently high price that yields welfare losses: Some developers may demand and hoard unfairly vast volumes of user data (using that information to exploit consumers, among other things), while some privacy-sensitive consumers may decline to use a beneficial machine-learning-based advance because of such demands.³⁹ In response, regulators might promulgate a scheme of rate regulation that both advances privacy interests and reflects public preferences by limiting data collection and use to only that which is reasonably necessary for the continued development of the regulated system, or by directly specifying the information that providers may collect and keep.⁴⁰ These alternatives—drawn from the leading modes of rate regulation, rate-of-return regulation and price-cap regulation—offer two concrete paths, steeped in an existing regulatory tradition, for addressing the varied privacy-related concerns attending to the volumes of personal information amassed by private companies. Such public law approaches, moreover, are better suited to the relational and interconnected nature of data than privacy law’s traditional focus on individual rights and harms.

Second, monopolists may have private incentives to underinvest in their systems in pernicious ways—say, by underinvesting in safety and security protocols or by overlooking the social costs of discrimination.⁴¹ Though

www.nytimes.com/2020/11/12/technology/2020-hot-gadget-is-the-computer.html [https://perma.cc/RL7Z-SVND] (explaining that Google’s service “has never really been free” because “Google uses our photos to train its software systems”).

39. See, e.g., Maurice E. Stucke, *Should We Be Concerned About Data-opolies?*, 2 GEO. L. TECH. REV. 275, 285–86, 302 (2018) (explaining that a monopolist in a data-intensive industry “has the incentive to reduce its privacy protection below competitive levels and collect personal data above competitive levels” giving rise to privacy harms and deadweight losses); Ryan Calo, *Digital Market Manipulation*, 82 GEO. WASH. L. REV. 995, 1018 (2014); Shira Ovide, *Just Collect Less Data, Period*, N.Y. TIMES (July 16, 2020), <https://www.nytimes.com/2020/07/15/technology/just-collect-less-data-period.html> [https://perma.cc/4FQX-9WTG] (“For the companies, there’s no downside to limitless data collection, and there’s little to prevent them from doing so in the United States.”); see also Jennifer Valentino-DeVries, *Coronavirus Apps Show Promise but Prove a Tough Sell*, N.Y. TIMES (Dec. 7, 2020), <https://www.nytimes.com/2020/12/07/technology/coronavirus-exposure-alert-apps.html> [https://perma.cc/W88U-UE8N] (explaining that privacy concerns have depressed interest in coronavirus-related contact-tracing applications); Cynthia Cole, Brooke Chatterton & Natalie Sanders, *The Safety of Privacy: Increased Privacy Concerns May Prevent Effective Adoption of Contact Tracing Apps*, LAW.COM (Aug. 18, 2020, 7:00 AM), <https://www.law.com/legaltechnews/2020/08/18/the-safety-of-privacy-increased-privacy-concerns-may-prevent-effective-adoption-of-contact-tracing-apps> [https://perma.cc/CH4L-MC79] (“Privacy concerns, which lead to low adoption rates, are a barrier to the success of a contact tracing app.”).

40. See, e.g., *infra* note 234 and accompanying text (describing HIPAA as an example of such a provision); cf. CAL. CIV. CODE § 1798.100 (West 2020) (explaining that personal information may be used for business purposes only to the extent such use is “reasonably necessary and proportionate”).

41. See, e.g., Tejas N. Narechania, *The Secret Life of a Text Message*, 120 COLUM. L. REV. F. 198, 205 (2020) (suggesting that telephone providers “may have private incentives . . . to underinvest in safety-related infrastructure”). In some cases, providers might even discriminate intentionally. Cf. Elena Botella, *TikTok Admits It Suppressed Videos by Disabled, Queer, and Fat Creators*, SLATE (Dec.

competition can cause providers to offer ever-better services and to ensure that their products are available to the widest possible market, such competition is unlikely to discipline more highly concentrated markets.⁴² Hence, regulatory authorities have instead set service-quality standards for such providers to account for the difference between private incentives and the public good.⁴³ In the machine learning context, such regulation could take several forms. For one, regulators might require that systems satisfy some regulatorily-prescribed standard for accuracy (notwithstanding computational or market-based measures of optimization), to prevent the harms of a faulty medical application, or to prevent errant arrests based on defective facial recognition systems. Or regulators might use their authority to address bias and algorithmic redlining by requiring that these systems do not discriminate against historically marginalized and disadvantaged populations. In short, regulators might substitute social metrics—metrics that fare better along dimensions of accountability, legitimacy, and transparency, among other values—for computational ones.⁴⁴

Finally, because competition in natural monopoly markets can be wasteful and inefficient,⁴⁵ regulators will sometimes instead seek to induce efficient competition—and the concomitant competition-based outcomes (to avoid prescribing rates and service standards). For example, regulators sometimes require putative monopolists to compete by bidding for a franchise to serve a natural monopoly market (rather than compete within that market), thereby avoiding duplicative and wasteful fixed investments.⁴⁶ In

4. 2019, 5:07 PM), <https://slate.com/technology/2019/12/tiktok-disabled-users-videos-suppressed.html> [<https://perma.cc/NN93-JP4E>] (“TikTok, a social network video app with more than 1 billion downloads globally, admitted Tuesday to a set of policies that had suppressed the reach of content created by users assumed to be ‘vulnerable to cyberbullying.’”).

42. See Restoring Internet Freedom, 35 FCC Rcd. 12328, para. 48, 12351 n.170 (order on remand) (Oct. 27, 2020) (suggesting that competition among providers will improve safety-related services and applications).

43. See *infra* notes 80–89 and accompanying text.

44. See *Munn v. Illinois*, 94 U.S. 113, 124–25 (1876) (affirming regulators’ power to regulate private providers “for the public good”); K. Sabeel Rahman, *The New Utilities: Private Power, Social Infrastructure, and the Revival of the Public Utility Concept*, 39 CARDOZO L. REV. 1621, 1634–38 (2017); see also SOLON BAROCAS, MORITZ HARDT & ARVIND NARAYANAN, FAIRNESS AND MACHINE LEARNING: LIMITATIONS AND OPPORTUNITIES 13 (2021), <https://fairmlbook.org/pdf/fairmlbook.pdf> [<https://perma.cc/2Q6G-VZWW>] (introducing “formal methods for characterizing . . . problems [relating to objectionable results of machine learning systems] and assess[ing] various computational methods for addressing them”); cf. Seth Katsuya Endo, *Technology Opacity & Procedural Injustice*, 59 B.C. L. REV. 821, 862–63 (2018) (highlighting the importance of participation values (as against accuracy and efficiency concerns) when considering whether to use predictive coding in litigation discovery contexts).

45. Posner, *Natural Monopoly*, *supra* note 8, at 548 (suggesting that sustained competition in natural monopoly markets “produces inefficient results”).

46. See, e.g., *Cnty. Commc’ns Co. v. City of Boulder*, 660 F.2d 1370, 1373 (10th Cir. 1981) (“[T]he City . . . concluded that cable systems are natural monopolies. Consequently, the City became concerned that CCC, because of its headstart, would always be the only cable operator in

other contexts, policy makers aim to induce efficient competition by lowering the barriers to market entry by, say, allowing new competitors to share in existing investments,⁴⁷ or by imposing sharing requirements on market participants. Similar outcomes might result from a shift toward franchising for machine-learning-based applications deployed for public purposes (like law enforcement),⁴⁸ by mandating federated learning or otherwise requiring data-sharing by putative competitors,⁴⁹ and by closely regulating the terms of shared computing infrastructure providers to ensure neutrality.

In short, machine-learning-based applications make extractive and inefficient demands of their users, and they generate results that are both inaccurate and systematically biased. I do not, to be sure, claim to be the first to identify the accuracy, bias, and general privacy difficulties that attend to these applications.⁵⁰ Rather, this Article draws together two parallel debates that have informed these findings—one regarding ethics and transparency in machine-learning contexts, and another regarding power and consolidation

Boulder if allowed to expand . . . The City decided to place a moratorium on CCC's expansion in order to provide other companies the opportunity to make bids to service the remaining parts of Boulder . . .").

47. See, e.g., 47 U.S.C. § 251(c)(3) (2018); see also NUECHTERLEIN & WEISER, *supra* note 3, at 58 (describing the purpose of § 251(c)(3)).

48. See Deirdre K. Mulligan & Kenneth A. Bamberger, *Procurement as Policy: Administrative Process for Machine Learning*, 34 BERKELEY TECH. L.J. 773, 785–98 (2019).

49. Charles I. Jones & Christopher Tonetti, *Nonrivalry and the Economics of Data*, 110 AM. ECON. REV. 2819, 2820 (2020) (explaining that firms have incentives to “hoard data in ways that are socially inefficient”). In response, both Mark A. Lemley & Bryan Casey, *Fair Learning*, 99 TEX. L. REV. 743 (2021) and Amanda Levendowski, *How Copyright Law Can Fix Artificial Intelligence's Implicit Bias Problem*, 93 WASH. L. REV. 579 (2018) generally contend that copyright law should be interpreted so as to limit infringement liability for training data, thereby reducing the training costs associated with application development. But, as I elaborate below, copyright is only one limit among several on the availability and use of training data. See *infra* Section IV.B.3; cf. *Ruckelshaus v. Monsanto Co.*, 467 U.S. 986, 1018–19 (1984) (finding that, even though data used for product development and regulatory approval may be protected as trade secret and under the Constitution's Taking Clause, statutory schemes for arbitrating a price for the use of that data by competitors satisfy constitutional scrutiny). Federated learning may offer one way out of this quandary. See *TensorFlow Federated: Machine Learning on Decentralized Data*, TENSORFLOW, <https://www.tensorflow.org/federated> [<https://perma.cc/L5PW-DG54>] (defining federated learning as “an approach to machine learning where a shared global model is trained across many participating clients that keep their training data locally”).

50. On bias and accuracy, see, for example, Batya Friedman & Helen Nissenbaum, *Bias in Computer Systems*, 14 ACM TRANSACTIONS ON INFO. SYS. 330, 332–37 (1996) (describing preexisting, technical, and emergent bias); see also Joy Buolamwini & Timnit Gebru, *Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification*, 81 PROC. MACH. LEARNING RSCH. 1, 6–11 (2018) (evaluating three commercial gender classification systems and finding that darker-skinned females are misclassified in up to 34.7 percent of test cases, while lighter-skinned males are misclassified in only 0.8 percent of test cases). On privacy and autonomy, see, for example, Julie E. Cohen, *What Privacy Is for*, 126 HARV. L. REV. 1904, 1906, 1912–18 (2013); see also M. Ryan Calo, *The Boundaries of Privacy Harm*, 86 IND. L.J. 1131, 1151–52 (2011) (describing and classifying the privacy harms that attend to informational systems, including Google's Gmail (email) service).

in the information technology industry. Computer scientists, for example, have begun to investigate computational (and other) solutions to the problems of accuracy and bias in machine learning.⁵¹ And legal and social science scholars such as Safiya Umoja Noble, Frank A. Pasquale, and K. Sabeel Rahman, among others, have argued in favor of greater regulation of our informational infrastructure in view of the growing concentration in these markets, contending that companies such as Google and Facebook resemble modern utility providers.⁵² Meanwhile, others have responded with cautions against regulation, suggesting that innovation and market competition are better suited to address such problems.⁵³ But if these applications do indeed act as natural monopolists—or even, perhaps, as mere monopolists or oligopolists—then market power may contribute to these problems of accuracy, bias, and privacy. Hence, the legal traditions of natural monopoly regulation, and market regulation more generally, offer a ready, but perhaps overlooked, framework to help address these problems.⁵⁴ This Article consequently explores the extent to which this legal framework applies to machine-learning-based applications and considers how policy makers may draw from that public regulatory tradition in this new technological space. It

51. See BAROCAS ET AL., *supra* note 44, at 13 (“[I]ntroduc[ing] formal methods for characterizing these problems [relating to objectionable results of machine learning systems] and assess[ing] various computational methods for addressing them.”); MICHAEL KEARNS & AARON ROTH, *THE ETHICAL ALGORITHM: THE SCIENCE OF SOCIALLY AWARE ALGORITHM DESIGN* 16 (2019) (“[O]ur book . . . dives headfirst into the emerging science of designing social constraints directly into algorithms”); see also RUSSELL, *supra* note 4, at 171–83 (2019) (discussing “beneficial machines”); Benjamin Kuipers, *Perspectives on Ethics of AI: Computer Science*, in *THE OXFORD HANDBOOK OF ETHICS OF AI* 421, 433–35 (Markus D. Dubber, Frank Pasquale & Sunit Das eds., 2020) (discussing risks of bias); cf. LAWRENCE LESSIG, *CODE: VERISON 2.0* 1–9 (2006) (describing how the design of technical systems constrains and guides behavior); Joel R. Reidenberg, *Lex Informatica: The Formulation of Information Policy Rules Through Technology*, 76 TEX. L. REV. 553, 556–65 (1998) (similarly explaining how the technical systems can implement policy choices and values).

52. See SAFIYA UMOJA NOBLE, *ALGORITHMS OF OPPRESSION: HOW SEARCH ENGINES REINFORCE RACISM* 156–60 (2018); Rahman, *supra* note 44, at 1640–41; Frank A. Pasquale, *Internet Nondiscrimination Principles Revisited* 8–10 (Brook. L. Sch. Legal Stud., Working Paper No. 655, 2020); see also Bracha & Pasquale, *supra* note 18, at 1152, 1201–09 (outlining “some possible directions for effective regulation”).

53. E.g., Andrea O’Sullivan & Adam Thierer, *Counterpoint: Regulators Should Allow the Greatest Space for AI Innovation*, 61 COMM’NS ACM 33, 33 (2018) (“[A]rtificial intelligence technologies should largely be governed by a policy regime of permissionless innovation so that humanity can best extract all of the opportunities and benefits they promise.”); see also Alicia Solow-Niederman, *Administering Artificial Intelligence*, 93 S. CAL. L. REV. 633, 653 (2020) (counseling “against ‘command-and-control’ regulation”); cf. JULIE E. COHEN, *BETWEEN TRUTH AND POWER: THE LEGAL CONSTRUCTIONS OF INFORMATIONAL CAPITALISM* 90–92 (2019) (summarizing such arguments).

54. Cf. Rahman, *supra* note 44, at 1634–35 (considering regulatory possibilities through the “more expansive” lens of public utility regulation); *id.* at 1674–78 (considering regulatory possibilities for Uber, Airbnb, and like competitors after concluding that such “businesses may not be conventional natural monopolies” but are nevertheless likely to be entrenched and therefore require a regulatory response); John M. Newman, *Antitrust in Zero-Price Markets: Foundations*, 164 U. PA. L. REV. 149, 190–94, 198–200 (2015).

does so by building on the important prior work identifying issues with machine-learning-based applications and calling for greater regulation of those applications by casting those issues as problems of market power and by drawing on the tradition of natural monopoly and market regulation to address these specific concerns.

This Article proceeds in three parts. First, I begin with a brief description of the legal tradition of natural monopoly regulation. As noted above, policy makers have regulated apparent natural monopolies since at least the development of railway networks, both to target the specific and reasonably anticipated harms caused by the exercise of monopoly power (such as high prices and poor service), as well as to implement a range of social policies related to the natural monopoly condition (such as universal service). Second, I consider whether machine-learning-based applications fall into this class of natural monopolies. I find that at least some do, in part because of several features particular to machine-learning-based applications, including the means of training such applications to learn from masses of data. Third, given that at least some machine learning systems may—like their natural monopoly predecessors—be appropriate candidates for regulatory scrutiny, I begin to consider the possibilities for such a regime. Other scholars have noted myriad problems with algorithmic systems, including those related to accuracy, bias, and privacy. I suggest that some such problems may derive from a provider’s monopoly status and can therefore be addressed with tools found within the tradition of natural monopoly regulation.

II. NATURAL MONOPOLIES AND REGULATION

I begin with a brief overview of the history and tradition of natural monopoly regulation. Understanding natural monopoly regulation, however, demands an understanding of natural monopolies, and so I start by briefly elaborating this unique market condition before turning next to a description of paradigmatic regulatory responses.

A. UNDERSTANDING NATURAL MONOPOLIES

Scale characterizes natural monopolies: If the benefits of scale (i.e., decreases in the average cost to serve a customer) accumulate across an entire relevant market, then that market may be best served by a monopolist.⁵⁵ This may seem odd. The antitrust laws, for example, were originally enacted to counter the harmful effects of powerful monopolists. So why might some markets operate better under monopoly control? Again, the benefits of scale provide the answer: In natural monopoly markets, it is always more efficient (from the standpoint of productive efficiency) to add new users to the

55. See Baumol, *Quasi-Permanence*, *supra* note 22, at 11–12 (explaining “natural monopoly always has been associated loosely with the phenomenon of economies of scale”).

incumbent provider; it never makes sense to introduce a competitor.⁵⁶ And this is so no matter the number of competitors actually in the market—natural monopoly markets sometimes include (typically temporarily) multiple participants, yielding an initially inefficient allocation of market inputs (among other concerns described below).⁵⁷

Imagine, for example, a communications network with high fixed costs, including wire, pole, and labor costs, among others. Once the network is deployed, however, it is comparatively simple to connect a user to the network: In some instances, it is little more than flipping a switch.⁵⁸ In this example, it always makes sense for a user to connect to the existing network—and, inversely, it never makes sense for a user to connect to a different network. Why? The second, competing network duplicates those high fixed costs. Where the market is served by a single provider, it might spend only \$1,000,000 (as an example) to connect one million users—\$1 each. But if another provider puts up a competing network, then the competitors have spent \$2,000,000 combined to connect the same million users—\$2 each. With two providers instead of one, it costs more to do the same (assuming users of the two networks can even connect with each other). This gives rise to higher consumer prices, as well as increased costs for other companies competing for the same inputs—other companies that need copper wire or wooden poles. In short, from a production standpoint, one provider is better than two (because the market's cost function is, in a word, subadditive). Hence, even when multiple providers enter a natural monopoly market, one will typically win the competition for the market and consolidate the investments made by (now-former) competitors.

56. See *id.* at 11 (explaining that a market is “a natural monopoly if production by a single firm is the cheapest way to produce the combination of outputs supplied by the industry” and that, consequently, “if several firms can produce its output at least as cheaply as one, then” there is no natural monopoly); see also ROBERT BALDWIN, MARTIN CAGE & MARTIN LODGE, UNDERSTANDING REGULATION: THEORY, STRATEGY, AND PRACTICE 16–17, 444–46 (2d ed. 2012) (describing how natural monopolies occur and their effects); Baumol, *Proper Cost*, *supra* note 21, at 815–16 (discussing implications for evidence on natural monopolies).

Some readers might contend that a natural monopoly exists only where only single-firm control of an entire market can justify that firm's investment of the fixed costs required to serve the market. As I explain below, however, that's not quite right. As the following examples suggest, it may be profitable for a second firm to enter a market (either wholly or in part), either in the hopes of eventually winning the competition for the market, or perhaps to serve some lucrative market subset (e.g., cream-skimming). But such duplicative entry may nevertheless be inefficient (again, from the standpoint of productive efficiency).

57. See Posner, *Natural Monopoly*, *supra* note 8, at 548; cf. Hovenkamp, *supra* note 18, at 1971 (“While we sometimes use the term ‘natural monopoly’ to describe a firm or a market, natural-monopoly status actually applies to particular inputs or technologies.”).

58. See *U.S. Telecom Ass'n v. FCC*, 359 F.3d 554, 569 (D.C. Cir. 2004) (explaining that incumbent telecommunications providers typically need only make a “software change” in order to connect a new customer).

This is true, however, only where such scale economies extend throughout the relevant market. Compare, for example, some other complex commodity—an elevator, say. It is expensive to build an elevator factory and train elevator installers and inspectors. Once the factory is built and installers trained, however, it might be comparatively cheap to build and deploy an elevator. So is the elevator business, too, a natural monopoly? It seems unlikely. This is because the benefits of scale do not extend across the whole relevant market. It is unlikely that one factory and one cohort of elevator installers can support the market's demand for elevators. Hence, at some point, the market will demand a second factory and a second cohort of installers and inspectors—and it does not matter whether Otis Elevator Company or ThyssenKrupp makes that additional investment. The average cost per elevator is higher because of the new significant investments, and the fact of higher average costs does not depend on whether those investments are made by the incumbent or a new competitor. Scale economies do not extend across the entire relevant market because the market is greater than a single facility's capacity.⁵⁹

But if one provider—one communications network, to return to the earlier example—can support the market, then there is no need for a competitor. Indeed, such competition will typically give way to a monopoly, with higher prices for consumers and other industries in the interim. And monopoly power sometimes yields further benefits: AT&T devoted some of its capital to Bell Labs, which helped drive a vast amount of new, innovative activity. So policy makers must trade the benefits of consolidation against the negative consequences of monopoly power.

B. NATURAL MONOPOLIES AND REGULATORY POLICY

The possibility that some markets tend naturally toward monopoly control has long yielded some specific legal consequences. For example, “[t]he railroad industry represented a true natural monopoly when the Interstate Commerce Commission”—the nation's first regulatory agency

59. One competitor might, of course, be more efficient than the other—Otis, say, might be able to build a better factory more cheaply than ThyssenKrupp. But such differences are beside my point. It suffices to note that the market can support multiple providers, and that those providers will rationally compete amongst each other to provide the best elevators most cheaply, and that such market competition is more likely, we assume under this framework, to identify the better competitors as compared to, say, a central planner deciding who gets to build the second elevator factory.

Similarly, in intellectual property markets, the demand for books (or movies, or music) is greater than any one work of authorship. So while the ratio of fixed costs to average marginal costs for goods traditionally protected by intellectual property is similar to that seen in natural monopoly markets—indeed, that ratio is often cited as a primary justification for intellectual property protection—many such goods are not natural monopolies because of these greater market dynamics.

—“emerged to regulate it during the 1880s.”⁶⁰ Likewise, the Kingsbury Commitment—one of AT&T’s earliest encounters with regulation—was consummated by the Justice Department under its antitrust authorities to address AT&T’s apparent natural monopoly in long-distance telephone service.⁶¹ Over time, those regulatory interventions grew to encompass a range of state and federal oversight over an array of communications technologies.

Some aspects of these responses react to the possibility of monopoly control. Monopolists typically charge inefficiently high prices. Monopolists, that is, charge consumers more than they would under competitive conditions, thereby leaving some consumers unable to afford (or unwilling to buy) its socially valuable telephone or railway service.⁶² Similarly, monopolists can skimp on quality because they face no competitive pressure to improve.⁶³

Other regulatory responses flow from the view that these markets tend naturally toward monopoly control, sometimes entrenching the first provider to a market (without respect to whether the first is the best).⁶⁴ Communications regulators, for example, helped ensure affordable and reliable service across a wide range of locales.⁶⁵

In all, regulators have traditionally held some power to influence the rates, service standards, and number of providers in order to balance the unique benefits of consolidation in natural monopoly markets against these potential harms.

1. Rate Regulation

Rate regulation is perhaps the hallmark of natural monopoly regulation.⁶⁶ It is a direct response to the natural monopolist’s power over price, namely, its power to charge inefficiently high prices. What makes monopolists’ prices

60. THOMAS K. MCCRAW, *PROPHETS OF REGULATION* 306 (1984).

61. See VISCUSI ET AL., *supra* note 15, at 534–35. The Kingsbury Commitment is widely considered to be one of the first instances of federal telecommunications regulation. See NUECHTERLEIN & WEISER, *supra* note 3, at 5, 264–65.

62. Stated similarly, monopolists often find it profitable to charge prices above average marginal cost. As a result, some consumers—even consumers who are willing to pay more than it costs to produce a good—will be unable to pay the market price for the service. See *infra* note 71 and accompanying text (describing deadweight loss).

63. Such quality effects are analogous to the price effects described above. One way to restate this issue, for example, might be to say that the monopolist’s price is too high for such a low-quality good.

64. Cf. Jean Tirole, *Market Failures and Public Policy*, 105 AM. ECON. REV. 1665, 1669–70 (2015) (describing the role that luck can play in the acquisition of market power, including in natural monopoly and utility contexts).

65. See NUECHTERLEIN & WEISER, *supra* note 3, at 307–24.

66. *Verizon Commc’ns v. FCC*, 535 U.S. 467, 477 (2002) (“At the dawn of modern utility regulation, in order to offset monopoly power and ensure affordable, stable public access to a utility’s goods or services, legislatures enacted rate schedules to fix the prices a utility could charge.”).

(including natural monopolists' prices) inefficiently high?⁶⁷ Stated simply, it is the disconnect between a monopolist's incentives and the social good.

Ideally, under a standard rational choice framework, consumers buy goods that they think are worth more than their production costs. If a railway journey costs \$1 to complete, but is worth \$3 to the average consumer, then those consumers should buy train tickets: Such tickets are worth more than the paper they are printed on (together, of course, with the other costs incurred by the railroad to complete the journey). But this simple equation ($\$3 > \1) assumes that railroads will charge close to cost for their services. In competitive conditions, this is likely true: Different providers will compete to provide ever-better products and services at ever-lower prices (but they will not charge less than cost, because they are not interested in losing money on the enterprise). Hence, competition drives prices down to cost. So long as consumers value goods more than they cost, they'll buy them. And that's good, in general, for the world: Consumers get more out of the goods they buy than it costs to make those goods.⁶⁷

But what if there is no competition? In that case, the monopolist's price is not driven by competitive concerns—it faces no pressure to charge ever-lower prices. Instead, the monopolist chooses a price that it thinks will maximize its profits.⁶⁸ For our railroad monopolist, that might be, say, \$5. Some people will still make the journey: To them, the trip is especially important—worth \$5, \$6, or even \$7. And the monopolist maximizes profits from those consumers—fewer buyers, but higher margins. That is one problem: To the extent these higher rates exaggerate wealth inequity by consolidating capital and resources in a single entity, or limit access to a critical facility to only a selected portion of the population, or are otherwise extractive, monopoly pricing can undermine important distributive values.⁶⁹

67. My explanation of the basic economics here is, by necessity, a simplification, ignoring, e.g., the externalities of railway pollution. I put such complications to one side for now in order to explain one view of the basic premises underlying natural monopoly regulation. Regulators have sometimes directly addressed externality-related concerns in natural monopoly markets (among others), and I discuss such possibilities in the machine learning context *infra* notes 235, 238–39, 246–47 and accompanying text.

68. Cf. David Singh Grewal & Jedediah Britton-Purdy, *Liberalism, Property, and the Means of Production*, LPE PROJECT (Jan. 25, 2021), <https://lpeproject.org/blog/liberalism-property-and-the-means-of-production> [<https://perma.cc/V3EA-G6NS>] (describing the view that, in markets characterized by market power, “there are many possible [price] equilibria, none of them determinate”).

69. See, e.g., TIM WU, *THE CURSE OF BIGNESS: ANTITRUST IN THE NEW GILDED AGE* 33–44 (2018); see also Grewal & Britton-Purdy, *supra* note 68 (describing “the problem that scale economies presented to a conception of freedom and equality”). This may be especially true in classical natural monopoly markets. In competitive markets, “[t]he opportunity to charge monopoly prices”—i.e., the desire to dominate the market—supplies a powerful incentive for firms to deploy the “business acumen” that gives rise to “innovation and economic growth” that may help offset these distributive concerns (assuming these concerns are commensurable at all). Compare *Verizon Commc’ns Inc. v. L. Offs. of Curtis V. Trinko LLP*, 540 U.S. 398, 407 (2004) (distinguishing monopolies that result from such “acumen” from those that result from

Moreover, other consumers might conclude that \$5 is too steep and forgo the purchase. That's another problem. The monopolist's incentives are to maximize profits by charging some price above cost.⁷⁰ But that means that some consumers—even consumers who are willing to pay more than cost—are left out. There is unrealized potential when a consumer is willing to pay \$3 for a \$1 ride—but cannot: \$3 is more than cost, but less than the monopolist's profit-maximizing price.⁷¹ And that has further effects on populations that are left unable to travel—to, say, visit family or seek out new employment opportunities.

Enter rate regulation. Unregulated, a natural monopolist is tempted to charge profit-maximizing prices, leaving some consumers back at the station. So rate regulation attempts to mimic competitive conditions by requiring that the monopolist charges prices, set by regulators, that approach cost (thereby offering a more equal opportunity to travel by rail, or use the telephone, across populations).⁷² Mimicking competitive conditions is important, and so some statutes allow regulators to set rates only in the absence of competition (revoking such authority whenever a provider's prices are already disciplined by “effective competition”).⁷³ In short, the monopolist is required to forgo most monopoly profits in exchange for some legally recognized status as the monopoly provider, more easily facilitated access to necessary inputs (e.g., public rights-of-way), and a regulatorily set reasonable profit.

Assessing such reasonable rates, however, is not simple.⁷⁴ It requires that regulators gain access to information typically exclusive to the monopolist: Regulators need to know how much it costs to operate in the regulated market,

anticompetitive practices), with David Singh Grewal & Jedediah Purdy, *Inequality Rediscovered*, 18 THEORETICAL INQUIRIES L. 61, 78 (2017) (contesting the view that “attempt[s] to use state power to shape economic outcomes directly is an error,” noting that such “arguments . . . ten[d] to conceal the distributive choices inherent in market-making policy”). But because natural monopoly markets will typically tend to a single winner, natural monopolists have a somewhat diminished incentive to bring their acumen to bear on the market, and may thereby have less to offer by way of offsetting innovation and growth. See Posner, *Natural Monopoly*, *supra* note 8, at 548 (suggesting that, in natural monopoly markets, “firms will quickly shake down to one through mergers or failures”).

70. See, e.g., BALDWIN ET AL., *supra* note 56, at 447–48.

71. This is because of the loss in the difference between cost and value (here, \$1 and \$3, respectively), known more formally as surplus. In this example, our consumer would realize \$2 in surplus by paying \$1 for a service that is worth \$3 to her. That loss in surplus is more formally known as deadweight loss. See, e.g., VISCUSI ET AL., *supra* note 15, at 82 (defining deadweight loss).

72. See, e.g., *id.* at 421–22 (explaining that regulators aim “to ensure that the monopolist's revenues equal its costs,” though regulators sometimes prioritize social goals when establishing a rate structure); see also 47 U.S.C. §§ 201–02 (describing the Federal Communications Commission's authority to set reasonable rates); Mary A. Wallace, *Interstate Commerce Commission*, 56 GEO. WASH. L. REV. 937, 938 & n.7 (1988) (explaining that the Interstate Commerce Commission was designed to address “abuses of monopoly power” including “rate abuse”).

73. See 47 U.S.C. § 543(a)(2) (2018) (revoking federal and state rate setting authority wherever “a cable system is subject to effective competition”).

74. See *infra* text accompanying note 76.

they need to estimate necessary upkeep and repairs, they must anticipate changes in input costs (e.g., steel tracks or copper wires), and they account for some profits, too.

Historically, regulators have set such rates using a “rate-of-return” method: After accounting for all relevant past and future costs, the regulators add some reasonable rate-of-return (profit) on those investments and set prices accordingly.⁷⁵ This method, however, is susceptible to manipulation. Since monopolists typically have better information than regulators, they may make unnecessary improvements that inflate actual costs, persuade regulators of inflated estimates of their actual costs, exaggerate their future difficulties, and convince regulators to award higher rates of return.⁷⁶ If any of these strategies are successful, rates increase—yielding profits for the monopolist at the expense of the public regulator’s view of the social good (since, as explained, the higher rates leave some consumers out).

Such difficulties with rate-of-return rate setting led regulators to experiment with a different approach—price-cap regulation. Under this approach, regulators set an upper limit on rates (based on, e.g., historical data or benchmarks derived from other firms) and sometimes lower that ceiling decrementally.⁷⁷ In the first year under price-cap regulation, a monopolist may be allowed to charge \$5; in the second year, \$4; in the third, \$3.50; and so on.⁷⁸ Under this approach, regulated monopolists have comparatively little opportunity to manipulate the baseline rate, and have ever-increasing incentives to improve internal efficiency in order to maximize profits under the price cap. But no matter the method—rate-of-return or price cap—the goal of rate regulation is to prevent the social losses that arise from monopoly pricing.

2. Service Specification

As noted, competition pushes providers in most markets to offer better products at better prices. But natural monopolists face limited, if any, competition. And so, just as rate regulation addresses concerns that prices are not disciplined by competition, service specification aims to address quality-related concerns. Indeed, price-cap regulation itself gives rise to further risks

75. See CHRISTOPHER DECKER, *MODERN ECONOMIC REGULATION: AN INTRODUCTION TO THEORY AND PRACTICE* 104 (2015) (describing rate-of-return regulation); see also NUECHTERLEIN & WEISER, *supra* note 3, at 33–34 (describing rate-of-return regulation).

76. See NUECHTERLEIN & WEISER, *supra* note 3, at 33–34 (describing problems attending to rate-of-return regulation).

77. See DECKER, *supra* note 75, at 114–29 (describing price-cap regulation).

78. See Implementation of Sections of the Cable Television Consumer Protection and Competition Act of 1992: Rate Regulation and Adoption of a Uniform Accounting System for Provision of Regulated Cable Service, 9 FCC Rcd. 4527, 4531 (1994) (report and order) (explaining that “[t]he Commission adopted a benchmark and price cap approach to serve as the primary regulatory mechanism for setting initial regulated rates and for governing rates” for cable companies).

regarding the service quality of natural monopolists: Incentives to improve internal efficiency (and maximize profit under a price cap) might also give rise to incentives to skimp on quality, yielding poor customer experiences or even grave risks to public health and safety.⁷⁹ Monopolists might, for example, forego necessary repairs or prudent investments in safety features both to maximize returns under a price-cap methodology and because they face little competitive pressure to do so. Monopoly regulators have therefore focused their attention not only on price, but also quality.

If it is difficult for regulators to predict, counterfactually, what prices could be under conditions of competition, then it might be near impossible for regulators to anticipate the sorts of dynamic service quality improvements that competition could bring. As a result, service regulation has tended to focus on two primary matters, one more technocratic, the other more concerned with matters of democratic governance. First, service regulation has aimed to ensure some baseline standard of service: Where rate regulation sets a price ceiling, service specification sets a quality floor (sometimes set, again, by reference to relevant benchmarks).⁸⁰ Second, service specification has also tended to focus on issues of public concern—public safety, universal service, and nondiscrimination—that relate to the natural monopoly condition.⁸¹ In the context of communications regulation, for example, this has included rules relating to 911 service and emergency backup power, nudging providers into adopting certain technical standards—for better or for worse. It has also encompassed carrier-of-last-resort requirements (i.e., rules requiring carriers to serve customers that might otherwise be left off a communications network), as well as prohibitions against redlining in the construction of network facilities.⁸² Service rules thus sometimes implement social goals—sometimes as a substitute for market pressures that might normally induce a provider to satisfy such standards, sometimes to ensure that the benefits of market consolidation (e.g., network effects) accrue to the whole population, and sometimes with the effect of entrenching a particular

79. See, e.g., NUECHTERLEIN & WEISER, *supra* note 3, at 34–35 (describing incentives to skimp under price-cap regulation).

80. Cf. Connect America Fund, 26 FCC Rcd. 17663, para. 85 (2011) (report and order) (requiring that subsidized monopoly providers of broadband service meet standards of “reasonable comparability” with competitively offered service).

81. See, e.g., WILCOX, *supra* note 10, at 304, 309 (showing that, even in the earliest days of telephone regulation, municipalities imposed conditions to protect public safety operations).

82. See, e.g., 47 U.S.C. § 541(a)(3) (2018) (prohibiting income-based redlining); Hinton Tel. Co., 30 FCC Rcd. 2308 (2015) (forfeiture order) (fining telephone carrier \$100,000 for failing to properly route 911 calls); Connect America Fund, 26 FCC Rcd. 17663, para. 15 (2011) (report and order) (explaining that carrier-of-last-resort requirements help ensure that subsidized monopoly providers “exten[d] voice and broadband service, both fixed and mobile, where it is lacking, to better meet the needs of their consumers.”).

technical pathway.⁸³ And so service (and rate) regulation is often a participatory process, one that includes both the regulated entity as well as interested members of the public.⁸⁴

Consider, for example, the New York Public Service Commission's exercise of regulatory power over Verizon in the wake of Superstorm Sandy. Sandy destroyed a significant portion of Verizon's landline telephone network on Fire Island, New York.⁸⁵ When Verizon, the monopoly provider of wireline telephone service there, began to rebuild its infrastructure, it sought regulatory approval to replace its copper-based infrastructure with "Voice Link," a wireless substitute. This wireless solution was, from Verizon's perspective, cheaper to deploy and maintain. But the New York Public Service Commission worried that Voice Link was a fundamentally unreliable service—one that fell below a reasonable standard for phone service.⁸⁶ Voice Link was, for example, susceptible to inclement weather and network congestion, resulting in higher volumes of dropped calls (far more so than the previous copper-based service).⁸⁷ Moreover, these quality problems gave rise to specific public safety concerns: "Verizon itself conceded that congestion could block or slow completion of 911 calls via Voice Link," and the Voice Link terms of service purported to limit Verizon's "liability for failing to complete 911 calls."⁸⁸ These limits gave rise to both intense public dissatisfaction with Verizon's plan, as well as severe regulatory scrutiny. Verizon relented: As the Public Service Commission readied to deny Verizon's request, Verizon withdrew its application and agreed to build a wireline fiber-based network on the island.⁸⁹

83. See Rahman, *supra* note 44, at 1637 (describing the moral and social dimensions of progressive-era utility regulation).

84. See, e.g., *Major Rate Case Process Overview*, N.Y. STATE DEP'T OF PUB. SERV. (Sept. 23, 2011, 3:55 PM), <https://www3.dps.ny.gov/W/PSCWeb.nsf/o/364Do704BEEC5B7D85257856006C56B3?OpenDocument> [<https://perma.cc/PUJ3-KPYE>] ("Rate cases are a primary instrument of government regulation of these industries. Interested persons may intervene and become parties in a utility company's rate case. Typical intervenors include: industrial, commercial and other large-scale users of electricity; public interest groups; representatives of residential, low-income and elderly customers; local municipal officials; and, dedicated advocacy groups . . . Rate cases proceed in an entirely public and open process."); see also BREYER, *supra* note 16, at 341, 351–53 (discussing the importance (and shortcomings) of participation and representation to the legitimacy of regulatory action).

85. *The Phone Network Transition: Lessons from Fire Island*, PUB. KNOWLEDGE (Mar. 7, 2014), <https://www.publicknowledge.org/blog/the-phone-network-transition-lessons-from-fire-island> [<https://perma.cc/U33M-YQRB>].

86. See Letter from State of N.Y. Dep't of Pub. Serv., to Marlene H. Dortch, Fed. Comm'ns Comm'n 2 (July 29, 2013), <https://ecfsapi.fcc.gov/file/7520934180.pdf> [<https://perma.cc/79RY-JWBL>].

87. *Id.*

88. Tejas N. Narechania & Erik Stallman, *Internet Federalism*, 34 HARV. J.L. & TECH. 547, 580 (2021).

89. Patrick McGeehan, *Verizon Backing Off Plans for Wireless Home Phones*, N.Y. TIMES (Sept. 12, 2013), <https://www.nytimes.com/2013/09/13/nyregion/verizon-abandons-plans-for-wireless-home-phones-in-parts-of-new-york.html> [<https://perma.cc/28CQ-947X>].

In short, the regulator's authority over communications services providers encompassed some power to dictate the terms of the service provided, including standards related to matters of both private concern (such as call quality) and public interest (such as emergency communications).

3. Induced Competition

Both rate regulation and service specification aim, in part, to substitute for the benefits—better prices for better products—that typically come by way of market competition, on the assumption that competition is unlikely or undesirable inside natural monopoly markets. But some modes of regulation aim to induce efficient competition in these markets all the same.

The tradition of utility franchising offers one important, longstanding example.⁹⁰ Rather than attempt to directly mimic the outcomes of intramarket competition, regulators control market entry by establishing an arranged competition: Several firms, each vying to serve some market, submit offers to provide their best service at their best price, and whichever firm offers the best value is endowed with a monopoly franchise—essentially, the license legally required to serve the area.⁹¹ Harold Demsetz, who was among the first to argue in favor of franchising over direct regulation, reasoned that competition among bidders could achieve the same ends as rate regulation and service specification, but without the substantial regulatory structure that typically attended to natural monopoly regulation.⁹²

In other contexts, policy makers have lowered barriers to market entry by, for example, allowing new competitors to share in existing investments, thereby reducing duplicative fixed investments and the concomitant consumer price increases.⁹³ In short, the rules of the market are specially structured to sustain competition. Consider the Telecommunications Act of

90. See WILCOX, *supra* note 10, at 28–33.

91. Paul L. Joskow, *Regulation of Natural Monopolies*, in 2 HANDBOOK OF LAW AND ECONOMICS 1227, 1244–47 (A. Mitchell Polinsky & Steven Shavell eds., 2007); Robin A. Prager, *Franchise Bidding for Natural Monopoly: The Case of Cable Television in Massachusetts*, 1 J. REGUL. ECON. 115, 126–27 (1989). *But see* Oliver E. Williamson, *Franchise Bidding for Natural Monopolies—in General and with Respect to CATV*, 7 BELL J. ECON. 73, 76–78 (1976) (describing the difficulties that attend to evaluating ex ante what constitutes the best value).

Notably, a state's decision to award an exclusive franchise to a natural monopolist (or a locality's similar decision, made under duly delegated authority) has traditionally posed no antitrust law concerns. See *Cnty. Commc'ns Co. v. City of Boulder*, 455 U.S. 40, 54–56 (1982) (citing *City of Lafayette v. La. Power & Light Co.*, 435 U.S. 389, 413 (1978)); *see also* *Town of Hallie v. City of Eau Claire*, 471 U.S. 34, 43–47 (1985) (finding that a delegation of sewage services by the state and cities was not an antitrust concern).

92. See Harold Demsetz, *Why Regulate Utilities?*, 11 J.L. & ECON. 55, 63–65 (1968); *see also* VISCUSI ET AL., *supra* note 15, at 469 (suggesting that “franchise bidding [could] result[] in average cost pricing and the most efficient firm operating” with the added benefit “that [bidding] imposes no informational requirements on a government agency”).

93. See *supra* notes 58–59 and accompanying text (explaining that duplicative fixed investments lead to higher consumer and input prices).

1996. Before the 1996 Act, policy makers had long regulated local communications markets as natural monopolies by setting rates and imposing service conditions (such as carrier-of-last-resort obligations).⁹⁴ But the 1996 Act sought to change the dynamics of entry into these local markets by giving putative competitors various rights to lease capacity on existing telephone networks.⁹⁵ Competitors could thus avoid building out duplicative networks before offering competing service, relying instead on the existing infrastructure (until those new providers found it more desirable to invest in new facilities—perhaps when existing infrastructure could no longer support the total demand for communications services, thus requiring new investments).⁹⁶

Similarly, antitrust's essential facilities doctrine requires that natural monopolists make their facilities available to competitors in adjacent markets: Companies that own power lines, for example, must make those lines available to a range of electricity producers (e.g., power plants, wind farms), so that competition among producers may thrive without requiring that each build its own transmission grid.⁹⁷

Just as both rate regulation and service specification have proved to be imperfect substitutes for market competition in some respects, so too have modes of inducing competition sometimes come up short. Franchising, for example, did not upend all other natural monopoly regulation (as Demsetz originally envisioned).⁹⁸ Commentators responding to Demsetz, such as Oliver Williamson, reasoned that natural monopolists might behave opportunistically after having won a franchise, charging prices or offering service inconsistent with their promises.⁹⁹ And public choice concerns likely

94. NUECHTERLEIN & WEISER, *supra* note 3, at 40–44.

95. 47 U.S.C. §§ 251–252 (2018); *see also* NUECHTERLEIN & WEISER, *supra* note 3, at 52 (briefly summarizing Congress's goals for these new provisions).

96. *See* NUECHTERLEIN & WEISER, *supra* note 3, at 58–59; *see also supra* notes 58–59 and accompanying text (explaining that natural monopolies do not exist—that competition is sustainable—where costs do not fall over the relevant market, but rather new fixed investments are required as demand grows).

97. *Otter Tail Power Co. v. United States*, 410 U.S. 366, 378 (1973) (“When a community serviced by Otter Tail decides not to renew Otter Tail’s retail franchise when it expires, it may generate, transmit, and distribute its own electric power. We recently described the difficulties and problems of those isolated electric power systems. Interconnection with other utilities is frequently the only solution There were no engineering factors that prevented Otter Tail from . . . wheeling the power Otter Tail’s refusals to . . . wheel were solely to prevent municipal power systems from eroding its monopolistic position.” (citations omitted)).

98. *See* Priest, *supra* note 7, at 304–08 (explaining that Demsetz’s view of franchising presupposed the existence of the very regulatory apparatus he sought to undermine).

99. *See* Williamson, *supra* note 91, at 76–79; *see also* Robin A. Prager, *Firm Behavior in Franchise Monopoly Markets*, 21 RAND J. ECON. 211, 223–24 (1990) (finding evidence of opportunistic behavior that varies in severity across sectors).

played a role in entrenching existing regulators, too.¹⁰⁰ Hence, even after monopolists were selected competitively, they remained subject to regulatory supervision to ensure that providers adhered to their franchising commitments, and to negotiate the terms of the provider's ongoing franchise.¹⁰¹ In all, even competitively designated franchisees remained subject to rate and service supervision.

Commentators likewise "agree that the [1996 Act] failed to fulfill its promise of leading to more vibrant and competitive markets," casting blame on both incumbent providers, who resisted the legal mandates to share their infrastructure, and on the new competitors, who often failed to make the hoped-for investments in new communications facilities (riding instead on the incumbents' facilities).¹⁰²

* * *

Industries characterized predominantly by scale can sometimes present natural monopolies, where such scale economies extend over the entire relevant market. Such monopolists—including, historically, railways, electric grid operators, and telephone carriers—have traditionally been subject to regulatory supervision over rates and standards. Indeed, such regulation is sometimes carried forward into other monopolistic or oligopolistic markets (even if the natural monopoly condition is not strictly met). Underlying this regulatory scrutiny is a view that something is necessary to discipline a dominant provider's prices and services—but, in these markets, regular competition will not do it. And so state and federal authorities substituted competition with regulatory control over rates and service, thereby mimicking the effects of competition while also retaining the gains of consolidation and implementing other important social goals, or by restructuring the market to sustain competition. These regulatory substitutes were sometimes imperfect. But regulators and policy makers—responsible to elected officials and a voting public—seemed to conclude that the costs of such regulatory shortcomings were less than the costs of monopoly control, and that public accountability for these decisions (and missteps) was better than private obscurity.¹⁰³ However, once newer technological advances proved these

100. See Priest, *supra* note 7, at 293; cf. NUECHTERLEIN & WEISER, *supra* note 3, at 11–12 (discussing public choice theory).

101. See Victor P. Goldberg, *Regulation and Administered Contracts*, 7 BELL J. ECON. 426, 444–46 (1976); see, e.g., *Denver Area Educ. Telecomm. Consortium, Inc. v. FCC*, 518 U.S. 727, 772 (1996) (Stevens, J., concurring) (“[P]ublic, educational, and governmental access channels . . . owe their existence to contracts forged between cable operators and local cable franchising authorities.”)

102. Stephen Labaton, *Phone Start-Ups Win the Latest Round in Court*, N.Y. TIMES (May 14, 2002), <https://www.nytimes.com/2002/05/14/business/phone-start-ups-win-the-latest-round-in-court.html> [https://perma.cc/2QZM-GEJ9].

103. See BREYER, *supra* note 16, at 357–60.

markets capable of sustaining (evermore) competition on their own—once, for example, the railroads faced competition from an array of long-haul truckers (among other newly developed transportation options), or once wireless telephone service proved to be an adequate substitute for wireline service—these industries were slowly deregulated, perhaps, in part, because this error cost analysis changed.¹⁰⁴

III. MACHINE LEARNING AS NATURAL MONOPOLY

Many markets historically regulated as natural monopolies—the railroads and the telephone, among others—have since proved capable of sustaining competition. But that need not imply the theory of natural monopoly regulation set out above is dead. Other natural monopolies may surface, and hence deserve regulatory attention (even assuming that these prior examples are now, in view of technological change, capable of sustaining competition).

Indeed, scholars across disciplines have recently hinted at the possibility of a newly emergent natural monopoly (or, more accurately, series of natural monopolies) in applications of machine learning technologies.¹⁰⁵ Oren Bracha and Frank Pasquale, for example, have suggested that the search engine market exhibits the characteristics of a natural monopoly.¹⁰⁶ Likewise, Steven Weber and Gabriel Nicholas have identified “a tendency toward natural monopolies in data platform businesses.”¹⁰⁷ Of course, neither search engines nor data platforms necessarily implement their services by way of machine learning algorithms: Google’s search results, for example, are determined by a range of factors, only some of which require machine learning.¹⁰⁸ But Steven Weber and Gabriel Nicholas focus their attention on

104. See Wallace, *supra* note 72, at 938 (railroads); NUECHTERLEIN & WEISER, *supra* note 3, at 137–41 (telephones); see also BREYER, *supra* note 16, at 315–21 (airlines); KAHN, *supra* note 16, at xv–xvii (electricity and gas); Biber et al., *supra* note 2, at 1591–92 (power utilities); Levine, *supra* note 13, at 191–93 (airlines).

As I noted above, I readily acknowledge this is only one among several plausible explanations for the deregulatory movement of the late twentieth century. It does not undermine my point here that some other theory also helps explain changes in the regulatory regime. See *supra* note 12 and accompanying text.

105. In addition to the sources described *infra* notes 106–07, see the sources cited *supra* note 18 and accompanying text.

106. Bracha & Pasquale, *supra* note 18, at 1180–81 (“It is unclear whether search engines fall under the strict definition of a natural monopoly, but they exhibit very similar characteristics.” (footnote omitted)).

107. Weber & Nicholas, *supra* note 18, at 25.

108. See Jack Clark, *Google Turning Its Lucrative Web Search Over to AI Machines*, BLOOMBERG (Oct. 26, 2015), <https://www.bloomberg.com/news/articles/2015-10-26/google-turning-its-lucrative-web-search-over-to-ai-machines> [https://perma.cc/YRS3-G5CA]; see also LAWRENCE PAGE, SERGEY BRIN, RAJEEV MOTWANI & TERRY WINOGRAD, *THE PAGERANK CITATION RANKING: BRINGING ORDER TO THE WEB* 4–10 (1998), <http://ilpubs.stanford.edu:8090/422/1/1999->

machine learning systems in particular. And both sets of authors identify data feedback loops—one facet of machine learning (sometimes termed continual learning)—as critical to their findings.¹⁰⁹ Continual learning, however, is only part of the story, since such positive feedback loops, standing alone, need not necessarily imply natural monopoly conditions.¹¹⁰

Hence, I closely analyze the claim that machine learning systems are natural monopolies. I begin with a brief description of machine-learning-based applications, and then turn to a general (if informal) analysis of the natural monopoly claim focusing on three primary features. First, I describe the scale economies that inhere to the development of such systems. Second, I describe the ways in which the costs associated with training certain machine-learning-based applications appear to be, from a computational perspective, subadditive—a sufficient condition for natural monopoly.¹¹¹ And third, I elaborate on the positive feedback effects of continual learning. Viewed together, I find it likely that, for at least some applications, a natural monopoly exists.

A. UNDERSTANDING MACHINE LEARNING

Computers do what they are instructed to do. As I elaborate *infra*, a software developer might, for example, ask a computer to sort an array of numbers—and, in so doing, give the computer explicit directions on how, exactly, to sort that list.¹¹² That’s useful because the computer now has some new functionality—each new set of instructions provides a new tool: a list sorter, a number adder, a game player, and so on.

But in each of these examples, the software’s output is some *thing*—a sorted list, a numerical sum, or a path to victory over Bowser.¹¹³ Machine

66.pdf [<https://perma.cc/SRA4-LFVN>] (describing Google’s original ranking methodology, which did not employ machine learning, but rather weighted link analysis).

109. See Weber & Nicholas, *supra* note 18, at 25–26 (describing the “positive feedback loop that creates a tendency toward natural monopolies in data platform businesses”); Bracha & Pasquale, *supra* note 18, at 1178–80 (citing, among others, Pasquale, *supra* note 34, at 130, which describes a similar feedback effect in searches).

110. See *infra* notes 176–77 and accompanying text (describing the difference); see also, e.g., Mark A. Lemley & David McGowan, *Legal Implications of Network Economic Effects*, 86 CALIF. L. REV. 479, 595–96 (1998) (drawing distinction between supply-side natural monopoly effects and demand-side network effects); David McGowan, *Networks and Intention in Antitrust and Intellectual Property*, 24 J. CORP. L. 485, 488 (1999) (explaining how network theory does not necessarily imply a monopoly); *David L. Aldridge Co. v. Microsoft Corp.*, 995 F. Supp. 728, 754 (S.D. Tex. 1998) (describing how natural monopolies can occur in industries that have a large initial fixed cost but also have declining marginal production costs); CARL SHAPIRO & HAL R. VARIAN, *INFORMATION RULES: A STRATEGIC GUIDE TO THE NETWORK ECONOMY* 13–14 (1998) (discussing the relationship between positive feedback and network effects).

111. See, e.g., Baumol, *Quasi-Permanence*, *supra* note 22, at 12–18.

112. See *infra* notes 152–54 and accompanying text.

113. Bowser is a primary adversary in the video game franchise *Super Mario Bros.* See, e.g., *SUPER MARIO BROS.* 3 (Nintendo 1990).

learning algorithms operate differently: With machine learning, the output is itself “another algorithm that can be directly applied to further data.”¹¹⁴ In this way, machine learning offers a sort of “meta-algorithm,” one that draws on existing data to derive a new prediction model.¹¹⁵

An illustration may help clarify. In *The Ethical Algorithm*, Michael Kearns and Aaron Roth develop an example related to high school performance and college graduation: Given an array of students’ high school coursework, grades, test scores, and college graduation statuses, we can calculate a number of functions: the average test score for graduates and nongraduates, the frequency with which graduates and nongraduates took certain classes, and so on.¹¹⁶ But using that database to manually “derive a model [for] predicting the likelihood of graduation for future applicants” could be “quite difficult and subtle.”¹¹⁷ A machine-learning-based application, however, might absorb that mass of data—information on hundreds or thousands (or more) student profiles—and thereby derive a model for predicting the likelihood of graduation for some *new* applicant, i.e., a student absent from the original dataset. Once the application has learned which students graduated and which did not, it can attempt to predict whether other high school students are likely to also graduate from college.¹¹⁸

Such applications are hardly limited to this college graduation example. Machine learning algorithms have been deployed to help solve a wide range of complex problems operating across distinct markets, including medicine, law enforcement, and online search.¹¹⁹

W. Nicholson Price, among others, has detailed how medical professionals use machine-learning-based applications to diagnose ailments and to develop treatment plans for patients.¹²⁰ Such applications, “derived from large datasets[,] . . . reflect complex underlying biological relationships” and are used to diagnose and treat a range of ailments and conditions—from migraines

114. KEARNS & ROTH, *supra* note 51, at 6.

115. *Id.*

116. *See id.*

117. *Id.*

118. *Cf.* Tom Simonite, *Meet the Secret Algorithm That’s Keeping Students Out of College*, WIRED (July 10, 2020, 7:00 AM), <https://www.wired.com/story/algorithm-set-students-grades-altered-futures> [<https://perma.cc/gTCN-XUCD>] (describing the International Baccalaureate program’s decision to cancel testing during the COVID-19 crisis and replace actual test scores with algorithmically-predicted results, to much controversy).

119. *See, e.g., supra* note 5 and accompanying text; Emmanuel Gbenga Dada, Joseph Stephen Bassi, Haruna Chiroma, Shafi’I Muhammad Abdulhamid, Adebayo Olusola Adetunmbi & Opeyemi Emmanuel Ajibuwa, *Machine Learning for Email Spam Filtering: Review, Approaches and Open Research Problems*, HELIYON, June 2019, at 1, 2 (“To effectively handle the threat posed by email spams, leading email providers such as Gmail, Yahoo mail and Outlook have employed the combination of different machine learning (ML) techniques such as Neural Networks in its spam filters.”).

120. *See, e.g.,* W. Nicholson Price II, *Regulating Black-Box Medicine*, 116 MICH. L. REV. 421, 423 (2017); *see also* Rai et al., *supra* note 32, at 2 (examining “AI-enabled clinical decision software”).

to autism and beyond.¹²¹ Specifically, these algorithms rely on the stores of data contained in medical health records to identify both patterns suggestive of a particular illness and treatments likely to succeed: Some providers claim “to use ‘170 million data points to “transform” data from insurance claims, electronic medical records, medical sensors and other sources into information . . . to predict the best ways to treat individual patients and conditions.”¹²² Google’s Project Nightingale, for example, draws from the medical records “of millions of people across 21 states” and may incorporate data from personal health tracking devices such as Fitbit.¹²³ Many such algorithms are “locked”—i.e., once trained on an existing dataset, they do not incorporate new information from their in-market use.¹²⁴ But while most FDA-approved applications are, so far, locked, developers continue to experiment with continual learning models that learn from the new outcomes generated by the application’s use in medical settings.¹²⁵

Likewise, facial recognition technology has, among other uses, gained significant traction in law enforcement, using machine-learning-based advances to identify criminal suspects. Using collections of digital photos, these applications learn to recognize the visual identity of a person through identified photographs—photos tagged in Facebook, associated with a drivers’ license, or taken while processing an arrest—and use that information to predict the identity of someone in a new image—say, a frame from a security camera recording. Many of these systems do continually learn and adapt, both by incorporating new known images as they become available and by learning from erroneous outcomes.¹²⁶

121. Price II, *supra* note 120, at 430.

122. *Id.* at 427 (quoting Chris Rauber, *Lumiata Nabs \$6 Million for Personalized Medical Care Software*, S.F. BUS. TIMES (Sept. 11, 2014, 8:38 AM), <http://www.bizjournals.com/sanfrancisco/blog/2014/09/lumiata-6-million-funding-personalized-health-data.html> [<https://perma.cc/3PEG-FQZV?type=image>]).

123. Rob Copeland, *Google’s ‘Project Nightingale’ Gathers Personal Health Data on Millions of Americans*, WALL ST. J. (Nov. 11, 2019, 4:27 PM), <https://www.wsj.com/articles/google-s-secret-project-nightingale-gathers-personal-health-data-on-millions-of-americans-11573496790> [<https://perma.cc/E5EV-J8RK>].

124. Rai et al., *supra* note 32, at 5.

125. Price II, *supra* note 120, at 427–28; Rai et al., *supra* note 32, at 8.

126. Apple’s Face ID technology, for example:

[I]mprove[s] match performance and keep[s] pace with the natural changes of a face and look, [by] augment[ing] its stored mathematical representation over time. Upon a successful match, Face ID may use the newly calculated mathematical representation—if its quality is sufficient—for a finite number of additional matches before that data is discarded. Conversely, if Face ID fails to recognize a face but the match quality is higher than a certain threshold and a user immediately follows the failure by entering their passcode, Face ID takes another capture and augments its enrolled Face ID data with the newly calculated mathematical representation. This new Face ID data is discarded if the user stops matching against it or after a finite number of matches. These augmentation processes allow Face ID to keep up with

And a vast range of internet companies use machine-learning techniques to provide consumers with search results and recommendations.¹²⁷ Search engines, for example, use machine-learning-based applications, trained on text, to better understand search queries and identify relevant results—by, say, identifying synonyms and antonyms, and by using metrics suggestive of a successful (or not) search result to improve training over time.¹²⁸ Similarly, companies like Netflix and Spotify use machine-learning-based applications to recommend new content to consumers based on their knowledge of every user’s content preferences (to derive a prediction model), and each user’s individual content consumption habits (to offer tailored predictions).¹²⁹

In total, machine-learning-based applications—together with the volumes of data now available—are remaking operations across sectors.¹³⁰ To some, these new tools are a welcome development: College admissions officers, for example, might be interested in tools to help better assess students’ applications. Patients want better health outcomes. Law enforcement officials want to reliably identify criminal suspects. And most consumers would like Netflix to reliably recommend good content. But such software may also seem potentially problematic, implicating concerns regarding recommendation accuracy and bias, the incentives of the application’s developer, and user privacy, among others. Before addressing such concerns, I examine the nature of these machine-learning-based applications in more detail, focusing on the possibility that any one application may be a natural monopolist in its own market, whether medicine, law enforcement, or search, among other possibilities.

dramatic changes in a user’s facial hair or makeup use while minimizing false acceptance.

APPLE, APPLE PLATFORM SECURITY 24 (2021), https://manuals.info.apple.com/MANUALS/1000/MA1902/en_US/apple-platform-security-guide.pdf [<https://perma.cc/DK7N-XS75>].

127. See, e.g., Clark, *supra* note 108.

128. For example, the verb “execute” means one thing in computer science contexts, and quite another thing in legal contexts. So a search application may begin by thinking that “execute” means “kill” or “terminate”—but learn, over time, that when associated with other computer science terms (e.g., program, computer), it should use different synonyms, like “run” or “start.” See, e.g., Kyle Wiggers, *Search Engines Are Leveraging AI to Improve Their Language Understanding*, VENTUREBEAT (May 20, 2020, 8:40 AM), <https://venturebeat.com/2020/05/20/search-engines-google-bing-are-leveraging-ai-to-improve-their-language-understanding> [<https://perma.cc/W4EG-23N6>].

129. See, e.g., *Machine Learning*, SPOTIFY, <https://research.atspotify.com/machine-learning> [<https://perma.cc/RHC3-98HJ>] (“Machine learning touches every aspect of Spotify’s business. It is used to help listeners discover content via recommendations and search . . .”).

130. RUSSELL, *supra* note 4, at 62–101 (cataloguing near-term and long-term changes based on artificial intelligence and suggesting that, cumulatively, “[t]hese developments . . . could change the dynamic of history”); RUSSELL & NORVIG, *supra* note 6, at 28 (explaining that “many thousands of AI applications are deeply embedded in the infrastructure of every industry” (quoting KURZWEIL, *supra* note 6, at 204)).

B. MACHINE LEARNING AS NATURAL MONOPOLY

At least three features of machine-learning-based applications are suggestive of natural monopoly conditions (or natural monopoly effects): the costs of application development; the costs of training and optimization; and the potential for network effects. I consider each in turn below.

1. Application Development

Software markets mimic natural monopolies in at least one important respect: It can cost a lot to develop software, but it is comparatively trivial to distribute that software to consumers.¹³¹ Developing software can be an exceptionally complex task, requiring significant resources in terms of expertise and labor (i.e., software developers), as well as computing power.¹³² But once a program is built, it is easy to burn copies onto discs or, more commonly (and more cheaply), transmit them over the internet. Because software has such relatively high fixed and low marginal costs, it has—roughly consistent with the traditional understanding of natural monopolies—declining average marginal costs.¹³³

Machine learning presents an extreme example of this dynamic, as the fixed costs that attend to developing a machine-learning-based application are significantly greater than those associated with the average software development project (while marginal costs remain comparatively low).

First, for so long as machine learning remains a relatively new and novel technology, we might expect that the labor costs for machine learning engineers—their living wages—to be higher, given the demand for their specialized expertise.¹³⁴

131. See, e.g., *Microsoft Corp. v. AT&T Corp.*, 550 U.S. 437, 439 (2007) (“[C]opying software . . . is indeed easy and inexpensive . . .”).

132. See, e.g., Bryan H. Choi, *Software as a Profession*, 33 HARV. J.L. & TECH. 557, 572 (2020) (describing “the exponential growth of software complexity”).

133. See David McGowan, *Regulating Competition in the Information Age: Computer Software as an Essential Facility Under the Sherman Act*, 18 HASTINGS COMM’NS & ENT. L.J. 771, 813, 841 (1996); see also *supra* Part II (explaining the importance of scale to the definition of natural monopoly). But see Mark A. Lemley, *Antitrust and the Internet Standardization Problem*, 28 CONN. L. REV. 1041, 1056 (1996) (“The Internet software industry is not a natural monopoly”). I explore this apparent disconnect further *infra* notes 176–82 and accompanying text.

134. E.g., Stacy Stanford, *Artificial Intelligence (AI): Salaries Heading Skyward*, MEDIUM: TOWARDS AI (May 20, 2020), <https://medium.com/towards-artificial-intelligence/artificial-intelligence-salaries-heading-skyward-e41b2a7bba7d> [<https://perma.cc/L6Q8-YF4B>]; Louis Columbus, *Machine Learning Engineer Is The Best Job In The U.S. According To Indeed*, FORBES (Mar. 17, 2019, 10:35 AM), <https://www.forbes.com/sites/louiscolumnbus/2019/03/17/machine-learning-engineer-is-the-best-job-in-the-u-s-according-to-indeed/?sh=59d5879d7bbo> [<https://perma.cc/Q3XR-JRG8>].

Though I describe machine learning as “new” and “novel” above, it is worth noting that the premises that underlie modern machine learning technology have been long in development. See RUSSELL & NORVIG, *supra* note 6, at 16–28 (describing the history of artificial intelligence, beginning in 1943). Despite this long history, only recent advances in computer

But even if this effect is only transitory, other more durable aspects of machine learning development exaggerate the differences between fixed and marginal costs. In addition to standard software development costs, machine-learning-based applications rely on vast volumes of data—data used to train the learning algorithm to, say, recognize someone’s face or identify a potentially successful therapeutic regimen. In some cases, such data may be freely (if controversially) available: IBM, for example, scraped millions of publicly available images from image-sharing website Flickr to train its then-nascent facial recognition software, and Clearview AI employed a similar strategy, drawing criticism from privacy advocates.¹³⁵ Nevertheless, good data is typically expensive, no matter whether the proprietor chooses to “build it” or “buy it.”¹³⁶

Moreover, the physical computing resources required to design and train a machine learning algorithm can be expensive, requiring specialized equipment that can cost millions of dollars.¹³⁷ Some applications—particularly

processing and the availability of large data sets have made possible the practical and increasingly ubiquitous applications described here. *See id.* at 27–28; *see also* Amy Kapczynski, *The Law of Informational Capitalism*, 129 *YALE L.J.* 1460, 1468 (2020) (describing the recent developments that have enabled machine-learning-based advances).

135. *See, e.g.*, Olivia Solon, *Facial Recognition’s ‘Dirty Little Secret’: Millions of Online Photos Scraped Without Consent*, *NBC NEWS* (Mar. 17, 2019, 10:25 AM), <https://www.nbcnews.com/tech/internet/facial-recognition-s-dirty-little-secret-millions-online-photos-scraped-n981921> [<https://perma.cc/9WNY-YR8A>]; Kashmir Hill, *The Secretive Company That Might End Privacy as We Know It*, *N.Y. TIMES* (Nov. 2, 2021), <https://www.nytimes.com/2020/01/18/technology/clearview-privacy-facial-recognition.html> [<https://perma.cc/4LFZ-4GSD>].

These examples underscore the importance of other features of machine learning development, in addition to data. While obtaining training data is important, obtaining *good* training data can be more difficult and costly. And even if one regards publicly available digital images as good training data, the computing infrastructure required to train a deep learning facial recognition system is extensive and costly. *Cf. infra* notes 137 & 157 (discussing, respectively, the infrastructural and computational costs associated with training).

136. Levendowski, *supra* note 49, at 606–10 (explaining that acquiring data “to use as training data for AI systems can be exhaustingly resource intensive,” no matter whether it is built (by, say, a social media platform) or purchased (from specialized data keepers, such as hospitals as to patient case histories)); *see* C. Scott Hemphill, *Disruptive Incumbents: Platform Competition in an Age of Machine Learning*, 119 *COLUM. L. REV.* 1973, 1978–79 (2019).

It matters, of course, that data is different from other inputs in that data is not strictly rivalrous. When one company uses certain information to train its machine-learning-based applications, it does not diminish that data’s availability for future competitors. But, as other scholars have noted, intellectual property rights over these data can mimic the effects of rivalry, limiting their use to particular rightsholders. *See, e.g.*, Lemley & Casey, *supra* note 49, at 756–70; Levendowski, *supra* note 49, at 590–610. And intellectual property rights are not the only concerns at issue. I explore this further *infra* Section IV.B.3.

137. *See* Edd Gent, *Powerful AI Can Now Be Trained on a Single Computer*, *IEEE SPECTRUM* (July 17, 2020), <https://spectrum.ieee.org/tech-talk/artificial-intelligence/machine-learning/powerful-ai-can-now-be-trained-on-a-single-computer> [<https://perma.cc/DXZ7-D6DV>] (explaining that some training approaches “have relied on servers packed with hundreds of CPUs and GPUs” while others use “[s]pecialized hardware . . . with a price tag running into the millions”); *see also* Hemphill, *supra* note 136, at 1978 (2019) (discussing how Google and other companies have invested

those developed using relatively simple prediction algorithms (on which I elaborate in the next Section) can be developed on commodity hardware—the sorts of regular computers that consumers might purchase for home use.¹³⁸ But others, such as GPT-3—a popular natural language process application—require a massive computing infrastructure that is likely out of reach of new competitors.¹³⁹

In short, human, data, and hardware resources are all more expensive in machine learning contexts (as compared to the average software development project), yielding an even wider disparity between fixed development costs and marginal distribution costs suggestive of natural monopoly conditions.¹⁴⁰

In the general software context, scholars and commentators have generally concluded that these sorts of scale economies are, standing alone, insufficient to give rise to a natural monopoly in need of regulatory intervention (even if natural monopoly effects arise). This is for two reasons. First, though the fixed costs of software development are much higher than the marginal costs of software distribution, those fixed costs are, in absolute terms, not all that high. Compared, say, to the fixed costs of building a power plant and an electric grid, developing a software program may seem trivial.¹⁴¹ And so competition in these software markets, even if properly characterized as natural monopolies, yields less waste than in traditionally regulated

in custom hardware to support machine-learning) (citing Norman P. Jouppi, Cliff Young, Mishant Patil & David Patterson, *Motivation for and Evaluation of the First Tensor Processing Unit*, IEEE MICRO, May/June 2018, at 10, 14 tbl.2, 16 tbl.4).

138. See Aleksei Petrenko, Zhehui Huang, Tushar Kumar, Gaurav Sukhatme & Vladlen Koltun, *Sample Factory: Egocentric 3D Control from Pixels at 100000 FPS with Asynchronous Reinforcement Learning*, PROCEEDINGS OF THE 37TH INTERNATIONAL CONFERENCE ON MACHINE LEARNING (2020); see also Gent, *supra* note 137 (explaining that the Petrenko et al. paper, *supra*, “describe[s] how [the authors] were able to use a single high-end workstation to train AI . . . using a fraction of the normal computing power”).

139. BROWN ET AL., *supra* note 30, at 39 (describing the computational costs of training GPT-3); see also Kyle Wiggers, *OpenAI’s Massive GPT-3 Model Is Impressive, But Size Isn’t Everything*, VENTURE BEAT (June 1, 2020, 1:05 PM), <https://venturebeat.com/2020/06/01/ai-machine-learning-openai-gpt-3-size-isnt-everything> [<https://perma.cc/8X75-L2EB>] (estimating that GPT-3 required “over 350GB of memory and \$12 million in compute credits,” which was feasible because it was developed by a “a well-capitalized company that teamed up with Microsoft to develop an AI supercomputer,” but is “potentially beyond the reach of AI startups . . . which in some cases lack the capital required”).

140. See Hemphill, *supra* note 136, at 1977 (“Advances in machine learning . . . ha[ve] a high fixed cost and low marginal cost . . .”); cf. DUCCI, *supra* note 18, at 57 (“It is therefore plausible that cost subadditivity in delivering a general search service applies over the relevant demand output region and average cost may be decreasing over the same range, which makes the presence of a single dominant search engine the most cost-efficient way to serve the market. The cost structure is indeed similar to traditional natural monopolies.”).

141. Cf. Transcript of Oral Argument at 5, *Alice Corp. Pty. Ltd. v. CLS Bank Int’l*, 573 U.S. 208 (2014) (No. 13-298), 2014 WL 1279672 (statement of Kennedy, J.) (suggesting that it is trivial for an inexperienced software developer to write a program covered by the challenged (and ultimately invalid) patent).

markets.¹⁴² Moreover, the possibility of a longer, ongoing competition for the market (facilitated by these comparatively lower barriers to entry) yields some important quality gains in the software product.¹⁴³ Altogether, scholars have concluded that the possibility of quality gains through competition (together with avoiding the downside risk of regulatory error) seem, in this context, to outweigh the efficiency gains of avoiding competition in a natural monopoly (or natural-monopoly-like) space.¹⁴⁴

But these justifications operate with less force in the machine learning context. As noted, the fixed data and equipment costs that attend to developing machine-learning-based applications are significantly higher than for traditional software development, thus suggesting both that the losses from wasteful competition are higher and that the possible gains from competition may be lower (due to the higher barriers to market entry).¹⁴⁵

In total, the fixed costs of developing a single machine-learning-based application for a single market (e.g., facial recognition or diagnostics) seem to significantly outpace the marginal costs of making those applications available to its consumers (e.g., law enforcement or health care providers), possibly giving rise to a natural monopoly in that application's specific market.¹⁴⁶ After all, one provider is likely to be able to serve the market—several law enforcement agencies, or several hospitals—without needing to build additional capacity after designing, developing, and training the application. And so, in such cases, a second provider will add no new capacity to the market but will duplicate these fixed costs—thereby increasing the average marginal cost across that market.¹⁴⁷

I am careful not to overstate this claim. As I explain *infra*, it may be possible to reduce these fixed costs and avoid the natural monopoly condition by, say, relying on shared infrastructure or public datasets.¹⁴⁸ Moreover, some of these fixed investments may yield improvements across markets or sectors, i.e., spillovers. Competitive investments in machine learning hardware may,

142. See Lemley, *supra* note 133, at 1059.

143. See *id.* at 1058–59; cf. Demsetz, *supra* note 92, at 57 n.7 (“[C]ompetition for the [market]’ [can mimic effects of] ‘competition within the [market]’” (emphasis removed) (quoting Edwin Chadwick, *Results of Different Principles of Legislation and Administration in Europe; of Competition for the Field, as compared with the Competition within the Field of Service*, 22 J. ROYAL STAT. SOC’Y. 381 (1859))).

144. See *supra* note 76 (describing some information asymmetries between providers and regulators that can give rise to regulatory error).

145. Hemphill, *supra* note 136, at 1975–81 (discussing “[m]achine [l]earning as a [b]arrier to [e]ntry”).

146. Hemphill, *supra* note 136, at 1977 (explaining that improving and deploying machine-learning-based software systems “has a high fixed cost and low marginal cost, a combination that tends to favor large firms that can spread the fixed cost over a large number of units”).

147. See *supra* text accompanying notes 58–59.

148. See *infra* notes 277–84 and accompanying text.

for example, yield positive externalities for other technological development.¹⁴⁹ Or investments in the development of a recommendation algorithm based on consumer preferences might find application in different consumer-facing markets (e.g., song, movie, and restaurant recommendations). Indeed, some markets may sustain several competitors through product differentiation: One song recommendation engine might prioritize similar musical genres while another emphasizes lyrical themes, and, given this differentiation, both might co-exist in a single market.¹⁵⁰ But, in general, the case for natural monopoly in machine-learning-based application markets is stronger than for traditional software development, particularly for deep-learning-based applications that are based on large text- and image-based datasets. Moreover, even where the natural monopoly condition is not strictly met, the market for particular applications may be highly concentrated, with such concentration giving rise to similar (if less pronounced) market effects.

2. Training and Optimization

Beyond the sorts of tangible and intangible inputs to machine-learning-based application development described above—human capital, hardware resources, and data—other machine-learning-specific features reinforce the preliminary conclusion in favor of natural monopoly. Specifically, the methods of optimizing machine-learning-based applications suggest that the time and computational resources dedicated to training amount to one more specific and significant fixed cost of development. And once an application is trained, it will often make sense to direct users (and their further inputs) to the existing application—much like, as in other natural monopoly contexts, it always makes sense to connect users to an existing network rather than to stand up a new one.

Computer scientists have often traditionally measured algorithmic efficiency using “Big O Notation,” a standard method for comparing different approaches to a computational problem (that discounts differences attributable to, say, hardware performance).¹⁵¹ Consider, for example, the familiar problem of sorting items in a list. There are more than a few ways to do this.¹⁵² You might, for example, look for the smallest item, set it aside, and then the next smallest, and so on. Or you might choose one random item in the middle

149. Cf. Colin Barker, *How the GPU Became the Heart of AI and Machine Learning*, ZDNET (Aug. 13, 2018), <https://www.zdnet.com/article/how-the-gpu-became-the-heart-of-ai-and-machine-learning> [<https://perma.cc/8XZ3-TRKC>] (describing how innovations in graphics processing have had spillover effects on machine learning).

150. Cf. VISCUSI ET AL., *supra* note 15, at 87; Hovenkamp, *supra* note 18, at 1996 (“Even if costs decline continuously as output increases or network effects are large, a digital platform is still not necessarily a natural monopoly. Another pervasive reason for [interplatform] competition is product differentiation.”).

151. See, e.g., RUSSELL & NORVIG, *supra* note 6, at 1053–55 (describing and explaining “O() notation”).

152. See KEARNS & ROTH, *supra* note 51, at 5 (listing sorting algorithms).

and push all the smaller elements to one side and all the larger elements to the other side, and repeat for each side, and so on, until everything is in place. From a computational perspective, these two approaches are significantly different.¹⁵³ The first approach—selection sort—is generally considered comparatively inefficient because as the list to sort gets longer, it takes significantly more time to finalize the sorting process: Each new element in the list adds one new round of comparisons and one additional comparison in each round. In formal terms, selection sort takes $O(n^2)$ time, where n represents the number of items in the list: As n gets larger, the algorithm's time complexity grows polynomially.¹⁵⁴ By contrast, the divide-and-conquer approach is more indifferent to increases in the size of the unsorted list: As the list gets longer, quicksort does not become unbearably slow. In formal terms, quicksort takes, on average, $O(n \log n)$ time—significantly better than selection sort's polynomial growth.

Such complexity measures suggest one way to detect further fixed and marginal costs (beyond hardware and data, among others) in computational spaces. An algorithm that operates in exponential time— $O(2^n)$ —becomes costly quickly. And such computational costs, moreover, are not mere hypothetical constructs: Rather, they represent real value in terms of hardware demands, electric power, climate effects, time, and opportunity costs, among others.¹⁵⁵ By contrast, an algorithm that operates in logarithmic time, e.g., $O(\log n)$, suffers its worst increases in complexity when n is small, but as n gets larger the incremental costs in computational complexity are comparatively slight.¹⁵⁶

153. *Id.* at 4–5 (noting “choices and trade-offs” in terms of memory and time).

154. One might be tempted to say, colloquially, that the complexity of this sorting algorithm grows “exponentially,” rather than “polynomially.” But it is more precise in this context to say that training runs in polynomial time ($O(n^2)$), since exponential runtime algorithms, e.g., $O(2^n)$, represent an even more costly class.

155. See generally Emma Strubell, Ananya Ganesh & Andrew McCallum, Energy and Policy Considerations for Deep Learning in NLP (June 5, 2019) (unpublished manuscript), <https://arxiv.org/pdf/1906.02243.pdf> [<https://perma.cc/2LH6-9YXV>] (describing costs in terms of power and shared computing resource prices); Kate Saenko, *Feed Me, Seymour!—Why AI Is so Power-Hungry*, ARSTECHNICA (Dec. 29, 2020, 6:38 AM), <https://arstechnica.com/science/2020/12/why-ai-is-so-power-hungry/?comments=1> [<https://perma.cc/XB8S-QQU2>] (citing the previous source and explaining that the power consumption demands of training and optimizing one machine-learning-based language model is equivalent to the cost of flying “315 passengers, or an entire 747 jet” on a “round trip between New York and San Francisco”).

156. Such algorithms operating in $O(\log n)$ time mimic, roughly, the effects of having large initial costs (high complexity costs for the first items in the list) and lower subsequent costs (comparatively lower complexity costs for the next items in the list).

These analogies are imperfect. Such computational algorithms (like the sorting example) typically operate at a single moment in time, with a complete set of inputs (say, the whole unsorted list). But a real market does not, of course, work in the same way. Nevertheless, as I describe in the following paragraphs, such distinctions are less salient (though still important) in the context of machine learning because of such applications' multistage nature.

These tools for measuring computational complexity can help assess the fixed and marginal technical costs associated with machine learning. In particular, machine-learning-based application development might be considered in at least two stages—one for designing, building, and training the application (which itself occurs in multiple stages), and a second for predicting future results based on such learnings.¹⁵⁷ Each of these stages has a different function and, importantly, a different computational cost structure. Application design and training more closely resembles software development: It is a critical input to the development of the machine-learning-based application. By contrast, prediction or classification more closely resembles distribution: Consumers of machine-learning-based applications send queries to the application and expect to receive some response—a proposed treatment regimen, the identity of a person in a photograph, or a set of search results.

Application training is computationally costly as compared to prediction or classification. In one example, computing the solution set—i.e., training—runs in polynomial time, while prediction runs in linear time.¹⁵⁸ I should clarify that my assessment here operates at a high level of generality, and, below, I use Bayesian classifiers as an exemplar. The genre of machine-learning-based applications encompasses a wide range of classifiers and techniques, each of which operates slightly differently and which may be better suited to different sorts of tasks and applications.¹⁵⁹ It is beyond my present

157. See, e.g., Lehr & Ohm, *supra* note 27, at 655 (“contend[ing] that legal scholars should think of machine learning as consisting of two distinct workflows” first, the design, development and training of an application, and second, an application “deployed and making decisions in the real world”); see also Antonio Rafael Sabino Parmezan, Vinicius M.A. Souza & Gustavo E.A.P.A. Batista, *Evaluation of Statistical and Machine Learning Models for Time Series Prediction: Identifying the State-of-the-Art and the Best Conditions for the Use of Each Model*, 484 INFO. SCIS. 302, 307–08 fig.4 (2019).

158. See MEHRVAR MOHRI, AFSHIN ROSTAMIZADEH & AMEET TALWALKAR, FOUNDATIONS OF MACHINE LEARNING 230 (2d ed., 2018); cf. Frank Hutter, Lin Xu, Holger H. Hoos & Kevin Leyton-Brown, *Algorithm Runtime Prediction: Methods & Evaluation*, 206 A.I. 79, 89 (2014).

I acknowledge that some researchers in the computer sciences disagree over the value of measuring algorithmic complexity in Big O terms for machine learning, preferring instead to rely on, say, real-time measurements or other constructs. See, e.g., Daniel Justus, John Brennan, Stephen Bonner & Andrew Stephen McGough, *Predicting the Computational Cost of Deep Learning Models 1* (Nov. 28, 2018) (unpublished manuscript), <https://arxiv.org/pdf/1811.11880.pdf> [<https://perma.cc/4SXA-7JJE>]. My point, however, does not turn on any one form of measurement. No matter whether one chooses Big O, real-time measurements on like hardware, or some other method, the basic point remains—training is significantly more costly than prediction or classification. See, e.g., Florian Scheidegger, Roxana Istrate, Giovanni Mariani, Luca Benini, Costas Bekas & Cristiano Malossi, *Efficient Image Dataset Classification Difficulty Estimation for Predicting Deep-Learning Accuracy*, 37 VISUAL COMPUT. 1593, 1597 fig.3 (2021); Abdiansah Abdiansah & Retantyo Wardoyo, *Time Complexity Analysis of Support Vector Machines (SVM) in LibSVM*, 128 INT’L J. COMPUT. APPLICATIONS 28, 33 (2015).

159. See generally Lehr & Ohm, *supra* note 27 (describing the machine learning process in order to elucidate the harms as well as benefits in machine learning and provide policy solutions);

scope to assess each of these sorts of classifiers individually, and so I paint with a somewhat broad brush, knowing that some of the conclusions suggested here may be overinclusive. Nevertheless, the process I describe offers a good general approximation for the comparative computational costs associated with training machine-learning-based applications and then later deriving predictions from such applications.¹⁶⁰ In all, training is significantly more costly than prediction.

How does training work? Consider Bayes' Theorem—foundational to much modern machine learning research and practice—which, in its simplest terms, estimates the probability of a given outcome in view of some known probabilities and available evidence.¹⁶¹ Training supplies the prior known probabilities. Some candies might be fruit-flavored (Skittles), others might be chocolate (M&Ms). If we have some known information about the frequencies of various colors in Skittles and M&Ms (by, say, opening bags of each sort of candy at random and counting the colors inside), then we might be able to predict the identity of an unknown candy, given that is brown (99 percent M&M, one percent Skittles) or orange (50 percent M&M, 50 percent Skittles). Of course, such predictions cannot be perfect, and so the Bayes error rate offers the lowest possible error for a Bayesian classifier under ideal conditions—the best it could do.¹⁶² Hence, when training a machine learning system to predict the identity of a candy (or, more seriously, of a criminal suspect), developers sometimes measure their application's error against this benchmark.¹⁶³

As these machine-learning-based applications are trained (e.g., through more training examples—data), they typically logistically (asymptotically, that is) approach the Bayes error rate. In short, the learning curve is steep: These

Parmezan et al., *supra* note 156 (describing different models); MOHRI ET AL., *supra* note 158 (describing many computational methods in detail).

160. See, e.g., JOEL HESTNESS, SHARAN NARANG, NEWSHA ARDALANI, GREGORY DIAMOS, HEEWOO JUN, HASSAN KIANINEJAD, MD. MOSTOFA ALI PATWARY, YANG YANG & YANQI ZHOU, BAIDU RSCH., DEEP LEARNING SCALING IS PREDICTABLE, EMPIRICALLY 13 (2017), <https://arxiv.org/pdf/1712.00409.pdf> [<https://perma.cc/WSSM-TDZD>] (demonstrating, across application types, that “model accuracy improvements from growing data set size and scaling computation are empirically predictable”).

161. See Thomas Bayes & Richard Price, *An Essay Towards Solving a Problem in the Doctrine of Chances*, 53 PHIL. TRANSACTIONS (1683-1775) 370, 370-418 (1763); see also, e.g., David Haussler, Michael Kearns & Robert E. Schapire, *Bounds on the Sample Complexity of Bayesian Learning Using Information Theory and the VC Dimension*, 14 MACH. LEARNING 83, 106 (1994) (“Perhaps the most important general conclusion to be drawn from the work presented here is that the various theories of learning curves based on diverse ideas from information theory [and] statistical physics . . . are all in fact closely related, and can be naturally and beneficially placed in a common Bayesian framework.”); Daniel Berrar, *Bayes' Theorem and Naive Bayes Classifier*, 1 ENCYC. BIOINFORMATICS & COMPUTATIONAL BIOLOGY 403, 403 (2018) (“Bayes' theorem is of fundamental importance for inferential statistics and many advanced machine learning models.”).

162. E.g., TREVOR HASTIE, ROBERT TIBSHIRANI & JEROME FRIEDMAN, *THE ELEMENTS OF STATISTICAL LEARNING: DATA MINING, INFERENCE, AND PREDICTION* 21 (2d ed. 2009).

163. See GOODFELLOW ET AL., *supra* note 28, at 113.

applications generally begin poorly, improve predictive accuracy with the first several thousands (or more) training examples and intense computational effort, and thereafter show some continued, though ever-decreasing, improvement as the volume of training examples increases.¹⁶⁴

Hence, the application's developers must select some point in the process to conclude the training process. Such decisions are often made by trading the costs of additional training—costs in terms of computing resources, additional data acquisition, and time—against the expected benefit of such training—benefits in terms of improved accuracy. Where those costs exceed benefits, training typically terminates.¹⁶⁵ What remains is an application that has maximized the trade (from the developer's view) between performance accuracy and training costs and is ready to respond to queries about a candy's flavor, a patient's treatment, a criminal's identity, or a website's location, as the case may be.

Importantly, moreover, the computational costs of responding to such individual queries are, as measured against the computational training costs, relatively small.¹⁶⁶

164. See, e.g., HASTIE ET AL., *supra* note 162, at 243 (2009); GOODFELLOW ET AL., *supra* note 28, at 114; see also SAMUEL L. SMITH & QUOC V. LE, A BAYESIAN PERSPECTIVE ON GENERALIZATION AND STOCHASTIC GRADIENT DESCENT 5–8 (2018), <https://arxiv.org/pdf/1710.06451.pdf> [<https://perma.cc/ZR4D-BTD8>]; TARUN YELLAMRAJU, JONAS HEPP & MIREILLE BOUTIN, BENCHMARKS FOR IMAGE CLASSIFICATION AND OTHER HIGH-DIMENSIONAL PATTERN RECOGNITION PROBLEMS 2 (2018), <https://arxiv.org/pdf/1806.05272.pdf> [<https://perma.cc/98PX-MH67>]; cf. Ivor W. Tsang, James T. Kwok & Pak-Ming Cheung, *Core Vector Machines: Fast SVM Training on Very Large Data Sets*, 6 J. MACH. LEARNING RSCH. 363, 380 (2005); Gary M. Weiss & Ye Tian, *Maximizing Classifier Utility When There Are Data Acquisition and Modeling Costs*, DATA MINING & KNOWLEDGE DISCOVERY (2007); Christopher Meek, Bo Thiesson & David Heckerman, *The Learning-Curve Sampling Method Applied to Model-Based Clustering*, 2 J. MACH. LEARNING RSCH. 397, 407 (2002).

It does not matter, for my purposes, precisely how computationally complex this training process is. It suffices, to note that training can be very costly in computational terms, and as I explained, that these training costs typically far outpace prediction costs. It is true that the nature of these costs can vary substantially, from exponential versus polynomial (as measured by the volume of the training data), depending on the sort of learning algorithm used, and so the conditions I describe here may not hold where the ratio between training costs and prediction costs is comparatively low. For more on the computational complexity of machine learning algorithms, see generally MICHAEL JUSTIN KEARNS, *THE COMPUTATIONAL COMPLEXITY OF MACHINE LEARNING* (1990) (exploring the computational complexity of machine learning). For a short summary of the complexity of various learning algorithms, see MOHRI ET AL., *supra* note 158, at 280, as well as RUser4512, *Computational Complexity of Machine Learning Algorithms*, KERNEL TRIP (Apr. 16, 2018), <https://www.thekerneltrip.com/machine/learning/computational-complexity-learning-algorithms> [<https://perma.cc/292X-HN4H>]; see also Hutter et al., *supra* note 158, at 89.

165. See, e.g., Weiss & Tian, *supra* note 164, at 4; Meek, et al., *supra* note 164, at 397. Indeed, developers must make these decisions about tradeoffs along a number of dimensions, including accuracy, bias, and robustness. See, e.g., IRENE Y. CHEN, FREDRIK D. JOHANSSON & DAVID SONTAG, *WHY IS MY CLASSIFIER DISCRIMINATORY?* 2 (2018), <https://arxiv.org/pdf/1805.12002.pdf> [<https://perma.cc/6J7M-JEVT>].

166. More precisely, for Bayesian classifiers, the prediction function is constant to the number of parameters to which the application is sensitive. In our candy example, for instance,

Some machine-learning-based applications may thus resemble a natural monopoly. This is especially true for large deep learning applications (such as GPT-3, a natural language processing application, or for facial recognition applications) though perhaps less true for simple regression models based on only a few thousand observations (such as COMPAS, an application used to conduct a risk-assessment of a criminal defendant, or even the college graduation example described above).¹⁶⁷ In such examples, the fixed computational costs of training the application will be comparatively high, growing with each training set and terminating when the costs of additional training outweigh the predicted accuracy benefits of such training. In contrast, the marginal costs of responding to queries after training will be comparatively low. Hence, assuming such a significant difference between the fixed computational costs of training and the marginal costs of prediction—an assumption that, as I noted above, may vary by the specifics of any given machine learning implementation—the average marginal computational cost of each query to the system seems likely to decrease continually. In short, machine-learning-based applications seem computationally subadditive.¹⁶⁸ Stated similarly, a second application, trained on a similar (though, perhaps, not identical) set of random bags of candy and taught to predict brand by color, would duplicate the efforts of the first application, adding no new capacity to the candy-detection market.¹⁶⁹ Indeed, research regarding model distillation suggests that machine-learning-based models tend to funnel

a prediction algorithm might try to discern a candy's identity based on color alone, or color and height, or color, height, and weight. Each added parameter increases the cost of the prediction function (as well as the training function). See Parmezan et al., *supra* note 157, at 332 (“[C]omputational complexity is empirically related to the number of parameters of a model.”). But, importantly for my present purposes, such changes in cost do not seem to regularly vary relative to the high fixed costs of training. No matter whether it takes \$1 to connect a new user to a telephone network or \$10, it is likely that the average marginal cost of connecting users to the network will decrease over the whole market, because of the wide disparity between fixed and marginal costs. Indeed (to draw out the example further), a \$10 connection cost (as compared to \$1) is indicative of even-higher fixed costs. This is because a complex prediction—one that turns on a number of parameters—also requires complex training along that same number of parameters. See MOHRI ET AL., *supra* note 158, at 280; *see also id.* at 230.

167. Compare *supra* note 139 (highlighting the costs of training GPT-3), with WILLIAM DIETERICH, EUGENIE JACKSON & CHRISTINA MENDOZA, RISK ASSESSMENT FACTSHEET: CORRECTIONAL OFFENDER MANAGEMENT PROFILING FOR ALTERNATIVE SANCTIONS (COMPAS) PRETRIAL RELEASE RISK SCALE II (PRRS-II), <https://www-cdn.law.stanford.edu/wp-content/uploads/2019/06/COMPAS-PRRS-II-Factsheet-Final-6.20.pdf> [<https://perma.cc/7R92-MWYN>] (describing COMPAS as a regression model based on fewer than 3,000 observations).

168. Baumol, *Quasi-Permanence*, *supra* note 22, at 12–13.

169. But, in a market that can sustain some forms of product differentiation, a second application trained to optimize along a different dimension—candy size, say—could add new (different) capacity to the market. *See infra* note 185–86 and accompanying text.

toward a single set of outcomes.¹⁷⁰ In short, the second provider is akin to a redundant set of railroad tracks or telephone wires.

3. Machine Learning's Virtuous Circle

As I noted above, developers of machine-learning-based applications typically terminate training when the costs of additional training (in terms of, say, acquiring additional training data) outweigh the benefits of such training (in terms of predictive accuracy). But what if it were nearly costless to obtain additional data and insights from that additional information could be cheaply incorporated into an application's prediction model? Various advanced machine learning techniques—deep and continual learning, for example—are aimed at accomplishing exactly that, allowing the model to digest new information on an ongoing basis. Such applications further resemble natural monopolies.

Machine-learning-based applications that continue to internalize new data, including information drawn from their practical deployments, may gain an insurmountable lead over putative competitors in their initial competition for the market. This is because the first application in the market gains access to more recent and more relevant training data—data from in-market consumers—before any competitor. Integrating those results into its prediction scheme thus gives rise to better results for the next query. And that next query, again, gives the provider even more recent and relevant data that may further improve its application—and so on. Industry experts and technologists have referred to this process of continual learning as the “virtuous cycle” of artificial intelligence.¹⁷¹ In these examples, the application's learning curve ascends steeply—it continually becomes more accurate—over a significant portion of the market (perhaps even the whole market), thus making it difficult, if not impossible, for putative competitors to ever catch up.¹⁷²

Scholars have suggested that this cycle helps explain the natural-monopoly-like tendencies exhibited by some machine-learning-based applications.¹⁷³ Frank Pasquale, for example, has suggested that search

170. See, e.g., Derrick Mwit, *Research Guide: Model Distillation Techniques for Deep Learning*, HEARTBEAT (Nov. 20, 2019), <https://heartbeat.comet.ml/research-guide-model-distillation-techniques-for-deep-learning-4a100801coeb> [<https://perma.cc/JE9S-3ESR>]. My thanks to Aileen Nielsen for highlighting this research.

171. See NG, *supra* note 33, at 4; see also Lee et al., *supra* note 33, at 51–52 (similarly describing the “virtuous cycle of AI”). From a putative competitor's perspective, this cycle may seem more vicious than virtuous.

172. See, e.g., Sumit Gupta, *Deep Learning Performance Breakthrough*, IBM (Jan. 16, 2018), <https://newsroom.ibm.com/IBM-systems?item=30076> [<https://perma.cc/3UCH-H82S>]; see also Andrew Ng, PowerPoint Presentation at Extract Conf., What Data Scientists Should Know About Deep Learning, at slide 30 (Nov. 24, 2015), available at <https://www.slideshare.net/ExtractConf/andrew-ng-chief-scientist-at-baidu> [<https://perma.cc/LAW4-YXXK>].

173. In addition to the sources described *infra* notes 174–75; see also Dipayan Ghosh & Ben Scott, *Digital Deceit II: A Policy Agenda to Fight Disinformation on the Internet*, SHORENSTEIN CTR. (Oct.

engines (which are sometimes implemented, at least partially, with machine learning technology) resemble natural monopolies, explaining that “[t]he ‘rich get richer’” and that the most popular search services gain access to “ever more data to refine and improve.”¹⁷⁴ Steven Weber and Gabriel Nicholas have similarly ascribed “a tendency toward natural monopolies in data platform businesses” to this very “positive feedback loop.”¹⁷⁵

Strictly speaking, this feedback loop does not give rise to the conditions that typically define a natural monopoly. As explained above, natural monopolies arise where a service (like telephone service) is most efficiently supplied by a single provider. But the cycle described here does not obviously bear on the supply-side costs of machine-learning-based applications.¹⁷⁶ Rather, this positive cycle seems to resemble a sort of demand-side network effect: Each use of the deep learning application increases the application’s value for future uses by improving the application’s predictive functions.¹⁷⁷ In short, demand begets demand.

Such network effects have, in this context, natural monopoly-like consequences: As with traditional natural monopolies, one provider may be better than two. If two competing providers were to split the market—and thus split access to the new data made available for continued improvement and retraining—then each application would benefit from only half the newly available data. By contrast, a single application in monopoly control of the market (and the resultant data) would have access to the entire corpus of information, thus performing better than either of the two competitors forced to split the new information.¹⁷⁸ As with traditional natural monopolies, one

2, 2018), <https://shorensteincenter.org/digital-deceit-ii-policy-agenda-fight-disinformation-internet> [<https://perma.cc/7CK9-V53F>] (noting “curation and targeting algorithms feed on one another to grow ever smarter over time—particularly with the forward integration of advanced algorithmic technologies, including AI” and that “[t]he most successful of the[se] companies are natural monopolies in this space; the more data they collect, the more effective their services”).

174. Pasquale, *supra* note 34, at 130; *see also* Bracha & Pasquale, *supra* note 18, at 1180–81 (2008) (“It is unclear whether search engines fall under the strict definition of a natural monopoly, but they exhibit very similar characteristics.”); *see also* Clark, *supra* note 108.

175. Weber & Nicholas, *supra* note 18, at 25–26.

176. *See* Lemley & McGowan, *supra* note 110, at 495 (“Network effects are demand-side effects—they result from the value that consumers place on owning what other consumers already own. By contrast, economies of scale are supply-side effects—they are a function of the cost of making the goods . . . regardless of positive utility payoffs among consumers.”); *see also* Daniel McIntosh, *We Need to Talk About Data: How Digital Monopolies Arise and Why They Have Power and Influence*, 23 J. TECH. L. & POL’Y 185, 196 (2019) (drawing similar distinction). By contrast, the two preceding subsections do focus on supply-side costs.

177. Indeed, these network effects accrue across a single *user*. Consider a police department that uses facial recognition technology to identify criminal suspects. Even if the police department is the only user of the deep learning system, accuracy might improve between the first, second, and third (and so on) uses by the single department, provided that each query contributes some new predictive information.

178. *See* SHAPIRO & VARIAN, *supra* note 110, at 173–226; DUCCI, *supra* note 18, at 5, 57–59 (“[T]he value of data and the economies of scale and scope that arise from larger datasets are

provider is more efficient than two (even though the source of this conclusion rests in demand-side, rather than supply-side, effects).¹⁷⁹ And that provider thereby wields significant power—“network power”—over the further development of the market.¹⁸⁰

It is possible, of course, that a second comer might leapfrog over the incumbent: Perhaps the incumbent’s training data or model design is deeply flawed, so as to permit a newcomer to do better, even given the disadvantage of being a late-mover. But even these such cases present some of the problems of natural monopoly discussed above: wastefully redundant efforts, higher input prices, and, during the time of the incumbent’s dominance, problematic rates and service. And, assuming the natural monopoly condition holds, one provider—in this case, perhaps the second-comer—is likely to dominate in the end.

Hence, given this similar preference for one provider over two, network effects often give rise to conditions that are “analogous in important ways to that of a natural monopoly.”¹⁸¹ In the different but related context of internet software, Mark Lemley has suggested that network effects may, like traditional natural monopolies, cause a market to “converge on a single product . . . over time.”¹⁸² Moreover, that result may “be highly durable and self-reinforcing, as

critical for the purpose of improving search algorithm predictions . . . where more data increase the efficiency of search results in terms of quality-adjusted costs.”).

This is especially true in instances of deep or continual learning where significant accuracy improvements accrue across the entire set of training data (including data used for ongoing training). *See supra* note 172 and accompanying text. In more traditional machine learning examples where, say, new training data is incorporated at regular intervals through retraining and redeployment, this is also true—but the effect may be far less pronounced: In such scenarios, accuracy seems bounded by minimum error rate—that is, the application learns quickly at first, but then slowly levels off indefinitely. *See supra* note 164 and accompanying text. Hence, if the total available training data is large enough such that even a fraction of it would allow an application to ascend the majority of the learning curve, then the market (and its data) may be split among multiple applications with comparatively slight accuracy effects (namely, the gradual improvements that would accrue to one provider after having ascended the steep part of the curve).

179. *See* SHAPIRO & VARIAN, *supra* note 110, at 173–226; Lemley & McGowan, *supra* note 110, at 549 (explaining, from the example of the telephone networks, that though “network effects did not dictate a telephone network run by a single firm as a regulated natural monopoly” such network effects did “dictate[] that one network was the efficient outcome” (emphasis omitted)); *see also* THOMAS HANNA, MATHEW LAWRENCE & NILS PETERS, A COMMON PLATFORM 6 (2020), <https://thenextsystem.org/sites/default/files/2020-12/Ownership-Futures-A-Common-Platform.pdf> [<https://perma.cc/V7PS-85ZW>] (“Network effects . . . are one of the reasons why some analysts consider the dominant platforms to essentially be natural monopolies.”).

180. *See generally* DAVID SINGH GREWAL, NETWORK POWER (2009) (describing the network effects of standard setting (to draw from the same example as in Lemley, *supra* note 133) as a form of power).

181. Lemley, *supra* note 133, at 1056.

182. *Id.*

consumers . . . focus their attention on the dominant standard to the exclusion of all others.”¹⁸³

* * *

Applications based on machine-learning technologies have been deployed in a wide range of industries and markets, from law enforcement to healthcare to software development itself. Users rely on these applications, themselves relying on existing stores of data, to classify new inputs and predict future outcomes: Given a vast repository of patient records, for example, an application might diagnose an ailment and suggest a course for treatment.

Such applications, moreover, resemble natural monopolies in several significant ways. First, the basic inputs necessary to build a robust machine learning architecture—expertise, data centers equipped with specialized computing equipment, and vast volumes of data—are costly. Second, the computational costs of converting those volumes of data into a predictive model—the costs of training—are similarly high. But the costs of prediction are comparatively low. Hence, this cost structure—computational and otherwise—may suggest that at least some of the markets for machine-learning-based applications are subadditive: It is most efficient to serve such markets with only one application. Third, even more recent advances in machine learning, like continual learning, suggest that these applications benefit from important network effects. In all, the putative developers of a machine-learning-based service face, in any given market, some important entry barriers—infrastructural expenses, computational costs, and network effects.¹⁸⁴ Of course, network effects are not, strictly speaking, suggestive of a natural monopoly. But they nevertheless imply a similar conclusion: In some markets for machine-learning-based applications, one provider can be better than many.

I am careful not to overstate the case: It is not necessarily the case that every machine-learning-based application operates in a natural monopoly market. Where, for example, an application is easily built on shared infrastructure using freely available information and operates in a market that might sustain a wide degree of product differentiation, the natural monopoly case is comparatively weak. Such applications may include some regression-based predictive models that are based on a limited (and cheaply obtained) dataset, such as criminal risk assessment tools, and also, perhaps, consumer-facing recommendation engines deployed by popular content providers

183. *Id.*

184. See Ghosh & Scott, *supra* note 173 (“The size of its physical infrastructure, the sophistication of its data processing algorithms (including AI and machine learning), and the quantity of data served on its infrastructure and feeding its algorithms (constantly making them smarter) constitutes an unassailable market advantage that leads inexorably to natural monopoly.”).

through shared computing infrastructures.¹⁸⁵ Though, depending on the strength of some of these features—where, say, multiple (but only a few) companies have the computational resources to develop some medical deep learning application, or where only a few dimensions for product differentiation are available—some markets may nevertheless exhibit concentration through oligopoly or monopoly control.¹⁸⁶ And where these features compound because applications are based upon a deep learning model whose costly training algorithms turn on several (billion) parameters and demand specialized expertise and equipment, as well as expensive or difficult-to-obtain data—e.g., GPT-3, facial recognition applications, and deep-learning-based medical applications—the natural monopoly condition seems increasingly likely to be satisfied. In short, the natural monopoly case depends on a close assessment of the features and inputs—computing hardware, data (and related network effects), and computational power—described above.¹⁸⁷

IV. MACHINE LEARNING AND REGULATORY POLICY

It thus seems likely that at least some machine-learning-based applications operate as natural monopolists. Moreover, empirical examples of consolidation in some application markets may offer some evidence suggestive of this natural monopoly condition.¹⁸⁸ Consider, for example, the use of facial recognition technology by law enforcement. Amazon, IBM, and Microsoft all recently exited that market.¹⁸⁹ In so doing, these companies cited concerns about law enforcement practices and racial bias, including, especially, unregulated bias in the technology itself.¹⁹⁰ Such explanations are certainly plausible. But (more cynically) it is also possible that these

185. See, e.g., *With Deep Learning, Disney Sorts Through a Universe of Content*, AMAZON WEB SERVICES, <https://aws.amazon.com/machine-learning/customers/innovators/disney> [<https://perma.cc/72j6-43TZ>] (explaining that Disney uses Amazon Web Services to enable deep learning algorithms categorizing its content repository).

186. See VISCUSI ET AL., *supra* note 15, at 87.

187. See Hovenkamp, *supra* note 18, at 1971 (explaining that “natural-monopoly status actually applies to particular inputs or technologies”).

188. See, e.g., Posner, *Natural Monopoly*, *supra* note 8, at 548 (suggesting that, in natural monopoly markets, “firms will quickly shake down to one”).

189. See, e.g., Matt O’Brien, *Microsoft Joins Amazon, IBM in Banning Police from Using Its Facial Recognition Tech*, CHI. TRIB. (June 11, 2020, 2:13 PM), <https://www.chicagotribune.com/business/ct-biz-amazon-facial-recognition-ai-20200611-at2lh7ub5bavzeukdrcsxk4ada-story.html> [<https://perma.cc/E4US-8K49>] (“Microsoft has become the third big tech company this week to say it won’t sell its facial recognition software to police, following similar moves by Amazon and IBM.”); see also Jeffrey Destin, *Amazon Extends Moratorium on Police Use of Facial Recognition Software*, REUTERS (May 18, 2021, 1:12 PM), <https://www.reuters.com/technology/exclusive-amazon-extends-moratorium-police-use-facial-recognition-software-2021-05-18> [<https://perma.cc/SN52-M388>].

190. See, e.g., O’Brien, *supra* note 189 (explaining that concerns related to law enforcement bias precipitated these companies’ withdrawal from the market).

companies recognized a market tending toward consolidation and that they are too far behind to make up the difference in view of the fact that “none [of these three companies] is a major player.”¹⁹¹ Instead, companies like NEC, which has relationships with law enforcement agencies “from Texas to Virginia” as well as international bodies (like London’s Metropolitan Police), and Clearview AI, which has relationships with “more than 600 law enforcement agencies,” might seem to have a distinct advantage in the battle for training data, and for the market.¹⁹² I do not, of course, mean to say definitively that these companies withdrew because they recognized a coming loss in the battle for the market. IBM, Microsoft, and Amazon might have been entirely sincere in their desire to exit after evaluating the social consequences of the technology.¹⁹³ But their behavior is also consistent with the dynamic of a natural monopoly market.

These observations help illuminate the basis—and possible need—for the regulation of machine-learning-based applications as natural monopolies. Here, the companies most willing to voice concern regarding bias-related accuracy problems withdrew from the market—possibly paving the way for another provider to consolidate its power over the market and, even more concerningly, perpetuate bias without competition.¹⁹⁴ In short, competition—competition from providers interested in producing and selling a less-biased service—may not effectively discipline these problems of accuracy and prejudice in the law enforcement market for facial recognition services. This pattern seems likely to repeat: Competition fails in natural monopoly markets, and so

191. *Id.* (“But while all three companies are known for their work in developing artificial intelligence, including face recognition software, none is a major player in selling such technology to police . . . Several other companies that are less well known dominate the market for government facial recognition contracts in the U.S., including Tokyo-based NEC and the European companies Idemia and Gemalto.”); *see also, e.g.*, Jared Council, *Facial Recognition Companies Commit to Police Market After Amazon, Microsoft Exit*, WALL ST. J. (June 12, 2020, 5:28 PM), <https://www.wsj.com/articles/facial-recognition-companies-commit-to-police-market-after-amazon-microsoft-exit-11591997320> [<https://perma.cc/A6W7-LX6W>] (quoting Jameson Spivack, a policy associate at the Center on Privacy and Technology at Georgetown Law, as stating that “[m]ost of the major companies that have the most contracts with law enforcement for facial recognition are smaller, specialized companies that most people have not heard of,” like NEC or Clearview AI); *see also* MORDOR INTEL., *GLOBAL FACIAL RECOGNITION MARKET (2020–2025)* 4 (2019) (market analysis report similarly focusing on NEC, Idemia, and Gemalto (among other companies) and excluding, notably, Amazon, IBM, and Microsoft).

192. *See* Julia Horowitz, *Tech Companies Are Still Helping Police Scan Your Face*, CNN BUSINESS (July 3, 2020, 8:36 AM), <https://www.cnn.com/2020/07/03/tech/facial-recognition-police/index.html> [<https://perma.cc/KFQ2-NYFB>] (describing NEC’s domestic and international clients); Hill, *supra* note 135 (describing Clearview AI’s reach).

193. *But see* O’Brien, *supra* note 189 (“MIT researcher Joy Buolamwini found racial and gender disparities in facial recognition software. Those findings . . . irked Amazon, which last year publicly attacked her research methods.”).

194. *See* Council, *supra* note 191 (finding that other providers had no intention of exiting the market over such concerns, explaining that their services “help identify suspects in criminal cases, find missing children and assist in other investigations”).

regulation has often substituted for the price and quality effects that market competition usually supplies.

But what sort of regulation? As described above, monopolists, including natural monopolists, often have incentives to charge inefficiently high prices or to underinvest in product quality. And these concerns persist even in other highly concentrated markets, including those characterized as oligopolies, so providers have historically faced forms of rate regulation, service specification, or specially induced competition in order to avoid these predictable problems, while retaining the efficiencies and network effects of consolidation. I begin by briefly describing these problems of machine-learning-based applications—problems that have been well-documented elsewhere in the legal, social science, and computer science literatures—before turning to a solution set drawing on the tradition of regulation described above.

A. SOME PROBLEMS OF MACHINE-LEARNING-BASED APPLICATIONS

Some developers of machine-learning-based applications stand accused of a wide range of infractions, from exploiting an application's users to generating results that create or perpetuate bias, as well as hoarding information about these application users. Such problems are sometimes cast as problems of personal privacy, of antidiscrimination law, or even of copyright.¹⁹⁵ But these might also be problems of market power: Consolidation has both enabled undesirable privacy-invading and data-hoarding practices and has caused providers to take a lax approach to problems of technical bias. It is, of course, naïve to think that mere market competition could resolve problems of, say, racial bias in law enforcement activity (including automated suspect identification), so I do not mean to suggest that either competition or regulation is a magic salve for the problems of discrimination and prejudice that can pervade society. But it is nevertheless true that industry consolidation can entrench these harmful practices, yielding technologies that enshrine and perpetuate problems of prejudice. Some discipline—perhaps through regulation—can help address the effects of such concerns.

I start with matters relating to privacy and private surveillance. Many machine-learning-based applications demand a price from their users beyond those imposed in dollar terms—they also extract information.¹⁹⁶ To the extent providers condition the use of a machine-learning-based application on access to some consumer information—to the extent, for example, that Google

195. See, e.g., Calo, *supra* note 50, at 1153–55 (personal privacy); Sonia K. Katyal, *Private Accountability in the Age of Artificial Intelligence*, 66 UCLA L. REV. 54, 59–62 (2019) (civil rights); Pauline T. Kim, *Data-Driven Discrimination at Work*, 58 WM. & MARY L. REV. 857, 883–92 (2017) (anti-discrimination); Lemley & Casey, *supra* note 49, at 744–50 (copyright); Levendowski, *supra* note 49, at 630 (copyright).

196. See, e.g., Hoofnagle & Whittington, *supra* note 38, at 626–28; Newman, *supra* note 54, at 203–06; Wu, *supra* note 38, at 296–99.

requires that consumers agree, in exchange for using Google to carry out internet searches, to allow Google to use information derived from those searches (e.g., search terms, the websites consumers visit, etc.) for its own purposes—those developers thereby extract some charge from consumers in exchange for using the application’s predictive capacity.¹⁹⁷

Where such providers are monopolists (including natural monopolists), they can charge prices—in terms of both data and dollars—that are undisciplined by competition, to two possible effects. One, these high prices can lead to accumulations of capital—in this case, of personal information—that undermine both distributive values and democratic participation.¹⁹⁸ As Julie Cohen has explained, people who feel themselves under constant surveillance—who are made to trade personal information to engage in modern society—lack the freedom to engage in the sorts of expressive and reflective conduct that is both “an indispensable structural feature of liberal democratic political systems” and a necessary element for further development.¹⁹⁹ In short, we feel the weight of these pervasive and invasive systems—and that burden restrains our ability to carry out our everyday lives. And two, even skeptics of such an account may concede that monopoly prices are likely to give rise to deadweight losses—i.e., an inefficient underutilization of a socially valuable service.²⁰⁰ Consumers may avoid using certain applications out of concerns that those services will extract too much information²⁰¹—even if those same consumers would be more willing to use a service with stronger privacy guarantees (e.g., a facial recognition application might collect either only that which is necessary for the effective performance of the application (like updated images of one’s own face), or it might collect updated images, as well as periodic location information, voice print information, unrelated usage statistics, and more).²⁰² And such selection effects

197. Cf. Ovide, *supra* note 38 (explaining that Google’s photo storage service “has never really been free [because] Google uses our photos to train its software systems”); see also Daisuke Wakabayashi, *A Former Google Executive Takes Aim at His Old Company with a Start-Up*, N.Y. TIMES (June 21, 2020), <https://www.nytimes.com/2020/06/19/technology/google-neeve-executive.html> [<https://perma.cc/B3Q8-RVWW>] (describing the process Google uses to target advertising based on user data).

198. See, e.g., Kapczynski, *supra* note 134, at 1489; see also SHOSHANA ZUBOFF, *THE AGE OF SURVEILLANCE CAPITALISM* 13 (2019) (describing the “new logic of [data] accumulation”).

199. Cohen, *supra* note 50, at 1905, 1912–18; see Julie E. Cohen, *Examined Lives: Informational Privacy and the Subject as Object*, 52 STAN. L. REV. 1373, 1426 (2000).

200. See Jones & Tonetti, *supra* note 49, at 2820 (describing firms’ incentives to “hoard data in ways that are socially inefficient”); Stucke, *supra* note 39, at 286 (explaining that a monopolist in a data-intensive industry “has the incentive to reduce its privacy protection below competitive levels and collect personal data above competitive levels” giving rise to deadweight losses).

201. Though, of course, this forces us to confront the question of how much is too much? And the answer may lie in the cognitive and expressive burdens of surveillance described *supra* text accompanying note 199.

202. Cf. Olivia Solon, *Insecure Wheels: Police Turn to Car Data to Destroy Suspects’ Alibis*, NBC NEWS (Dec. 28, 2020, 8:18 AM), <https://www.nbcnews.com/tech/tech-news/snitches-wheels->

may, in turn, further skew the results of continual learning systems.²⁰³ In sum, there is a robust concern among scholars, consumers advocacy organizations, and policy makers over the scope and extent of these natural monopolists' extractive data policies, both imposing significant privacy costs on individual citizens and undermining the use and continued refinement of these applications.²⁰⁴

Indeed, scholars and commentators have explained how machine-learning-based applications—applications built upon the stores of data generated by society (or some subset of society)—yield inaccurate and problematic results.²⁰⁵ In some instances, such inaccuracies are the byproduct of a poorly selected collection of training data—training data that fails to adequately represent the application's users.²⁰⁶ In others, these inaccuracies both mirror and propagate the systemic biases already endemic to society.²⁰⁷ And in still others, these errors results from flaws in the application's model design.

IBM, for example, reportedly trained a cancer-treatment application with “hypothetical patient data” yielding “unsafe and incorrect” treatment recommendations.²⁰⁸ But, apparently unconcerned about these results (as well as any competitive pressure from other providers), IBM sought to bring the product to market, notwithstanding these important and substantial inaccuracies.²⁰⁹ And this is not an isolated incident, as there are several

police-turn-car-data-destroy-suspects-alibis-n1251939 [https://perma.cc/Q4JM-KXE9] (interviewing one person who has declined to use connected vehicle services out of data extraction concerns); Valentino-DeVries, *supra* note 39 (explaining that privacy concerns have depressed interest in coronavirus-related contact-tracing applications); Cole et al., *supra* note 39 (“Privacy concerns, which lead to low adoption rates, are a barrier to the success of a contact tracing app.”).

203. Cf. Will Douglas Heaven, *The Way We Train AI is Fundamentally Flawed*, MIT TECH. REV. (Nov. 18, 2020), <https://www.technologyreview.com/2020/11/18/1012234/training-machine-learning-broken-real-world-health-nlp-computer-vision> [https://perma.cc/DA56-Z8JT] (describing this phenomenon as data shift).

204. See, e.g., Calo, *supra* note 50, at 1147–52; cf. Eleanor Fox & Harry First, *We Need Rules to Rein in Big Tech*, CPI ANTITRUST CHRON. (Oct. 27, 2020), <https://www.competitionpolicyinternational.com/we-need-rules-to-rein-in-big-tech> [https://perma.cc/PT28-RMCE] (“The platforms vacuum up huge amounts of data from users of the platforms, and use the data not only for efficiencies but also for exploitations and exclusions.”).

205. See, e.g., *Algorithms Behaving Badly: 2020 Edition*, MARKUP (Dec. 15, 2020, 8:00 AM), <https://themarkup.org/2020-in-review/2020/12/15/algorithms-bias-racism-surveillance> [https://perma.cc/SWF8-HGFG] (cataloging some of the most troublesome algorithmic deficiencies of 2020).

206. See Heaven, *supra* note 203 (describing underspecification).

207. See generally Friedman & Nissenbaum, *supra* note 50 (describing preexisting, technical, and emergent bias in computer systems).

208. Solow-Niederman, *supra* note 53, at 641 & n.34 (quoting Casey Ross & Ike Swetlitz, *IBM's Watson Supercomputer Recommended 'Unsafe and Incorrect' Cancer Treatments*, *Internal Documents Show*, STAT (July 25, 2018), <https://www.statnews.com/wp-content/uploads/2018/09/IBMs-Watson-recommended-unsafe-and-incorrect-cancer-treatments-STAT.pdf> [https://perma.cc/HZ4K-G54C]).

209. Solow-Niederman, *supra* note 53, at 641.

examples of troublesome, if not downright dangerous, artificially intelligent medical applications.²¹⁰ Some examples even suggest differences by skin color in an application's efficacy.²¹¹

Likewise, facial recognition applications and other “police surveillance technologies” have been shown both to exhibit increasing industry consolidation—due, at least in part, to the technological and market features elaborated above,²¹² and to be notoriously inaccurate, with the effects of these errors falling disproportionately on women of color. For example, in their widely cited study of several classification systems, Joy Buolamwini and Timnit Gebru found that darker-skinned females are misclassified in over one-third of all test cases—while lighter-skinned males are misclassified in less than one percent of tests.²¹³ And such results have, predictably, real-world effects. Law

210. See, e.g., Casey Ross, *Epic's AI Algorithms, Shielded from Scrutiny by a Corporate Firewall, Are Delivering Inaccurate Information on Seriously Ill Patients*, STAT (July 26, 2021), <https://www.statnews.com/2021/07/26/epic-hospital-algorithms-sepsis-investigation> [<https://perma.cc/46VE-UE82>] (noting “particular[] concern[] about Epic’s algorithm for predicting sepsis” finding that it both “routinely fails to identify the condition in advance, and triggers frequent false alarms”); see also Heaven, *supra* note 203 (describing similar examples regarding eye disease, skin cancer, and kidney failure).

211. See Todd Feathers, *Google's New Dermatology App Wasn't Designed for People With Darker Skin*, VICE: MOTHERBOARD (May 20, 2021, 8:40 AM), <https://www.vice.com/en/article/m7evmy/googles-new-dermatology-app-wasnt-designed-for-people-with-darker-skin> [<https://perma.cc/PN2T-V2U9>]. This app is a complex case. Reporting suggests a clear problem of sampling bias, one that may affect dermatology practice more generally. Google’s own tests suggest that it performs well for Black patients. But those results are based on ethnicity, which is distinct from skin type (i.e., color). Other researchers found Google’s training dataset to underrepresent patients whose skin color falls into the clinical definitions for “skin type V (brown)” and “skin type VI (dark brown or black).” *Id.*

212. See Elizabeth Joh & Thomas Joo, *The Harms of Police Surveillance Technologies*, DENVER L. REV. F. (forthcoming 2022) (manuscript at 1).

213. Buolamwini & Gebru, *supra* note 50, at 8–9 (evaluating three commercial gender classification systems and finding that darker-skinned females are misclassified in up to 34.7 percent of test cases, while lighter-skinned males are misclassified in only 0.8 percent of test cases); see also U.S. GOV'T ACCOUNTABILITY OFF., GAO-21-518, FACIAL RECOGNITION TECHNOLOGY: FEDERAL LAW ENFORCEMENT AGENCIES SHOULD BETTER ASSESS PRIVACY AND OTHER RISKS 24–26 (2021) [hereinafter GAO Report], <https://www.gao.gov/assets/gao-21-518.pdf> [<https://perma.cc/8MTW-YQUV>].

Facial recognition is only the beginning. Speech recognition applications, for example, also perform markedly worse for Black populations than for white. See Allison Koencke et al., *Racial Disparities in Automated Speech Recognition*, 117 PROC. NAT'L ACAD. SCI. U.S. 7684, 7685 (2020). Others have found evidence of gender bias in image search. See JIEYU ZHAO, TIANLU WANG, MARK YATSKAR, VICENTE ORDONEZ & KAI-WEI CHANG, MEN ALSO LIKE SHOPPING: REDUCING GENDER BIAS AMPLIFICATION USING CORPUS-LEVEL CONSTRAINTS 1 (2017), <https://arxiv.org/pdf/1707.09457.pdf> [<https://perma.cc/F674-KFB2>]. And in predictive text, too. Paresh Dave, *Fearful of Bias, Google Blocks Gender-Based Pronouns from New AI Tool*, REUTERS (Nov. 27, 2018, 12:06 AM), <https://www.reuters.com/article/us-alphabet-google-ai-gender/fearful-of-bias-google-blocks-gender-based-pronouns-from-new-ai-tool-idUSKCN1NW0EF> [<https://perma.cc/5CQ3-W4SH>]; see also Feathers, *supra* note 211 (noting concerns about a dermatology application).

There is a growing literature that details and examines these problems. See generally CATHY O'NEIL, WEAPONS OF MATH DESTRUCTION (2016) (describing the perils of “models [that]

enforcement officials in Detroit recently suggested that their facial recognition software misidentifies suspects in approximately 96 percent of cases, sometimes leading to false arrests or to tenuous criminal prosecutions.²¹⁴ Such inaccuracies in facial recognition systems (and other similar technologies) come with significant costs—but costs that are not borne by systems developers or their direct law enforcement purchasers: Citizens may, for example, be falsely accused of crimes, arrested and held in jail, and kept away from their families and jobs, with substantial consequences.²¹⁵ But such harms are borne only by private citizens—and, given the embedded bias in these systems, borne disproportionately by people of color—and are not internalized by either applications developers or their buyers in law enforcement.²¹⁶ Developers thus

encoded human prejudice, misunderstanding, and bias into software systems”); SARA WACHTER-BOETTCHER, *TECHNICALLY WRONG: SEXIST APPS, BIASED ALGORITHMS, AND OTHER THREATS OF TOXIC TECH* (2017) (describing how the culture of technology firms often lead to “shallow perceptions” of users that “at best, leave out huge percentages of users—and, at worst, take advantage of our personal data and encode bias into [software] systems”); VIRGINIA EUBANKS, *AUTOMATING INEQUALITY: HOW HIGH-TECH TOOLS PROFILE, POLICE, AND PUNISH THE POOR* (2018) (“systematically investigat[ing] the impacts of high-tech sorting and monitoring systems on poor and working-class people in America”); NOBLE, *supra* note 52 (describing how “algorithms . . . reinforce oppressive social relationships and enact new modes of racial profiling”).

214. Timothy B. Lee, *Detroit Police Chief Cops to 96-Percent Facial Recognition Error Rate*, ARSTECHNICA (June 30, 2020, 11:12 AM), <https://arstechnica.com/tech-policy/2020/06/detroit-police-chief-admits-facial-recognition-is-wrong-96-of-the-time> [<https://perma.cc/N9VU-2G4U>].

215. See O’NEIL, *supra* note 213, at 24–27; Kashmir Hill, *Another Arrest, and Jail Time, Due to a Bad Facial Recognition Match*, N.Y. TIMES (Jan. 6, 2021), <https://www.nytimes.com/2020/12/29/technology/facial-recognition-misidentify-jail.html> [<https://perma.cc/AS9V-F1PN>]; see also Garance Burke, Martha Mendoza, Juliet Linderman & Michael Tarm, *How AI-Powered Tech Landed Man in Jail with Scant Evidence*, AP NEWS (Aug. 19, 2021), <https://apnews.com/article/artificial-intelligence-algorithm-technology-police-crime-7e3345485aa668c97606d4b54f9b6220> [<https://perma.cc/RK9C-SUEU>] (describing the ShotSpotter application, used to identify the source of a gunshot sound, and an individual who was arrested on the strength of the application’s findings, but eventually released (after contracting COVID in jail) for insufficient evidence). These effects result not only from inaccuracies and bias in facial recognition (and other similar) applications, but also those embedded in applications used to assess whether criminal defendants should be eligible for bail or parole. See Julia Angwin, Jeff Larson, Surya Mattu & Lauren Kirchner, *Machine Bias*, PROPUBLICA (May 23, 2016), <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing> [<https://perma.cc/SXP3-6GAD>].

216. Hence, in this example, the market failure I have identified is not necessarily linked to the application providers’ status as a natural monopolist. Rather, the example here more closely resembles a principal-agent failure: Law enforcement agencies bargain on behalf of the public, but do not (and are not forced to) internalize all the costs of their decisions. See Mulligan & Bamberger, *supra* note 48, at 778–79. Hence, the agent’s decision may vary from the principal’s preferences. Cf. *Glob. Tel*Link v. FCC*, 866 F.3d 397, 404–05 (D.C. Cir. 2017) (describing a distinct but related principal-agent failure). To be sure, law enforcement agencies would certainly prefer more accurate systems to less accurate ones (though because they bear the costs of inaccuracies in different ways—police officers are rarely, if ever, jailed for misidentifying a suspect—they weigh accuracy differently than the general public). It is only the unavailability of more accurate systems that might be attributed to a failure of competition—while the procurement and use of these systems more closely result from the principal-agent problem. In all events, no matter the mechanism—whether through a principal-agent failure or natural monopoly—providers have reduced incentives to invest in training yielding accuracy concerns.

have little incentive to improve the accuracy of their systems: They may succeed in the market—they may earn contracts to serve law enforcement agencies in major metropolitan cities—even when they fail to accurately identify suspects in 96 percent of cases, with such errors falling disproportionately on those cities' citizens of color.

In short, many of the now well-known problems attending to machine-learning-based applications might be understood as problems of market power—and, in some cases, the predictable effects of the natural monopoly condition.²¹⁷ Where any given application will tend to consolidate its control over an application market—the market, say, for cancer diagnostic software—the application's provider is likely to tend towards higher user prices. And where such applications gain a respite from the competitive pressures that typically force quality improvements, application quality (in terms of, say, accuracy) may be likely to suffer—with the effects of such underinvestments likely to fall on historically marginalized groups, who are, among other things, underrepresented in available training data or represented discriminatorily (as a consequence of extant discrimination).

I do not mean to suggest that these applications *only* present problems. Rather, their invention suggests some existing problem in need of a solution.²¹⁸ Medical software and facial recognition technology, among others, all help to solve real, pressing problems. Moreover, to the extent these applications satisfy the natural monopoly condition, such problems seem most efficiently solved by a monopolist. But that solution offered by a natural monopolist gives rise to the further, different problems of monopoly power suggested above. We may thus consider—as prior scholars and policy makers have—how to retain the benefits of the machine-learning-based advance (and the benefits of consolidation), while mitigating the predictable negative consequences of its natural monopoly provision.²¹⁹

In future work, I plan to address a wider range of market failures, including both addressed in this footnote, that may require intervention in data-intensive industries.

217. Cf. Lyria Bennett Moses, *Regulating in the Face of Sociotechnical Change*, in THE OXFORD HANDBOOK OF LAW, REGULATION, AND TECHNOLOGY 573, 577–78 (Roger Brownsword, Eloise Scotford & Karen Yeung eds., 2017) (arguing that lawyers and regulators should focus on how they can best adjust the law and regulation in accordance with continual sociotechnical advancements rather than asking how to best regulate new technology).

218. Cf. PLATO, REPUBLIC 369 (Paul Shorey trans., Harv. Univ. Press 1969), <http://www.perseus.tufts.edu/hopper/text?doc=urn:cts:greekLit:tlg0059.tlg030.perseus-eng1:369> [<https://perm.a.c/FZ4J-K9S4>] (“Its real creator, as it appears, will be our needs.”).

219. Cf. *Proposal for a Regulation of the European Parliament and of the Council Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act) and Amending Certain Union Legislative Acts*, COM (2021) 206 final (Apr. 2, 2021) (noting that the “socio-economic benefits of AI can also bring about risks of negative consequences for individuals or the society”).

B. MACHINE LEARNING AND MONOPOLY REGULATION

Given the predictable consequences of monopoly control, we may turn to the traditional regulatory measures that policy makers have long deployed against natural monopolies. As described above, regulators have often required that natural monopolists and providers in concentrated markets offer reasonable service at reasonable rates in order to substitute for the price and quality effects typically supplied by competition.²²⁰ In some instances, policy makers have altered a market's landscape to make competition viable. Similar principles can apply to machine-learning-based applications, too.

As before, the suggestions elaborated below operate, by necessity, at some level of generality: Just as there are various types of machine learning implementations that vary in their technical details, so too are there numerous machine learning-based applications—medical and diagnostic, facial recognition, search, and so on—each presenting its own specific regulatory challenges. But, setting these differences to one side, we can nevertheless distill some general principles, founded on the natural monopoly tradition, for some modes of regulation that address several of the problems attending to machine-learning-based applications.

1. Informational Rate Setting for Privacy Harms

Given the likelihood that consumers will face monopoly prices, one straightforward possibility is to regulate the prices, in dollar terms, that service providers charge users. Consider, for example, a machine-learning-based diagnostic service that, as a natural monopolist, has captured its market. As with other monopolists, the diagnostic provider is likely to charge a profit-maximizing price—one that, as described above, may give rise to social losses by leaving some consumers (here, healthcare providers and their patients) unable to access the benefits of the new technology.²²¹ Hence, cost-based regulation of the prices charged by such services providers might help ensure comparatively wider adoption of the medical advance—and the concomitant realization of any associated health benefits.²²² Such price-oriented rate setting presents a standard, but important, regulatory response to natural monopolies

220. See *supra* Section II.B.1.

221. See W. Nicholson Price II, *Black-Box Medicine*, 28 HARV. J.L. & TECH. 419, 462 (2015) (“Payment is the dominant concern for adoption of a newly available medical technology.”).

222. Cf. *id.* at 462–63 (“[P]ublic policy could drive overall adoption [of machine-learning based medical tools] by encouraging reimbursement by public insurers, especially Medicare and Medicaid.”). Though Medicare is not, strictly speaking, a rate regulator, Medicare does, by setting reimbursement rates, have an important effect on market rates more generally. See, e.g., Eleanor D. Kinney, *Making Hard Choices Under the Medicare Prospective Payment System: One Administrative Model for Allocating Medical Resources Under a Government Health Insurance Program*, 19 IND. L. REV. 1151, 1164 (1986). Of course, this brief example is presented as simplified—pricing in medical markets is complicated, made even more so by the roles played by insurance networks—to highlight the threats of monopoly pricing and the role of responsive rate regulation.

in general, and need not turn on any feature specific to machine-learning-based natural monopolies.²²³

This approach, moreover, hints at another response that is more specific to the problem of data extraction described above.²²⁴ As other scholars have explained, and as I described above, some machine-learning-based applications extract value from their users in terms of data—data that is often repurposed for various ends, including advertising sales, the continued training of that application itself, and the cross-subsidization of new applications, among other purposes—at monopoly levels.²²⁵ In short, where applications providers are monopolists, there is cause to be concerned about the nature of that informational charge: It is not disciplined by competition and may be both extractive and inefficiently high.²²⁶

223. See Nicholas Bagley, *Medicine as a Public Calling*, 114 MICH. L. REV. 57, 62 (2015); Cf. Ramsi A. Woodcock, *The Contrasting Approaches to Power of the Modern State and the Antitrust Laws: Lessons for Regulation 66* (Oct. 3, 2020) (unpublished manuscript), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3704450 [<https://perma.cc/QQN5-6N2C>] (“In theory, price regulation is a complete solution to the problem of the exercise of power.”).

224. See *supra* notes 196–204 and accompanying text.

225. See, e.g., Hoofnagle & Whittington, *supra* note 38, at 609–10 (explaining that many firms—including, ostensibly, firms that rely on machine learning technologies—“collect valuable information about consumers” and “realize financial gains from the collection and use of [such] information”); Fox & First, *supra* note 204, at 2–3.

226. Several scholars have suggested that even where there are several participants in one market, such competition is unlikely to discipline such data “charges” because consumers are not particularly sensitive to such privacy-related concerns. See, e.g., Daniel J. Solove, *Introduction: Privacy Self-Management and the Consent Dilemma*, 126 HARV. L. REV. 1880, 1880 (2013) (“[E]mpirical and social science research demonstrates that there are severe cognitive problems that undermine privacy self-management.”); see also Cohen, *supra* note 50, at 1919 (“[R]egulators timidly opine that privacy harms result from ‘unexpected’ disclosures of personal information and that more robust guarantees of notice and choice therefore may be needed to ‘build[] consumer trust in the marketplace.’ This simplistic view of the relationship between privacy and innovation is wrong.” (second alteration in original) (footnote omitted)). I do not mean to dispute that point here, and, indeed, in future work I plan to consider the possibilities for regulation even under conditions of competition. For now, it suffices to note that, *even if* consumers are privacy-sensitive, see, e.g., Janice Y. Tsai, Serge Egelman, Lorrie Cranor & Alessandro Acquisti, *The Effect of Online Privacy Information on Purchasing Behavior: An Experimental Study*, 22 INFO. SYS. RSCH. 254, 255 (2011) (suggesting “that individuals are willing to pay a premium for privacy when privacy information is made prominent and intuitive”); see also Clark D. Asay, *Consumer Information Privacy and the Problem(s) of Third-Party Disclosures*, 11 NW. J. TECH. & INTELL. PROP. 321, 323 (2013) (noting “that notice and choice can and should play a significant role” in privacy law, even given the troubles that consumers face in fully understanding and internalizing privacy costs), regulation is warranted in machine-learning-based natural monopoly markets. Stated similarly, the model described here suggests that regulation is necessary because competition is exceptionally unlikely to have this important effect. Hence, like scholars who, drawing from behavioral literatures, conclude that regulation is necessary because consumers are insufficiently sensitive to these privacy risks, I similarly conclude that regulation is necessary because, no matter the extent of consumer sensitivity, competition is unlikely to discipline application developers. In either case, relying on usual market-based forces seems unlikely to succeed.

Regulators can set a more efficient and fair informational price.²²⁷ Of course, the analogy is not perfect: Information is neither truly rivalrous, nor a currency with which applications providers can pay salaries or purchase computing equipment. But, as explained, access to information can be limited or controlled, and this information supplies a direct input into these (and other) applications.²²⁸ And so regulators might limit information collection by setting an upper bound on the sorts of information that may be collected (including, perhaps, none at all), or they may allow developers to use only that which is reasonably necessary for the continued improvement of the regulated application. In short, regulators may adapt traditional rate-setting approaches—rate-of-return regulation and price-cap regulation—to the monopolists' information collection practices, setting out rules for data collection and use.

Under a rate-of-return-styled approach, regulators may set data collection rates and use-limitations by reference to the application's needs. Recall that standard rate-of-return regulation sets prices (in dollar terms) by reference to the real costs of production: Regulators attempt to discern the monopolists' total costs (including anticipated expenditures) and set prices to cover only those costs and some reasonable margin for profit.²²⁹ Likewise, in data contexts, regulators might attempt to discern an application monopolists' real information needs—what information is useful to, say, an application's continued learning (as opposed to information that is instead used to cross-subsidize the development of some other, unrelated product)—and limit data collection and use to those terms: They might, for example, allow diagnostics developers to collect relevant medical information, but preclude them from accessing other patient details, including names, addresses, or billing details.²³⁰ In short, regulators can require both data minimization (i.e., reductions in the amount of data collected and retained) and data protection practices (including,

227. Cf. 47 U.S.C. § 201(b) (2018) (requiring the authority to set “just and reasonable” rates).

228. See Katharina Pistor, *Rule by Data: The End of Markets?*, 83 L. & CONTEMP. PROBS. 101, 103–04 (2020).

229. See *supra* notes 75–76 and accompanying text (describing rate-of-return regulation).

230. See, e.g., Rob Copeland & Sarah E. Needleman, *Google's 'Project Nightingale' Triggers Federal Inquiry*, WALL ST. J. (Nov. 12, 2019, 11:13 PM), <https://www.wsj.com/articles/behind-googles-project-nightingale-a-health-data-gold-mine-of-50-million-patients-11573571867> [<https://perma.cc/575N-TEBR>] (explaining that Google's agreement to obtain medical records for health application development purposes included “personally identifiable information on millions of patients, such as names and dates of birth . . . and some billing claims,” among other details); see also Roger Allan Ford & W. Nicholson Price II, *Privacy and Accountability in Black-Box Medicine*, 23 MICH. TELECOMM. & TECH. L. REV. 1, 32–35 (2016) (proposing some such limits); Christina Farr, *Facebook Sent a Doctor on a Secret Mission to Ask Hospitals to Share Patient Data*, CNBC (Apr. 6, 2018, 11:46 AM), <https://www.cnbc.com/2018/04/05/facebook-building-8-explored-data-sharing-agreement-with-hospitals.html> [<https://perma.cc/UE3D-WJB2>].

e.g., retention limits and anonymization, as appropriate),²³¹ thereby helping to relieve the pressures of constant surveillance and to ensure wider adoption of these advances.

Of course, rate-of-return regulation can sometimes be prone to error. As noted earlier, regulated entities have an important advantage over regulators: In order to set rates correctly, regulators need access to information about the regulated entities' internal operations.²³² Regulated entities in turn have incentives to inflate their requirements, to exaggerate the difficulties associated with existing limits, and to entangle separate products and markets with the regulated one. In short, regulated entities might persuade regulators to set rates too high, at the expense of both welfare and distributional values (and it is important to consider both). Applications that are too costly to use will result in medical benefits that are inefficiently allocated (and, moreover, that internalize new data from only a selected population set), and that extract more information than necessary, further burdening our individual and collective privacy.

When faced with such problems of rate-of-return regulation, regulators have often switched to a mode of price-cap regulation. We can adapt a similar approach to the machine learning context, too: Rather than setting rates from applications developers' own statements of their needs, price-cap regulators tell providers what consumers should be made to offer (within reason and based, perhaps, on benchmarks), leaving it to those providers to work within those constraints.²³³ Indeed, some existing privacy statutes can be conceptualized as price-cap rules: The Health Insurance Portability and Accountability Act, for example, defines certain classes of regulated entities, including healthcare providers and insurers, sets caps on the sorts of information those entities may collect, and limits how such information may be used.²³⁴

Hence, this mode of rate setting is analogous to calls for federal privacy legislation—federal rules that would directly constrain data collection and use by vast range of companies, including, presumably, machine learning-based applications developers. But while much privacy regulation (both in Congress and through the courts) has traditionally focused on individual rights and harms, the tradition of monopoly regulation takes a wider view, considering

231. Cf. Commission Regulation 2016/679 of Apr. 27, 2016, General Data Protection Regulation, art. 5 (L 119) 1 (EU). Of course, some anonymization practices can be reverse-engineered. But some protection is likely better than none.

232. See *supra* note 84 and accompanying text (describing such problems).

233. Cf. John M. Newman, *Regulating Attention Markets* 45–48 (July 22, 2020) (unpublished manuscript), https://papers.ssrn.com/sol3/Papers.cfm?abstract_id=3423487 [<https://perma.cc/F3EE-NM3U>] (proposing advertisement caps—i.e., caps on charges imposed on consumers in attention markets).

234. See, e.g., 45 C.F.R. pt. 160–64 (2021); see also CAL. CIV. CODE § 1798.100 (West 2020) (explaining that personal information may be used for business purposes only to the extent such use is “reasonably necessary and proportionate”).

the broad range of market harms and benefits that accrue from certain pricing (and service) practices.²³⁵ A similar approach to informational rate setting can help to balance the privacy costs—both to individuals as well as to the public at large—against the benefits of various machine-learning-based applications.

Moreover, an agency model—one which vests authority over the specifics of a data-rate regime with executive departments—can be made more precise to context, both across sectors and over time. Consumers may, for example, be more willing to trade more sensitive personal information for superior medical diagnostic services than they are for better music content: Some might give up their entire medical history for a better diagnosis but would not trade their favorite color for a better song recommendation. And a regulatory model may be better able to capture such nuances than omnibus legislation can. Moreover, a regulatory model is likely to be better able to respond to changes over time: An unexpected pandemic, for example, may shift both the need and willingness for information related to a disease, its symptoms, and its spread.²³⁶ Regulators may be able to react to such requirements faster, and with greater expertise, than legislators. Though a complete description of the institutional design of data rate setting is beyond the scope of this project, it is possible, if not likely, that the need for context, expertise, and agility all suggest an approach that vests decisional authority with expert agencies (over one that relies upon direct congressional action).²³⁷

In addition to the benefits set out above—a more efficient allocation of the advances offered by machine-learning-based applications, and fewer concerns about developer fairness—price-cap regulation seems likely also to

235. Cf. Salomé Viljoen, *A Relational Theory of Data Governance*, 131 YALE L.J. 573, 592 (2021) (describing “privacy law’s individualism” and explaining how that approach is insufficient in light of the “population-level insights” that application developers attempt to discern).

236. *But see* Valentino-DeVries, *supra* note 39 (explaining that privacy concerns have depressed interest in coronavirus-related contact-tracing applications); Cole et al., *supra* note 39 (“Privacy concerns, which lead to low adoption rates, are a barrier to the success of a contact tracing app.”).

237. Indeed, in a similar but distinct context, I have argued in favor of an agency-based approach. See Tejas N. Narechania, *Patent Conflicts*, 103 GEO. L.J. 1483, 1529–32, 1534–36 (2015). Cf. Andrew Tutt, *An FDA for Algorithms*, 69 ADMIN. L. REV. 83, 90 (2017) (“a dedicated agency charged with the mission of supervising the development, deployment, and use of algorithms will soon be highly desirable, if not necessary”); RYAN CALO, CTR. FOR TECH. INNOVATION AT BROOKINGS, *THE CASE FOR A FEDERAL ROBOTICS COMMISSION* 3 (2014), https://www.brookings.edu/wp-content/uploads/2014/09/RoboticsCommissionR2_Calo.pdf [<https://perma.cc/3UQN-4BXX>] (“[T]entatively conclud[ing] that the United States would benefit from an agency dedicated to the responsible integration of robotics technologies into American society.”). I am happy to align myself with such calls for greater administrative oversight over these systems and applications. Whether, however, such additional oversight should take the form of some *new* agency, or, instead, additional capacity at existing agencies—particularly in view of the wide range of sectors, many of which are already regulated by agencies with relevant domain expertise, affected by machine-learning-based applications—is well beyond my present scope and left to future work. Cf. Narechania, *supra* note 237, at 1529–32, 1534–36.

fare better along dimensions of transparency, accountability, and legitimacy. Rate setting is a participatory and political process: It invites public comment and debate. And such public participation can help ensure that rates reflect public input, and that, if they do not, regulators may be held to account by elected leaders or the voting public. So too could data rate caps be set participatorily, by reference to shared public values rather than private developers' demands.²³⁸ While I set to one side, for now, the institutional design specifics that might attend to such data rate setting, it suffices for present purposes to note that, under a price-cap regime, decisions about the amount of private control over personal information—the extent of the incursion into private information spaces society is willing to tolerate—could be made by the public's representatives, with direct public input, in a participatory and democratically accountable process (rather than by a closed set of monopoly or oligopoly providers).

Hence, though defining the scope of a reasonable intrusion into our privacy (in exchange for some technological advance) remains an unresolved and perhaps interminable problem in the abstract, our democratic processes offer one way of addressing the commensurability challenges that come with informational rate setting. Any regulatory intervention comes with some costs and benefits: As developers access more information, they may be able to generate more accurate results, or even new applications altogether—but these informational concessions come at a great cost to our collective privacy.²³⁹ These competing concerns present a difficult commensurability problem, requiring us to trade privacy values against the uncertain potential of innovations in health, safety, and convenience—but the democratic processes and political checks that attend to public rate setting help address this quagmire.²⁴⁰ Of course, any rate setting approach carries some risk either that

238. See Kapczynski, *supra* note 134, at 1480 (calling for an account of law that allows the public “to channel outrage into a platform for democratic change”); Paul M. Schwartz, *Internet Privacy and the State*, 32 CONN. L. REV. 815, 816–17 (2000) (arguing in part “that the state has a special role in . . . developing privacy norms”); see also Elettra Bietti, *Locked-In Data Production: User Dignity and Capture in the Platform Economy 1* (Oct. 14, 2019) (unpublished manuscript), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3469819 [<https://perma.cc/YB6Y-DQU9>] (contending that “individuals should have a say over how data is collected, used and stored, and not only a right to be compensated for uses of their data determined by others”); Ford & Price II, *supra* note 230, at 35–36; cf. Rahman, *supra* note 44 (considering regulatory possibilities through the “more expansive” lens of public utility regulation).

239. See, e.g., W. Nicholson Price II, *Problematic Interactions Between AI and Health Privacy*, 2021 UTAH L. REV. 925, 935 (2021) (“AI weakens protections for health privacy, and health privacy weakens the AI used in health.”).

240. See Paul M. Schwartz, *Privacy and Democracy in Cyberspace*, 52 VAND. L. REV. 1609, 1648 (1999) (advocating for democratic deliberation and legislative enactment to protect privacy in cyberspace); see also *Major Rate Case Process Overview*, *supra* note 84 (“Rate cases are a primary instrument of government regulation of these industries. Interested persons may intervene and become parties in a utility company’s rate case. Typical intervenors include: industrial, commercial and other large-scale users of electricity; public interest groups; representatives of residential, low-income and elderly customers; local municipal officials; and, dedicated advocacy

regulators will be too aggressive in limiting providers' access to resources, or, conversely, that regulators will be captured by developers and cater to their priorities.

Hence, I do not mean to suggest that rate setting is an appropriate response to every application that relies on machine learning—especially since, as noted above, such applications are beginning to pervade the economy widely. Rather, the appropriate regulatory response depends—as it has in previous settings—on a difficult and prospective, but necessary, error cost analysis: Given the possibility both that regulators may err (by setting prices incorrectly) and that developers may err (by setting prices inefficiently and in conflict with shared public values), which outcome is more likely, and which error is more severe?²⁴¹ Here, as before, public officials—accountable to a voting public—may decide when the downside risk of regulator error is preferable to the costs imposed by a provider. In cases where the natural monopoly condition seems likely satisfied (or where significant concentration is likely), yielding predictable price effects, and where the consequences of error are greater (in, say, medical or law enforcement applications), the case for regulatory intervention is comparatively stronger.

In sum, to the extent machine-learning-based applications developers are indeed natural monopolists, competition fails to discipline their prices, both in terms of dollars and in data. Our longstanding response to such power over price has been rate setting: From the railways to the telephones (and beyond), regulators have long set rates for natural monopolists to avoid both undue concentrations of capital and the social losses of monopoly pricing, the risk of regulator error notwithstanding. An analogous approach might yield similar benefits when it comes to machine-learning-based applications and to prices set in terms of information: No matter whether we choose a rate-of-return or a price-cap approach to data rate setting, regulators can reduce extractive, inefficient, and privacy-invading information collection through more public, participatory, and democratic processes.

groups Rate cases proceed in an entirely public and open process.”); Anton Korinek, *Integrating Ethical Values and Economic Value*, in THE OXFORD HANDBOOK OF ETHICS OF AI 475, 491 (2020) (“[W]e need a large and concerted public effort . . . to ensure we develop AI in a direction that is both economically beneficial and ethically desirable.”); Moses, *supra* note 217, at 581–83; Pistor, *supra* note 228, at 118; Ben Leonard, *Rep. Buddy Carter Signals Support for Federal Privacy Legislation*, POLITICO (Jan. 26, 2021, 1:58 PM), <https://www.politico.com/news/2021/01/26/buddy-carter-supports-federal-privacy-legislation-462679> [<https://perma.cc/T2N5-ZRPX>] (paraphrasing Representative Carter as saying that “privacy legislation . . . could ease concerns and build public confidence in AI”); *cf.* FED. TRADE COMM’N, COMMENT LETTER ON PROPOSED RULE FOR DEVELOPING THE ADMINISTRATION’S APPROACH TO CONSUMER PRIVACY 11–12 (2018), https://www.ftc.gov/system/files/documents/advocacy_documents/ftc-staff-comment-ntia-developing-administrations-approach-consumer-privacy/p195400_ftc_comment_to_ntia_112018.pdf [<https://perma.cc/44Q9-QEGK>] (explaining that the Commission already makes decisions about trade-offs among, say, innovation, privacy, and security).

241. See William P. Rogerson & Howard Shelanski, *Antitrust Enforcement, Regulation, and Digital Platforms*, 168 U. PA. L. REV. 1911, 1924 (2020).

2. Service Standards for Accuracy and Bias Harms

As described above, natural monopoly regulation has traditionally focused not only on price effects, but also on service quality.²⁴² In ordinary markets, competition spurs providers to offer not only better prices, but also better products. But because competition tends to fail in natural monopoly markets, regulators have often issued service specifications that aim both to set a quality floor and to advance related social goals related to natural monopoly condition, (e.g., democratic-inclusion-related values that also help to maximize the network effects). Communications and electricity systems regulators, for example, have historically required that monopoly providers meet basic service level requirements and that they act as “carriers-of-last-resort” to all residents in their service area, including those that are otherwise too expensive to serve (because, say, they live too far from the providers’ existing facilities), in order to ensure that such residents are not excluded from the advances offered by the regulated service and that they are connected to and made a part of the national conversations that run atop these communications networks.²⁴³

Regulators may likewise set service requirements for machine-learning-based applications along a variety of dimensions to reduce the harms that result from algorithmic error and to guard against the risk that such harms will exacerbate existing inequities.²⁴⁴ Consider accuracy. As I explained above, the accuracy of an application’s prediction model is set, in significant part, by the availability of (comparatively cheap) training data. Hence, machine learning engineers tend to terminate training where the costs of additional training—costs in terms of data, time, and other inputs—outweigh the benefits in terms of accuracy improvements.²⁴⁵ In a competitive market, such improvements can be quite valuable: Each provider has an incentive to offer a product that is better—e.g., more accurate—than its competitors. In short, market success may depend on predictive accuracy, and so providers have incentives to invest in a well-trained application. Such incentives, however, are decidedly diminished where competition is diminished, as a less accurate

242. See *supra* Section II.B.2 (describing service specification).

243. See, e.g., NUECHTERLEIN & WEISER, *supra* note 3, at 297–98. *Contra* Restoring Internet Freedom, 33 FCC Rcd. 311, para. 103 (2018) (declaratory ruling, report and order) (explaining that the Commission selected a statutory classification for broadband, even if it excludes low-income consumers from obtaining subsidized broadband internet access, because the benefits of “innovation” outweigh the costs of more limited broadband internet access for such consumers).

244. For more on the need for, and the possibility of, data security regulation in response to market failure, see generally Jeffrey L. Vagle, *Cybersecurity and Moral Hazard*, 23 STAN. TECH. L. REV. 71 (2020) (discussing issues with data security regulation in response to market failure and addressing the moral hazard).

245. See Weiss & Tian, *supra* note 164, at 31; Meek et al., *supra* note 164, at 399–401; see also Solow-Niederman, *supra* note 53, at 682 (“Choices about the design of algorithmic systems, often made for business or technical reasons, will effectively regulate the ways that this technology interacts with human values.”).

model may, as noted above, do as well as in the market as a more accurate one, so providers may forego some application training to reduce costs.²⁴⁶ Hence, in such cases, a provider may decide to terminate training early—or, stated slightly differently, may terminate training where the private costs of training exceed the private benefits of training—ignoring the public benefits of increased accuracy.

Regulatory intervention can align private incentives with public priorities.²⁴⁷ In the context of accuracy, for example, regulators may set output standards requiring that, say, facial recognition applications deployed in law enforcement contexts meet certain standards—that any indication of a “match” is correct in at least, say, 96 percent of cases (rather than incorrect in 96 percent of cases)—perhaps by passing audits for bias and fairness.²⁴⁸ Indeed, proponents agree that such “appropriate guardrails” should apply to law enforcement use of facial recognition technology in order to prevent the sorts of false arrests and the associated grave consequences described above.²⁴⁹ Moreover, such standards might be set independent of the provider’s private training incentives—where an unregulated provider might prefer to train applications on hypothetical data or easily obtained public information scraped from the internet, and to terminate training once such cheap data runs out—a regulated provider must continue to source data, and train and optimize its application, until it satisfies regulatory standards.²⁵⁰ More generally, public standards can force developers into particular technical pathways that help satisfy public goals in addition to private aims.²⁵¹ And, as

246. See *supra* notes 164, 205–16 and accompanying text.

247. Cf. Yafit Lev-Aretz & Katherine J. Strandburg, *Regulation and Innovation: Approaching Market Failure from Both Sides*, 38 YALE J. ON REGUL. 1, 26 (2020) (“Regulation’s traditional goal is to bring market demand into better alignment with individual and social preferences and values . . .”).

248. See Lee, *supra* note 214. Accuracy in machine learning’s computational context has a specific meaning, referring to the ratio of true positives and true negatives as compared to all outcomes (including false positive and false negatives). Here, however, I do not mean to necessarily imply that regulators should set a precise standard for this sort of accuracy (which would treat false positives and false negatives equally). Rather, I use accuracy in a more colloquial sense, to underscore the need for some collective determination on the range and rate of acceptable errors (including the need to protect against errors that are biased against certain populations).

249. See Shira Ovide, *A Case for Facial Recognition*, N.Y. TIMES (Nov. 11, 2020), <https://www.nytimes.com/2020/11/11/technology/facial-recognition-software-police.html> [<https://perma.cc/BHA6-P4JV>] (“James Tate, a member of Detroit’s City Council . . . believe[s] th[at] facial recognition software—with appropriate guardrails, including multiple steps for approval[—]was an imperfect but potentially effective tool . . . for law enforcement in Detroit.”); see also *supra* notes 213–16 and accompanying text.

250. See, e.g., Price II, *supra* note 120, at 459 (“FDA cannot and should not abandon its command-and-control role in directly regulating premarket access for at least some forms of algorithmic medicine.”).

251. See Lehr & Ohm, *supra* note 27, at 669–70 (describing various choices that developers may make, including choices about data inputs, data cleaning, acceptable errors, and so on). See

before, those public standards can be set through public and participatory processes—ones that can lend greater public confidence and legitimacy to the application and its use (especially for public purposes, as in law enforcement).²⁵²

Likewise, regulatory intervention may help avoid the problems of preexisting and technical bias that have been documented in many machine-learning-based applications.²⁵³ As described above, many of these harms fall disproportionately on people of color, women, and other historically underrepresented and disadvantaged populations.²⁵⁴ In short, these systems are biased in ways that both reflect and perpetuate existing biases in society (and in the data that society produces), sometimes because reliable training data on minority populations can be comparatively expensive or difficult to procure, or because such reliable data cannot exist in a society that produces intrinsically biased outcomes. Regulators, however, can require that developers employ practices tempering the extent to which continually learning applications propagate or magnify biased outcomes, and that they obtain sufficient (if costly) training data encompassing historically underrepresented and disadvantaged populations (either by directly regulating data inputs or by testing an applications output), and that they remain current by training on new and updated datasets. Indeed, where intellectual property concerns might limit or cabin access to such data, policy makers may be able to overcome such limits through the exercise of regulatory authority or even eminent domain.²⁵⁵

Some scholars and commentators have suggested that regulators are ill-suited to such interventions, both because regulators lack the expertise necessary to evaluate machine-learning-based applications, and because it is difficult to apply such output-oriented standards to dynamic, continually learning applications.²⁵⁶ I agree that such service specification will require

generally CHEN ET AL., *supra* note 165 (contending that discriminatory applications should be remedied through additional data collection, rather than model constraints).

252. See, e.g., *Major Rate Case Process Overview*, *supra* note 84 (“Rate cases are a primary instrument of government regulation of these industries. Interested persons may intervene and become parties in a utility company’s rate case. Typical intervenors include: industrial, commercial and other large-scale users of electricity; public interest groups; representatives of residential, low-income and elderly customers; local municipal officials; and, dedicated advocacy groups Rate cases proceed in an entirely public and open process.”); see also BREYER, *supra* note 16, at 351–52 (discussing how public interest participation in agency decisions protects the interests of various populations); cf. FRANK PASQUALE, *THE BLACK BOX SOCIETY: THE SECRET ALGORITHMS THAT CONTROL MONEY AND INFORMATION* 216–18 (2015) (describing a need to reclaim institutions of democratic governance).

253. See *generally* Friedman & Nissenbaum, *supra* note 50 (describing taxonomy of bias encompassing preexisting bias, technical bias, and emergent bias). The forms of regulations described here are most likely to address forms of technical bias and to mitigate the effects of preexisting bias, but are more limited in their applicability to emergent bias.

254. See *supra* notes 205–16 and accompanying text.

255. Cf. Narechania, *supra* note 237, at 1529–32; *Ruckelshaus v. Monsanto Co.*, 467 U.S. 986, 987–88 (1984).

256. See, e.g., Solow-Niederman, *supra* note 53, at 638–39.

that some agencies and regulators build new capacities and expertise, and that developing standards for dynamic systems may prove difficult. But neither obstacle is, in my view, insurmountable—nor do I think we should abandon the challenge.²⁵⁷ Some public agencies, for example, already conduct periodic audits of machine-learning-based applications to detect bias.²⁵⁸ Others, like the FDA, already mandate an approval process that includes such testing.²⁵⁹ Therefore it is plausible to imagine a wider range of similar requirements that both set a quality floor and implement public antidiscrimination goals (among other policy possibilities). Whether it is prudent to do so depends, again, on a cost-benefit analysis that turns on the comparative likelihood and severity of both regulator error (including dynamic effects on innovation) and developer error (including the predictable effects of natural monopoly power, where it exists).²⁶⁰ It seems overwhelmingly likely that, in at least some instances, the disconnect between public and private incentives has yielded applications that impose significant costs on the public.²⁶¹ And while I do not discount the possibility that regulators will seek to entrench themselves, we might nevertheless prefer that these decisions about tradeoffs between private value and public good—or between competing public concerns—be made by

257. See Virginia Dignum, *Responsibility and Artificial Intelligence*, in THE OXFORD HANDBOOK OF ETHICS OF AI 215, 224–25 (Markus D. Dubber, Frank Pasquale & Sunit Das eds., 2020); Lehr & Ohm, *supra* note 27, at 716–17 (“Over time, we imagine the boundary between [the] ‘art’ and [the] ‘science’ [of machine learning] will continue to shift, with more processes becoming understandable science and fewer remaining inscrutable art.”); see also Ryan Calo, *Artificial Intelligence Policy: A Primer and Roadmap*, 51 U.C. DAVIS L. REV. 399, 420–22 (2017) (proposing solutions to the expertise gap).

258. See Cathryn Virginia, *Visionaries: Inioluwa Deborah Raji*, MIT TECH. REV. (Sept. 23, 2020), <https://www.technologyreview.com/innovator/inioluwa-deborah-raji> [<https://perma.cc/5KST-98LL>] (describing Raji’s work and explaining that it has spurred “the US National Institute of Standards and Technology [to] update[] its annual audit of face recognition algorithms to include a test for racial bias”); see also GAO Report, *supra* note 213, at 7–8.

259. *Artificial Intelligence and Machine Learning (AI/ML)-Enabled Medical Devices*, FDA (Sept. 22, 2021), <https://www.fda.gov/medical-devices/software-medical-device-samd/artificial-intelligence-and-machine-learning-aiml-enabled-medical-devices> [<https://perma.cc/RG5T-TM5H>]; see, Price II, *supra* note 120, 423–24 (describing the FDA’s approval processes); see also Rai et al., *supra* note 32, at 5 (describing and critiquing the FDA’s approval process as both “not always . . . asking for highly rigorous performance information with respect to the ML-CD software it reviews” and failing to publicly disclose “information about either the machine learning model used or the training data”).

260. See *supra* note 251 and accompanying text (suggesting a similar analysis in the context of ratesetting); see also Rogerson & Shelanski, *supra* note 241, at 1924; *id.* at 1920–21 (noting benefits of expertise in making such a calculation). Price II, *supra* note 120, at 451 (suggesting that “a stricter regulatory approach . . . could be overregulatory”). *But cf.* Lev-Aretz & Strandburg, *supra* note 247, at 17–22 (suggesting a more nuanced approach to the view that regulation stifles innovation); Kapczynski, *supra* note 134, at 1492–93 (same).

261. See, e.g., the examples described *supra* notes 208–16 and accompanying text.

regulators and policy makers who can be held to account to elected leadership and a voting public.²⁶²

In short, regulators have long set service rules for natural monopolists, both to ensure a minimum quality standard and to implement important social goals. Communications providers, for example, are subject to service quality requirements that pertain to call completion and emergency communications—they must, for example, make best efforts to complete every call sent over their network, with special rules applying to emergency 911 calls.²⁶³ Moreover, communications providers must serve as carriers-of-last-resort to all citizens in the regions they serve (including remote and difficult-to-reach clients), and are prohibited from redlining their service areas—drawing the boundaries of their networks to exclude low-income or minority populations. Regulators might draw up similar rules for machine-learning-based natural monopolies. For example, regulators might set data security standards or mandatory accuracy benchmarks, thereby issuing rules that supply training incentives that the market fails to produce.²⁶⁴ And they may, as with existing bias audits,²⁶⁵ check for compliance against test data that

262. See, e.g., *Microsoft Recommends Government Lead AI Frameworks*, WASH. INTERNET DAILY (Sept. 25, 2020), <https://washingtoninternetdaily.com/news/2020/09/25/Microsoft-Recommends-Government-Lead-AI-Frameworks-2009240055> [<https://perma.cc/4JGH-AQGXX>] (explaining that a government-led effort can help ensure that embedded values are “accountable and transparent”).

263. Narechania, *supra* note 41, at 205 (describing such requirements).

264. Some applications providers may counter that such regulations amount to a speech regulation that is impermissible under the First Amendment. Providers, that is, might argue that they have a First Amendment right to “speak” through inaccurate or incomplete prediction models. See Kapczynski, *supra* note 134, at 1510–11; cf. Kashmir Hill, *Facial Recognition Start-Up Mounts a First Amendment Defense*, N.Y. TIMES (Mar. 18, 2021), <https://www.nytimes.com/2020/08/11/technology/clearview-floyd-abrams.html> [<https://perma.cc/TG5T-HVLW>]. But here, too, the natural monopoly condition provides a ready solution: Courts have long held that governments have more power to regulate providers, even in ways adjacent to the First Amendment, where they are natural monopolists or have market power. See, e.g., *Comm’ns Co. v. City of Boulder*, 455 U.S. 40, 48–51 (1982). My future work will more fully examine this apparent market power exception to the First Amendment. For now, it suffices to note that this exception empowers providers to regulate the speech, such as it is, of machine-learning-based natural monopolists. See, e.g., Tim Wu, *Machine Speech*, 161 U. PA. L. REV. 1495, 1496–97 (2013); Stuart Minor Benjamin, *Algorithms and Speech*, 161 U. PA. L. REV. 1445, 1446–47 (2013).

265. See REVA SCHWARTZ, LEANN DOWN, ADAM JONAS & ELHAM TABASSI, NAT’L INST. OF STANDARDS & TECH., A PROPOSAL FOR IDENTIFYING AND MANAGING BIAS IN ARTIFICIAL INTELLIGENCE 5 (2021) (“NIST has experience in creating standards and databases, and has been evaluating the algorithms used in biometric technologies since the 1960s”), <https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.1270-draft.pdf> [<https://perma.cc/MEF3-8TV4>]; see also Bran Knowles & John T. Richards, *The Sanction of Authority: Promoting Public Trust in AI*, in PROCEEDINGS OF THE 2021 ACM CONFERENCE ON FAIRNESS, ACCOUNTABILITY, AND TRANSPARENCY 262, 263 (2021), <https://dl.acm.org/doi/abs/10.1145/3442188.3445890> [<https://perma.cc/77EH-UXPF>] (“show[ing] that public trust in AI will arise only through development of a robust regulatory ecosystem that provides some guarantee that the public is protected from harmful consequences of AI.”).

encompasses a diverse range of populations in order to limit the harms of algorithmic redlining and to advance the goals of equality, nondiscrimination, and democratic inclusion that have long been a part of natural monopoly regulation.²⁶⁶

3. Toward Competition

As explained above, rate regulation and service specification aim to substitute for market competition, given that the absence of competition in a natural monopoly market yields predictable price and quality effects, and, hence, the effects of regulator error may be less severe than the effects of monopoly control. But, as noted above, some modes of regulation aim to minimize the effects of both regulatory missteps and monopoly power by inducing efficient competition through, e.g., franchise bidding or lowering the costs of market entry (thereby undermining the natural monopoly condition).²⁶⁷ Similar solutions may address these concerns in machine learning contexts, especially where the risks of regulator error seem especially grave, or the natural monopoly case seems borderline.²⁶⁸

I begin with the analogy to franchising. In essence, a franchise is a license to serve some designated population: In many locales, only franchised cable operators may deploy a cable network and offer service to residents.²⁶⁹ And such franchises were historically earned through a process of competitive bidding that enabled local regulators to select the provider that promised the best prices—or, in some cases, the best value, by comparing providers' promised prices alongside dimensions of service quality.²⁷⁰ In short, franchising aims to reproduce the effects of competition by shifting from costly and inefficient intramarket competition to a competition among bidders for the market. Such intramarket competition is likely costly from a privacy perspective, as several competing providers have independent incentives to collect and recollect personal information (as opposed to making such data available to single, trusted, and regulated provider).²⁷¹

Moreover, Ken Bamberger and Deirdre Mulligan have detailed the many shortcomings of public officials' current competitive processes for selecting machine-learning-based applications. Acting as procurers of technologies, public agencies often seek bids from technology providers that focus narrowly

266. See Rahman, *supra* note 44, at 1637.

267. See *supra* Part II.

268. See Hemphill, *supra* note 136, at 1981 (explaining that “[f]ostering competition against the leading platforms is socially desirable [because] competition encourages lower prices and higher quality on both sides of the platform, including lower prices to advertisers and greater privacy protection for users,” among other reasons).

269. Narechania & Stallman, *supra* note 88, at 584–86.

270. *But cf.* VISCUSI ET AL., *supra* note 15, at 421 (suggesting that a franchise bidding process that accounts for dimensions beyond price is a bad idea).

271. See DANIEL J. SOLOVE, UNDERSTANDING PRIVACY 133 (2008).

on matters such as price.²⁷² Bamberger and Mulligan urge a turn away from a narrow procurement mindset and towards an approach encompassing “expertise, transparency, participation and political oversight, and reasoned decision making.”²⁷³ And while traditional franchising processes sit somewhere between pure government procurement and public policymaking, Bamberger and Mulligan are correct to urge decisionmakers to more fully consider the ramifications of their decision. When public agencies select a provider for, say, a law enforcement application, they may well be helping to entrench a monopolist—and to entrench that monopolist’s service’s embedded biases and values. And so policy makers (which, as Bamberger and Mulligan rightly note, include public procurement officials) must take greater care to consider the range of concerns—from privacy to accuracy and bias—that are implicated in the procurement of machine-learning-based applications.²⁷⁴

Policy makers have also, as explained above, sometimes sought to change the economics of market entry to induce efficient competition. As explained above, the 1996 Telecommunications Act offers one example. There, Congress sought to induce entry into local telecommunications markets (sometimes considered a natural monopoly) by requiring that network owners share their wires with competitors on reasonable and fair terms. Competitors could thus avoid the significant fixed, upfront investments required to enter the market and instead choose instead to build out their networks incrementally (relying on incumbent providers to fill gaps in the interim).²⁷⁵ And some municipalities built public infrastructure that they leased to various competing private providers.

Similar tools might apply to the fixed investments that machine-learning-based applications demand, including both computational hardware and data. As explained above, designing and training machine-learning-based applications often demands specialized (and expensive) computing infrastructure.²⁷⁶ Such infrastructure is already available on shared terms: Amazon Web Services, for example, sells access to its massive computing resources to applications developers (among many others). But as concerns about exclusion and affiliation in such markets grow—as such shared

272. Mulligan & Bamberger, *supra* note 48, at 779–80.

273. *Id.* at 780.

274. *See id.*; Cary Coglianese & Erik Lampmann, *Contracting for Algorithmic Accountability*, 6 ADMIN. L. REV. ACCORD 175, 197 (2021) (suggesting that governments pay “greater attention to how they design and structure their contracts for services to develop and operate AI tools” noting that such a procurement contracting approach can help to balance the costs and benefits of such technologies).

275. Of course, policy makers might be rightfully concerned about the incumbents’ willingness to share existing infrastructure, and the new competitors’ appetite for making investments in new systems. *See supra* note 102 and accompanying text (describing how a similar scheme failed to achieve hoped-for results in the context of the 1996 Telecommunications Act).

276. *See supra* Section III.B.1.

infrastructure may not be available to developers building applications that compete with, say, Amazon's other businesses²⁷⁷—policy makers may decide to ensure that such “machine-learning-as-a-service” offerings (to the extent they are available from only one or a few providers) are available on fair and reasonable—i.e., neutral—terms. Or, beyond these behavioral remedies, policy makers may choose structural solutions, such as a separations rule.²⁷⁸ Policy makers may also consider investing in a public computing infrastructure, akin to some local public investments made in communications architectures: Such a public option for machine learning may likewise discipline existing infrastructure providers' terms and rates through competition (and gives applications developers another option).²⁷⁹

Other scholars, including Amanda Levendowski, as well as Mark Lemley and Bryan Casey, have argued against intellectual property protection for training data, thereby reducing the costs of such data acquisition.²⁸⁰ And while freeing such reliable data from the copyright's constraints can reduce these fixed costs, copyright is only one barrier. Other data is protected by trade secret.²⁸¹ Private agreements—e.g., exclusive data sharing agreements between, say, developers and medical providers—can also restrain the development of competing applications, so policy makers might consider an even more expansive approach. One possibility is to limit such exclusive agreements and mandate data sharing, through, say, models of federated learning—i.e., “an approach to machine learning where a shared global model is trained across many participating clients that keep their training data locally.”²⁸² Another possibility is to create public datasets, in order to induce

277. Robin Harris, *It's Time to Spin AWS Out of Amazon: They Regulate Utilities, Don't They?*, ZDNET (May 26, 2020), <https://www.zdnet.com/article/why-its-time-to-spin-aws-out-of-amazon> [<https://perma.cc/G6KA-EPDB>] (“Amazon's Prime video runs on AWS, and so does Netflix, Disney+, and Hulu” and so “AWS has deep insight into its competitors' businesses.”); see also Dana Mattioli & Joe Flint, *How Amazon Strong-Arms Partners Using Its Power Across Multiple Businesses*, WALL ST. J. (Apr. 14, 2021, 10:34 AM), <https://www.wsj.com/articles/amazon-strong-arms-partnersacr-oss-multiple-businesses-11618410439> [<https://perma.cc/Q5CR-MQFL>].

278. For some contrasting approaches to addressing this problem, see, e.g., Lina M. Khan, *The Separation of Platforms and Commerce*, 119 COLUM. L. REV. 973, 1065–90 (2019) (proposing and considering a separations rule), with Woodcock, *supra* note 223, at 66 (proposing price regulation); Rogerson & Shelanski, *supra* note 241, at 1930–33 (proposing a nondiscrimination regime akin to network neutrality rules); *id.* at 1934–36 (considering, cautiously, a separations rule or line-of-business restrictions). See Tim Wu, *Network Neutrality, Broadband Discrimination*, 2 J. TELECOMMS. & HIGH TECH. L. 141 (2003) (proposing a non-discrimination regime).

279. PASQUALE, *supra* note 252, at 208–10. Cf. Yotam Harchol, Dirk Bergemann, Nick Feamster, Eric Friedman, Arvind Krishnamurthy, Aurojit Panda, Sylvia Ratnasamy & Michael Schapria, *A Public Option for the Core*, in SIGCOMM '20 377, 378 (2020).

280. Levendowski, *supra* note 49, at 622–30; Lemley & Casey, *supra* note 49, at 748; see also Kapczynski, *supra* note 134, at 1509.

281. See, e.g., Sonia Katyal, *The Paradox of Source Code Secrecy*, 104 CORNELL L. REV. 1183, 1186–87 & n.13 (2019); Brenda M. Simon & Ted Sichelman, *Data-Generating Patents*, 111 NW. U. L. REV. 377, 378 (2017) (describing two such examples).

282. See *TensorFlow Federated*, *supra* note 49.

greater competition.²⁸³ One concern, of course, with public datasets is that these will push developers to some suboptimal but uniform standard. So a federated learning approach among competitors may be more attractive, especially because, as explained above, consolidating such information into a single model can give rise to accuracy-improving network effects.²⁸⁴

And finally, as explained above, antitrust's essential facilities doctrine limits the extent to which a provider may leverage a natural monopoly into a monopoly in an adjacent market (such as leveraging control over power transmission lines into a monopoly over power generation, or leveraging control over telecommunications networks into a monopoly over telephone devices).²⁸⁵ Similarly, in the context of machine-learning-based natural monopolies, providers might leverage their monopoly control over one market into another in at least two ways. First, such providers may adjust their predictive model to prefer affiliated services: Google searches, say, for "Xi'an Famous Foods reviews" might yield results from Google's own restaurant review service, notwithstanding a consumer preference for a competing provider (such as Zagat),²⁸⁶ or a diagnostics system provider might recommend an affiliated treatment, even where a generic or other therapeutic equivalent is available. Second, a monopoly provider might use the data it collects from consumers to cross-subsidize the development of some separate product. Indeed, the Fifth Circuit has already recognized an essential-facilities claim to data used by natural monopolists in this way (though in the somewhat

283. See, e.g., Peter K. Yu, *Beyond Transparency and Accountability: Three Additional Features Algorithm Designers Should Build Into Intelligent Platforms*, 13 NE. U. L. REV. 263, 290–95 (2020) (describing data interoperability as critical to sustaining competition); Solow-Niederman, *supra* note 53, at 689; Hovenkamp, *supra* note 18, at 2032–38. Public regulators, moreover, may be able to overcome the intellectual property concerns that would otherwise attend to forced data sharing. See, e.g., Kapczynski, *supra* note 134, at 1509 (describing some such possible intellectual property limits); Narechania, *supra* note 237, at 1483 (explaining how regulators have overcome similar objections in other policy contexts). Indeed, even where these datasets are protected by trade secret, see *supra* note 281 and text, regulators may require that incumbents negotiate rates for sharing such protected information with putative competitors. See generally *Ruckelshaus v. Monsanto Co.*, 467 U.S. 986 (1984) (finding that, even though data used for product development and regulatory approval may be protected as trade secret and under the Constitution's Taking Clause, statutory schemes for arbitrating a price for the use of that data by competitors satisfy constitutional scrutiny).

284. See, e.g., Mary D. Fan, *The Public's Right to Benefit from Privately Held Consumer Big Data*, 96 N.Y.U. L. REV. 1438, 1477–91 (2021); see also Rogerson & Shelanski, *supra* note 241, at 1927–29, 1933–34 (similarly arguing in favor of interoperability and portability across data platforms).

285. Even skeptics of monopoly leveraging theories might be persuaded by the possibility for leveraging in these contexts, especially if these natural monopolists are made subject to (informational) rate regulation. See NUCHESTERLEIN & WEISER, *supra* note 3, at 14–17 (explaining Baxter's Law).

286. See Michael Luca, Tim Wu, Sebastian Couvidat & Daniel Frank, *Does Google Content Degrade Google Search? Experimental Evidence* 13 (Harvard Business School, Working Paper No. 16-035, 2015); Pasquale, *supra* note 52, at 22.

antiquated context of telephone directories).²⁸⁷ In total, natural monopoly providers may be able to use their monopoly power in one market to help develop and distribute products in adjacent markets—so, as regulators seek to induce competition in putatively regulated markets, they must also be wary to safeguard competition in adjacent markets, too.²⁸⁸

* * *

If some machine-learning-based applications might be natural monopolies, then some such applications might behave as natural monopolists. Such applications might, for example, make unreasonable or inefficient demands of their consumers, extracting more information than is reasonably necessary for the application's continued training. Such developers might also fail to make appropriate investments in an application's training, yielding inaccurate or prejudiced results. Therefore, we may wish to regulate these markets in order to limit these effects of market power, all while retaining the benefits of the machine-learning-based advance (and of consolidation).²⁸⁹

I emphasize that I do not mean to declare open season on all applications developers, granting regulators a license to target any machine-learning-based system. Rather, as noted above, the nature and extent of any regulatory response depends on a close analysis of the application at issue, the likelihood that the natural monopoly condition is satisfied, and a comparative analysis of the costs and benefits of regulation against the costs and benefits of concentration. Monopolists have incentives to charge too much or to skimp on service. But regulators can make critical errors, too—either by setting prices too high and quality standards too low (thus replicating the effects of monopoly power), or also by setting prices too low and quality standards too high (thereby deterring investment and innovation). Therefore any regulatory intervention must trade the possibility for regulator error against the likelihood of the exercise of monopoly power and must consider the benefits of the regulation alongside the benefits of consolidation.²⁹⁰ In some cases

287. See *Great W. Directories, Inc. v. Sw. Bell Tel. Co.*, 63 F.3d 1378, 1384, 1386 (5th Cir. 1995) (concluding that current customer information controlled by a telephone provider was an “essential facility,” and that the phone company could thus not attempt to monopolize the separate “directory market”).

288. See generally Nikolas Guggenberger, *Essential Platforms*, 24 STAN. TECH. L. REV. 237 (2021) (stating that an essential facilities framework can help constrain “Big Tech”); see also Khan, *supra* note 278, at 1066–74.

289. See Symposium, *Artificial Intelligence in Federal Agencies: Government by Algorithm: Symposium on Artificial Intelligence in Federal Agencies*, ADMIN. CONF. U.S. 7 (2020) (statement of Hillary Brill, Geo. L. Inst. for Tech. Law and Pol’y) (noting such tradeoffs); Ovide, *supra* note 249.

290. Cf. Bagley, *supra* note 223, at 62 (“Public utility regulation is every bit as much a part of that tradition as *laissez-faire*. And if the market-oriented approaches that are ascendant today prove unsatisfactory, public utility regulation is an option worth exploring [T]he debate

—medical applications, say, where computational demands are high, data are expensive, and the costs of monopoly too risky—a more stringent approach may make sense. In other contexts, some basic quality standards and periodic audits may suffice. And in still others, the availability of shared computing infrastructure and cheap, reliable data, alongside possibilities for market differentiation, can suggest the possibility for a competitive market. In short, there is space for a pluralistic approach even within this domain of regulation. Moreover, such an evaluation must itself be dynamic: As commentators from Schumpeter to Breyer have remarked, natural monopolies may be transient, undermined by creative destruction, disruption, and innovation.²⁹¹ But so long as the natural monopoly condition is satisfied, policy makers should seriously consider the possibilities for rules founded on this longstanding regulatory tradition, protecting the public from the harmful distributional and welfare effects of monopoly power—the privacy-invading and data hoarding practices, the misalignment among private incentives and public preferences, and the entrenchment of bias.

V. CONCLUSION: PUBLIC REGULATION AND PRIVATE POWER

Since the advent of the railway system, policy makers have sought to regulate natural monopolies in order to balance the harms of monopoly power—e.g., inefficient pricing, substandard service—against the gains of consolidation and network effects in such markets. Now, a new class of natural monopoly is emerging. Just as prior natural monopolies helped to transform the nation, machine-learning-based applications are remaking the economy in countless substantial ways.²⁹² Moreover, the fixed costs of developing these applications—fixed costs in terms of computing equipment, data, and computational power—together with the virtuous cycle of artificial intelligence suggest that at least some markets may be best served by a single application (as opposed to multiple competitors).

Hence, as a range of commentators across disciplines have begun to document several problems with these applications—problems of accuracy, bias, privacy, and security, among others—we may both understand those problems as ones of monopoly power and draw from the tradition of market regulation to address them. If developers make unreasonable data demands of their consumers, undermining the availability (and accuracy) of the application, then regulators may place limits on the use and collection of personal information. And if applications—freed from the constraints of

between market-oriented and regulatory approaches should unapologetically examine the virtues and vices of both.”).

291. See, e.g., JOSEPH A. SCHUMPETER, *CAPITALISM, SOCIALISM AND DEMOCRACY* 89 (3d ed. 1950); BREYER, *supra* note 16, at 286–87.

292. RUSSELL & NORVIG, *supra* note 6, at 28 (explaining that “many thousands of AI applications are embedded in the infrastructure of every industry” (quoting KURZWEIL, *supra* note 6, at 204)).

competition—make unsafe or prejudiced recommendations, then regulators may subject them to accuracy standards.

I do not mean to suggest that making these regulatory determinations will be easy or straightforward: Indeed, it may well be difficult to set data rate caps that reconcile interests in innovation against those in individual privacy, and to set and enforce standards against dynamic and evolving systems. Hence, the case for regulation may be comparatively strong in some contexts, while weak or nonexistent in others. It depends both on the features described above—the fixed equipment and data costs, the computational costs of training, and data network effects—that suggest the existence of a natural monopoly, and on the likelihood and severity of regulator error and developer power—the costs to innovation and investment of regulatory intervention, the costs to privacy of intrusive information demands, and the costs to society of entrenching and propagating bias.

Alongside these concerns, at least one more dimension merits consideration. We must decide not only whether regulators or applications developers are more likely to impose costs on society, but also the extent to which each may be held to account for their inevitable blunders.²⁹³ In short, we must decide between public regulation and private power. Here, a regulatory approach offers much in the way of transparency, public participation, and legitimacy—while also helping to ensure that the advances made possible with machine learning yield reasonable service at reasonable rates.

293. See, e.g., BREYER, *supra* note 16, at 346–54.