Summer 2023

The HIPAA Privacy Rule at Age 25: Privacy for Equitable AI

Barbara J. Evans

THE HIPAA PRIVACY RULE AT AGE 25: PRIVACY FOR EQUITABLE AI

BARBARA J. EVANS, PH.D., J.D., LL.M.*

* Professor of Law and Stephen C. O'Connell Chair, University of Florida Levin College of Law; Professor of Engineering, and Glenn and Deborah Renwick Faculty Fellow in AI and Ethics, University of Florida Herbert Wertheim College of Engineering, evans@law.ufl.edu. Author has no conflicts to disclose. This work received support under the National Institutes of Health Common Fund's Bridge2AI "Patient-Focused Collaborative Hospital Repository Uniting Standards (CHoRUS) for Equitable AI" project (OT2OD0327-01, Eric S. Rosenthal, PI), but views expressed are the author's own and do not necessarily reflect positions of her institution, research collaborators, or funders. The author would like to thank Francis X. Shen, Karl Surkan, Jennifer K. Wagner and students in their graduate seminars who reviewed earlier drafts and provided detailed comments, and physicians Azra Bihorac, Tyler J. Loftus, Eric S. Rosenthal, and Michael J. Young for their valuable insights.

INTRODUCTION

President Bill Clinton signed the Health Insurance Portability and Accountability Act (HIPAA) into law on August 21, 1996.[1] Its twenty-fifth birthday passed largely unnoticed in August 2021 in a nation wracked by contagion and a rough exit from Afghanistan.[2] HIPAA was mainly an insurance statute best known among health care providers for its medical privacy regulation, the HIPAA Privacy Rule,[3] which took effect in 2003-2004 after a long, contentious rulemaking.[4] The Privacy Rule is simultaneously criticized for allowing too much access to patients' health information and too little.[5] Seemingly no one is happy with it.

The major criticism, both among scholars and members of the public, is that the Privacy Rule allows sensitive health information to be shared and used, potentially in identifiable formats, without individual consent.[6] This deviation from popular norms of informed consent (notice and consent) strikes many observers as unethical. An alternative view, advanced here, is that the Privacy Rule is ethically sound

---

1. Health Insurance Portability and Accountability Act of 1996, Pub. L. No. 104-191, 110 Stat. 1936 (codified as amended in scattered sections of 18, 26, 29 and 42 U.S.C.).

2. *See, e.g.*, Madeline Holcombe & Jason Hanna, *With More than 100,000 People in the Hospital with COVID-19 in the US, This August Is Worse than Last, Expert Says*, CNN (Aug. 26, 2021, 12:52 PM), https://www.cnn.com/2021/08/26/health/us-coronavirus-thursday/index.html [https://perma.cc/J2AY-2FD9] (reporting a surge in COVID cases); *see also* Michael D. Shear et al., *Miscue After Miscue: U.S. Exit Plan Unravels*, N.Y. TIMES (Aug. 31, 2021), https://www.nytimes.com/2021/08/21/us/politics/biden-taliban-afghanistan-kabul.html [https://perma.cc/X435-7JZB] (reporting difficulties with U.S. exit from Afghanistan).

3. *See* 45 C.F.R. pts. 160, 164 (2022).

4. *See* Barbara J. Evans, *Institutional Competence to Balance Privacy and Competing Values: The Forgotten Third Prong of HIPAA Preemption Analysis*, 46 U.C. DAVIS L. REV. 1175, 1213-15 (2013) [hereinafter Evans, *Institutional Competence*]; Grace Ko, *Partial Preemption Under the Health Insurance Portability and Accountability Act*, 79 S. CALIF. L. REV. 497, 505 (2006).

5. *See* COMM. ON HEALTH RSCH. & THE PRIV. OF HEALTH INFO.: THE HIPAA PRIVACY RULE, INST. OF MED., BEYOND THE HIPAA PRIVACY RULE: ENHANCING PRIVACY, IMPROVING HEALTH THROUGH RESEARCH 66 (Sharyl J. Nass et al. eds., 2009) [hereinafter IOM, PRIVACY REPORT], http://www.nap.edu/catalog/12458.html [https://perma.cc/6LUT-D9H2] (describing public concerns about unconsented access to data); William Burman & Robert Daum, *Grinding to a Halt: The Effects of the Increasing Regulatory Burden on Research and Quality Improvement Efforts*, 49 CLINICAL INFECTIOUS DISEASES 328, 328 (2009) (arguing that "the application of the Health Insurance Portability and Accountability Act to research has overburdened institutional review boards (IRBs), confused prospective research participants, and slowed research and increased its cost"); Fred H. Cate, *Protecting Privacy in Health Research: The Limits of Individual Choice*, 98 CALIF. L. REV. 1765, 1797 (2010) ("Consent requirements [imposed by the HIPAA Privacy Rule] not only impede health research, but may actually undermine privacy interests.").

6. *See infra* Table 1 (listing over twenty legal pathways for unconsented access to personal health information under the HIPAA Privacy Rule); *see also* IOM, PRIVACY REPORT, *supra* note 5, at 66 (noting that the HIPAA Privacy Rule has not eliminated the concerns of the public, which is "deeply concerned about the privacy and security of personal health information," and reporting that "[i]n some surveys, the majority of respondents were not comfortable with their health information being provided for health research except with notice and express consent").

but reflects a different balancing of competing moral principles, placing beneficence, justice, and equity on a more equal footing with individual autonomy. This balancing, while controversial, makes the Privacy Rule potentially well-tailored for novel ethical challenges that lie ahead in the age of AI-enabled health care.

The Privacy Rule was somewhat ahead of its time, designed in anticipation of health information technology that was not yet operational at the time the regulation was promulgated.[7] The Privacy Rule's drafters foresaw an increasingly diverse American population served by twenty-first-century health systems that, increasingly, would derive general medical knowledge from informational as well as clinical research.[8] This shift to informational research—large-scale data-driven discovery using people's health data and biospecimens, as opposed to experimenting on their bodies—was already underway in 1996, as information technology ushered in "an era of large volumes of data on platforms conducive to analyses."[9] Artificial intelligence/machine learning (AI/ML) clinical decision support (CDS) software, the focus of this Article, was one outgrowth of that trend.[10] It is here, now, contributing to workflows in today's health care system. As it does so, it poses a new set of ethical challenges that this Article explores.

In light of these challenges, the balance struck in the Privacy Rule may offer certain advantages. In particular, it offers legal pathways for assembling the diverse, inclusive health data sets that will be needed to tackle disturbing racial, gender, socioeconomic, and other biases observed in the current generation of AI/ML CDS tools.[11] This Article argues that unconsented access to data is not ethically problematic in and of itself. What has made it problematic is the weak framework of alternative protections that the Privacy Rule prescribes when data are used without individual consent. This Article proposes

7. *See* Nicolas P. Terry, *Regulatory Disruption and Arbitrage in Health-Care Data Protection*, 17 YALE J. HEALTH POL'Y L. & ETHICS 143, 166-67 (2017) (noting that HIPAA took "a pre-IT [information technology]" approach to data use at a time when electronic health records were anticipated but barely visible).

8. *See* Kayte Spector-Bagdady, *Governing Secondary Research Use of Health Data and Specimens: The Inequitable Distribution of Regulatory Burden Between Federally Funded and Industry Research*, J.L. & BIOSCIENCES, Jan.-June 2021, at 4 (discussing the shift from human subjects clinical research that studies people's bodies to informational "research with all the stuff [such as data and biospecimens] derived from them").

9. Telba Irony, *Evolving Methods: Evaluating Medical Device Interventions in a Rapid State of Flux*, *in* INST. OF MED., ROUNDTABLE ON EVIDENCE-BASED MEDICINE: THE LEARNING HEALTHCARE SYSTEM: WORKSHOP SUMMARY 93, 95 (LeighAnne Olsen et al. eds., 2007), https://www.nap.edu/catalog/11903/the-learning-healthcare-system-workshop-summary [https://perma.cc/JV2S-CU2H].

10. *See infra* Part II (defining and discussing CDS software).

11. *See infra* Section II.A (summarizing results from empirical studies of racial, gender, and socioeconomic disparities in how current CDS tools perform for various patient subpopulations).

specific measures to strengthen those protections so that the pursuit of greater health care equity need not imply a loss of meaningful privacy standards.

Part I describes two competing visions of how to protect data privacy, examining the roots of ongoing discontent with the Privacy Rule and tracing policymakers' original rationale for fashioning a major federal privacy regulation that allows so much unconsented access to health data.

Part II briefly introduces what AI/ML CDS tools are and why they are poised to occupy a central position in twenty-first-century AI-enabled health care. The growing use of these tools creates an unfamiliar landscape in which past insights about the "right" way to protect data privacy may need revisiting. Part II levels three critiques at popular, post-1970s privacy policies that rely on individual consent rights and simple data de-identification strategies as their main tools of data privacy protection. First, such policies can have disparate impacts that threaten to exacerbate health inequities in an AI-enabled health care system. Second, notice-and-consent privacy policies rest on philosophical and scientific assumptions that deny the reality of human diversity, completely at odds with a twenty-first-century health care system tasked with serving ever more diverse patient populations. The third and possibly most damning critique is that widely favored consent norms and data de-identification methods often fail at their central mission: they do not provide very strong privacy protection.

Part III identifies five legal pathways available under the Privacy Rule that could enhance access to diverse, inclusive data sets to train a new generation of more-equitable AI/ML CDS tools. Part IV explores why, twenty-five years after HIPAA's inception, these data access pathways continue to be underutilized, contributing to the observed pattern of CDS tools that tend to work better for cis-gendered white males treated at leading academic medical centers than for all the rest of us.[12] The Privacy Rule enables data acquisition practices that could enhance health equity while affording stronger privacy protections than patients enjoy today, yet gatekeepers of data hesitate to embrace these practices amid lingering concern about gaps in the Privacy Rule's privacy framework. Part IV concludes that these concerns are valid and proposes specific measures to address them.

I. TWO COMPETING VISIONS OF DATA PRIVACY

This Part describes two starkly different visions of data privacy and how to protect it. The first is the "control-over-information" theory, which enjoys wide support both among scholars and members of the

12. *See infra* notes 128-30 and accompanying text.

public at this time.[13] It is a rights-based model that "conceives of privacy as a personal right to control the use of one's data" through strong norms requiring individual consent, and it is said to be the "leading paradigm on the Internet and in the real, or offline world."[14] It stresses the ethical principle of "respect for persons," widely understood to mean respect for their autonomy and rights of self-determination, which is often tinged with Lockean assertions that individuals have property rights in their persons and, by extension, in information about themselves.[15]

The second vision of privacy is an older, duty-based approach that treats autonomy as important without elevating it above "other competing values in the hierarchy of ethical goods, such as beneficence, justice, dignity, and equality."[16] This duty-based approach has been called "privacy's other path" because, instead of emphasizing the data subject's right of control over data, it stresses data handlers' duty to treat data confidentially.[17] This approach, which survives in modern medical privacy laws, protects privacy by regulating specific social relationships (e.g., the physician/patient relationship) where data are generated and used.[18] It accepts that "autonomy as a construct cannot account for the ethical responsibilities of the caregiver."[19] In other words, consent rights alone cannot protect people's privacy unless those who handle their data have duties to treat the data with care. By this view, asking people to consent to do business with irresponsible data

13. *See* Daniel J. Solove, *Conceptualizing Privacy*, 90 CALIF. L. REV. 1087, 1110-11 (2002) (discussing control-over-information theory).

14. Paul M. Schwartz, *Internet Privacy and the State*, 32 CONN. L. REV. 815, 820 (2000).

15. *See* Solove, *supra* note 13, at 1112 (citing JOHN LOCKE, SECOND TREATISE ON GOVERNMENT § 27, at 19 (1980) (1690)); *see also* NAT'L COMM'N FOR THE PROT. OF HUM. SUBJECTS OF BIOMEDICAL AND BEHAV. RSCH., THE BELMONT REPORT: ETHICAL PRINCIPLES AND GUIDELINES FOR THE PROTECTION OF HUMAN SUBJECTS OF RESEARCH pt. B (1979) [hereinafter BELMONT REPORT], https://www.hhs.gov/ohrp/regulations-and-policy/belmont-report [https://perma.cc/6526-X9Z9] (listing, in a foundational 1979 work in the field of bioethics, "respect for persons" as the first "basic ethical principle[]," followed by beneficence/nonmaleficence and justice).

16. O. CARTER SNEAD, WHAT IT MEANS TO BE HUMAN: THE CASE FOR THE BODY IN PUBLIC BIOETHICS 71 (2020) (criticizing the rights-based model of bioethics that is reflected in control-over-information theory); *see also* TOM L. BEAUCHAMP & JAMES F. CHILDRESS, PRINCIPLES OF BIOETHICS (5th ed. 2001) (discussing the weight given to autonomy in modern bioethics after 1970); Paul Root-Wolpe, *The Triumph of Autonomy in American Bioethics: A Sociological View*, *in* BIOETHICS AND SOCIETY: CONSTRUCTING THE ETHICAL ENTERPRISE 39, 43 (Raymond DeVries & Janardan Subedi eds., 1988) (same).

17. *See generally* Neil M. Richards & Daniel J. Solove, *Privacy's Other Path: Recovering the Law of Confidentiality*, 96 GEO. L.J. 123 (2007) (discussing various contexts, such as medical privacy and attorney-client privacy, where law relies fully or partly on duty-based approaches to protect privacy).

18. *See* Barbara J. Evans, *Rules for Robots, and Why Medical AI Breaks Them*, J.L. & BIOSCIENCES, Jan.-June 2023, at 12-15 [hereinafter Evans, *Rules for Robots*] (describing the duty-based privacy protections of medical privacy law and tracing them back 2,400 years to Hippocrates).

19. ALFRED I. TAUBER, PATIENT AUTONOMY AND THE ETHICS OF RESPONSIBILITY 18 (2005).

controllers is not as effective as placing the controllers under clear legal responsibilities to do the right thing. Medical privacy law presumes that patients in the health care setting might be too vulnerable to protect their own privacy by controlling access to their health data, and it instead relies on a duty-based ethic of beneficence ("ethic of responsibility") that places the burden on medical data handlers to "act as entrusted fiduciaries" for them.[20]

This Part discusses the two approaches and then turns to the question of which is better tailored to the challenge of protecting privacy in AI-enabled health care.

A. The Leading Privacy Paradigm, and How the Privacy Rule Violates It

Bioethicists embraced control-over-information theory after 1970, and a major federal research regulation from that era, the Common Rule, grants people a right of informed consent before identifiable health data and biospecimens are used in biomedical research.[21] The more recent "Information Privacy Law Project" endorses similar notice-and-consent norms to protect privacy in the modern "surveillance society," where retailers, employers, social media providers, law enforcement, private security services, and many other actors constantly collect and analyze personal data.[22] Control-over-information theory views privacy as something autonomous individuals, empowered by strong consent norms, can protect for themselves by vetoing unwanted access to personal data.[23] The phrase "de-identify or get

20. *See id.*

21. *See* THE PRIV. PROT. STUDY COMM'N, PERSONAL PRIVACY IN AN INFORMATION SOCIETY 280 (1977) (finding, in a study authorized under the Privacy Act of 1974, that health data were widely used without consent in medical research and public health studies during the 1970s and recommending that it would be ethical to seek consent before such uses); Basic HHS Policy for Protection of Human Research Subjects, 45 C.F.R. §§ 46.101-.124 (2022) (the Common Rule); *see also* Spector-Bagdady, *supra* note 8 (discussing regulations affecting "secondary use" of data—that is, use of information for a purpose other than the one for which it was originally collected—in biomedical research).

22. *See, e.g.*, Neil M. Richards, *The Information Privacy Law Project*, 94 GEO. L.J. 1087 (2006) (assessing the accomplishments and potential of the Information Privacy Law Project); DAVID LYON, SURVEILLANCE SOCIETY: MONITORING EVERYDAY LIFE 33-35, 114-18 (2001) (describing the modern surveillance society); FRANK PASQUALE, THE BLACK BOX SOCIETY: THE SECRET ALGORITHMS THAT CONTROL MONEY AND INFORMATION (2015) (same); Julie E. Cohen, *Privacy, Visibility, Transparency, and Exposure*, 75 U. CHI. L. REV. 181, 181-82, 186 (2008) (discussing the pervasive collection and use of data in modern surveillance societies).

23. *See* Solove, *supra* note 13, at 1109-10; *see also* Ferdinand Schoeman, *Privacy: Philosophical Dimensions of the Literature*, *in* PHILOSOPHICAL DIMENSIONS OF PRIVACY: AN ANTHOLOGY 1, 3 (Ferdinand David Schoeman ed., 1984); Deborah C. Peel, *Written Testimony Before the HIT Policy Committee*, ELEC. PRIV. INFO. CTR. (Sept. 18, 2009), https://epic.org/wp-content/uploads/privacy/medical/Peel_PPR%20Written%20testimony%20HIT%20Policy%20Committee.pdf [https://perma.cc/Y7NT-27QR] (framing privacy as "control of personal information"); Schwartz, *supra* note 14, at 820 (noting that individual control over one's data is central to the modern concept of data privacy).

consent" (DOGC) encapsulates what such norms typically require: get consent if the data are in a format that identifies the individual the data describe.[24]

Ethicists' major discontent with the Privacy Rule is that it is not a DOGC privacy scheme. The Privacy Rule is a federal medical privacy regulation administered by the U.S. Department of Health and Human Services (HHS). It regulates a narrowly defined class of "covered entities"—basically, private-sector actors that provide or pay for clinical health care services (physicians, clinics, hospitals, and health insurers), plus their "business associates."[25] Business associates are parties that obtain identifiable data from covered entities while performing professional or informational services for them.[26] Business associates include, for example, data processing companies hospitals hire to analyze patient data to look for ways to improve hospital efficiency or a law firm that receives patient data from a doctor who hired the firm to defend a malpractice suit.[27] Business associates become covered entities and must comply with the Privacy Rule when handling the data they receive while working for other covered entities.

The Privacy Rule "is exclusively mapped to and calibrated for the traditional health-care domain."[28] Its coverage leaves out many modern businesses commonly thought of as health-related, such as companies selling fitness trackers, direct-to-consumer genetic and other testing, and pharmaceutical or medical device companies that sell medical products as opposed to traditional health care services.[29] It governs how covered entities can use and disclose "protected health

24. *See* 45 C.F.R. § 46.102(e) (defining a "human subject" as a person about whom a researcher obtains identifiable information, implying that if the person's data are provided to the researcher in a de-identified format, the person is not a human subject from whom consent would be required under the Common Rule, *see id.* § 46.116(a)(1), so that the requirement is to de-identify data *or* obtain consent); *see also* 45 C.F.R. § 164.502(d)(2) (2022) (providing that the HIPAA Privacy Rule's requirements, including its authorization requirements, *id.* § 164.508, do not apply to information that has been de-identified, thus creating an either/or requirement to de-identify data *or* obtain authorization). *See generally* Kobbi Nissim & Alexandra Wood, *Is Privacy Privacy?*, PHIL. TRANSACTIONS ROYAL SOC'Y, Sept. 2018, at 11 (noting that "many privacy regulations require data providers to protect information that can be linked to an individual [i.e., identifiable information]," thus setting up an either/or choice to de-identify data *or* provide the required protections).

25. *See* 45 C.F.R. § 160.102 (2022) (providing that the HIPAA regulations, including the Privacy Rule, apply to health care providers such as physicians, clinics, hospitals, laboratories, and various other entities, such as insurers, that transmit "any health information in electronic form in connection with a transaction covered by this subchapter [the Administrative Simplification provisions of HIPAA]" and to their business associates); *see also id.* § 160.103 (defining the terms "covered entity" and "business associate").

26. *Id.* § 160.103.

27. Maggie Hales, *Lawyers as HIPAA Business Associates*, HIPAA E-TOOL (Mar. 23, 2021), https://thehipaaetool.com/lawyers-as-hipaa-business-associates/ [https://perma.cc/ME35-8V5R].

28. *See* Terry, *supra* note 7, at 202.

29. *See* 45 C.F.R § 160.103.

information" (PHI), which is the category of data, defined in the HIPAA statute, that the Privacy Rule protects.[30]

The Privacy Rule lulls casual observers into thinking it is a DOGC privacy scheme by allowing covered entities to disclose PHI pursuant to an individual authorization (HIPAA's name for consent) or if the data have been de-identified.[31] Then comes the betrayal: the regulation goes on to list twenty-five additional legal pathways for moving patients' health data into a wide variety of secondary uses without individual authorization and potentially in identifiable format, as seen in Table 1. The fact that unconsented data sharing is legal does not make it ethical in many people's minds.[32]

Each of the norms allowing unconsented data disclosure sets out alternative privacy protections to apply in lieu of individual authorization.[33] How well the Privacy Rule protects privacy depends on whether these alternative protections are adequate. Some of the norms include meaningful alternative protections, and for those, the alarm expressed about HIPAA's unconsented data disclosures seems overwrought. For other norms, the alternative privacy protections are weak and are a worthy focus for reforms.[34] The next Section briefly surveys the Privacy Rule's alternative protections.

30. *See id.* (defining "protected health information," the information that the HIPAA Privacy Rule protects, as "individually identifiable health information" and defining the term "health information" for purposes of the HIPAA Privacy Rule); Genetic Information Nondiscrimination Act (GINA) of 2008, 42 U.S.C. § 1320d(4) (reflecting the original 1996 HIPAA statute's definition of "health information" as "any information, whether oral or recorded in any form or medium, that—(A) is created or received by a health care provider, health plan, public health authority, employer, life insurer, school or university, or health care clearinghouse; and (B) relates to the past, present, or future physical or mental health or condition of an individual, the provision of health care to an individual, or the past, present, or future payment for the provision of health care to an individual"); *id.* § 1320d-9(b)(1) (stating, in a new section introduced by GINA, that Congress deems "genetic information," as broadly defined by GINA at 42 U.S.C. § 300gg-91, to be health information for purposes of making it subject to HIPAA's privacy protections); *id.* § 1320d-9(a) (expanding the definition of "health information" that HIPAA protects to include genetic information).

31. *See* 45 C.F.R. § 164.502(a)(1)(iv) (2022) (allowing PHI to be released with individual authorization); *see also id.* § 164.508(c) (describing requirements for a valid individual authorization, which is HIPAA's term for a consent); *id.* § 164.502(d) (allowing de-identified data to be used and disclosed without individual authorization).

32. *See* IOM, PRIVACY REPORT, *supra* note 5, at 81-86 (reporting the results of various surveys of patients' attitudes about health data privacy which suggest patients are not entirely comfortable with some of the unconsented data uses that HIPAA allows).

33. *See infra* Section I.B (summarizing these alternative protections).

34. *See infra* Section IV.B (recommending specific reforms).

Table 1. The Privacy Rule's 27 Norms Allowing PHI to Be Used and Disclosed[35]

Norms for disclosure with de-identification or consent
1. Can use and disclose ("share") with individual authorization
2. Can share data that have been de-identified
a. Safe harbor de-identification
b. Statistical de-identification
Norms on disclosures to patients/executors
3. MUST disclose designated record set to the individual upon request, under HIPAA's right of access to one's own data
4. Can disclose additional data to the individual
5. Can disclose to patient's legal representative after death
Seven norms allowing unconsented disclosure and use, not subject to the minimum necessary standard but subject to alternative protections
6. Can share patient data with a health care provider for use in treating a patient—*any* patient
7. Can share data with HHS for regulatory compliance purposes
8. Can share as required for HIPAA compliance
9. Can share with agencies that detect abuse and neglect
10. Can share for judicial and regulatory proceedings
11. Can share for law enforcement purposes
12. Can share if required by law *cont'd next page*

35. *See* 45 C.F.R.§ 164.502(a)(1)(iv); *see also id.* § 164.508 (describing requirements for a valid authorization) [Norm 1]; *id.* §§ 164.502(d), .514(a); *id.* § 164.514(b)(1) (statistical de-identification); *id.* § 164.514(b)(2)(i) (safe harbor de-identification) [Norm 2]; *id.* § 164.524 (providing an individual access right); *id.* § 164.501 (defining the "designated record set" that is subject to mandatory disclosure to the individual or to a third party the individual specifies) [Norm 3]; *id.* § 164.502(a)(1), (b)(2)(ii) [Norm 4]; *id.* §§ 160.103, 164.502(f), (g)(1) [Norm 5]; *id.* § 164.502(a)(1)(ii); *512-May a Provider Disclose Information About an Individual to Another Provider*, U.S. DEP'T HEALTH & HUM. SERVICES (Jan. 13, 2009) [hereinafter HHS, FAQ 512], https://www.hhs.gov/hipaa/for-professionals/faq/512/under-hipaa-may-a-health-care-provider-disclose-information-requested-for-treatment/index.html [https://perma.cc/W3R7-CYUQ] (clarifying that, except for psychotherapy notes, a HIPAA-covered doctor may disclose a patient's information to another doctor without individual authorization for use in treating "another patient"—not necessarily a family member of the person whose data is disclosed) [Norm 6]; 45 C.F.R. § 164.502(b)(2)(iv) [Norm 7]; *id.* § 164.502(b)(2)(vi) [Norm 8]; *id.* § 164.512(c) [Norm 9]; *id.* § 164.512(e) [Norm 10]; *id.* § 164.512(f) [Norm 11]; *id.* § 164.512(a) [Norm 12].

Table 1, *cont'd*. The Privacy Rule's 27 Norms Allowing PHI to Be Used and Disclosed[36]

Fifteen norms allowing unconsented disclosure and use, subject to the minimum necessary standard*
13. Can share pursuant to waiver approved by IRB/privacy board
14. Can share data for payment and health care operations and health care quality improvement studies
15. Can share with public health authorities and their contractors
16. Can share with FDA-regulated entities to aid their compliance with FDA-required activities
17. Can share with health oversight agencies
18. Can share a limited data set subject to data use agreement
19. Can share with people exposed to communicable disease
20. Can share with employers for workplace safety/exposures
21. Can share to facilitate dignified burial of the deceased
22. Can share to facilitate organ transplants
23. Can share for fundraising, but must allow an opt-out
24. Can share for certain insurance underwriting purposes
25. Can share to avert serious threats to health or safety
26. Can share for special governmental functions (e.g., military)
27. Can share for use in workers' compensation cases
*** Minimum necessary disclosures can include identifiers if they are necessary to fulfill the purpose of the disclosure**

36. *See* 45 C.F.R. § 164.512(i) [Norm 13]; *id.* §§ 164.502(a)(1)(ii), .506. *But see* Proposed Modifications to the HIPAA Privacy Rule to Support, and Remove Barriers to, Coordinated Care and Individual Engagement, 86 Fed. Reg. 6446 (proposed January 21, 2021) (controversially proposing to exclude these disclosures from HIPAA's minimum necessary standard) [Norm 14]; 45 C.F.R. §§ 164.512(b)(1)(i), (ii), 164.514(d)(3)(iii)(A), 164.514(h)(2)(ii), (iii) [Norm 15]; *id.* § 164.512(b)(iii) [Norm 16]; *id.* §§ 164.501, .512(d) (defining oversight agencies) [Norm 17]; *id.* § 164.514(e)(3)(i), 164.514(e)(4) [Norm 18]; *id.* § 164.512(b)(iv) [Norm 19]; *id.* § 164.512(b)(v) [Norm 20]; *id.* § 164.512(g) [Norm 21]; *id.* § 164.512(h) [Norm 22]; *id.* § 164.514(f) [Norm 23]; *id.* § 164.514(g) [Norm 24]; *id.* § 164.512(j) [Norm 25]; *id.* § 164.512(k) [Norm 26]; *id.* § 164.512(l) [Norm 27].

B. The Privacy Rule's Alternative Privacy Protections for Unconsented Data Flows

An alternative protection common to many of the norms (Norms 13-27) in Table 1 is HIPAA's "minimum necessary" standard, which allows covered entities to use or disclose only the least amount of information needed to support the purpose for which the data were requested.[37] This standard traces back to a Code of Fair Information Practices (FIPs) set out in a 1973 report by the U.S. Department of Health, Education, and Welfare Secretary's Advisory Committee on Automated Personal Data Systems.[38] The Federal Privacy Act of 1974[39] codified these FIPs, which also influenced privacy laws around the world including the EU's GDPR, which applies this same standard but calls it "data minimisation."[40] The minimum necessary standard prevents information from being shared unless it is truly needed for the intended data use. Still, identifiers can be shared under this standard when necessary to the task for which the data were requested—for example, if identifiers are needed to let clinical data be correlated with other important sources of data, such as data on social determinants of health.[41]

---

37. *See* 45 C.F.R. § 164.502(b); *see also id.* § 164.514(d) (further explaining how to comply with the minimum necessary standard).

38. *See* U.S. DEP'T OF HEALTH, EDUC., & WELFARE, RECORDS, COMPUTERS, AND THE RIGHTS OF CITIZENS 41 (1973) (announcing an influential set of "safeguards for personal privacy," commonly known as fair information practices (FIPs), based on five principles); *see also Confidentiality of Individually Identifiable Health Information: Recommendations of the Secretary of Health and Human Services, Pursuant to Section 264 of the Health Insurance Portability and Accountability Act of 1996*, U.S. DEP'T HEALTH & HUM. SERVICES (Sept. 10, 1997) [hereinafter *HHS, 1997 Recommendations*], https://aspe.hhs.gov/report/confidentiality-individually-identifiable-health-information [https://perma.cc/D5RQ-RLBH] (referring to this principle from the 1973 HEW Code of FIPs in the roadmap for the HIPAA Privacy Rule).

39. *See* 5 U.S.C. § 552a(d).

40. *See* 5 U.S.C. § 552a(a), (d) (2012), *amended by* 5 U.S.C. § 552a (Supp. III 2016) (Federal Privacy Act); Fred H. Cate, *The Failure of Fair Information Practice Principles*, *in* CONSUMER PROTECTION IN THE AGE OF THE 'INFORMATION ECONOMY' 341, 346 (Jane K. Winn ed., 2006) (tracing subsequent development of FIPs, including access rights, after the 1973 HEW Code of FIPs); Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the Protection of Natural Persons with Regard to the Processing of Personal Data and on the Free Movement of Such Data, and Repealing Directive 95/46/EC (General Data Protection Regulation), art. 5, 2016 O.J. (L119) [hereinafter Regulation 2016/679] (stating five principles of processing of personal data, including the principle that data should be "adequate, relevant and limited to what is necessary in relation to the purposes for which they are processed ('data minimisation')").

41. *See* Barbara J. Evans, *Much Ado About Data Ownership*, 25 HARV. J.L. & TECH. 69, 93-94 (2011) [hereinafter Evans, *Data Ownership*] (discussing why identifiers are important to link data from multiple sources); 45 C.F.R. § 164.502(b) (providing that the minimum necessary standard applies as a general rule, *id.* § 164.502(b)(1), and then enumerating specific situations in which it does not apply, *id.* § 164.502(b)(2)); *id.* § 164.514(d) (outlining what covered entities must do to comply with the minimum necessary standard when it does apply); *id.* § 164.514(d)(3)(ii) (requiring covered entities to develop criteria for assessing whether requested information is "reasonably necessary to accomplish the purpose for which disclosure is sought" and "[r]eview requests for disclosure[s] on an individual basis in accordance with such criteria"); *id.* § 164.514(d)(3)(ii)(A)-(B).

Some norms set clear substantive or procedural limitations on what can be disclosed. Thus, Norm 11 on disclosures for law enforcement purposes restricts both the *types* of medical information that can be provided (e.g., blood type, day and time of treatment, and time of death) and the *purposes* for which data can be provided without individual authorization.[42] Norm 10 on disclosures for judicial and regulatory proceedings allows covered entities to release PHI only if there is a court order, subpoena, or other process affording due process to the person whose data are shared without consent.[43]

Several of the norms protect privacy by restricting who can receive the data, directing disclosures to parties under independent legal duties of confidentiality (i.e., apart from any duties the Privacy Rule imposes). Norm 6 broadly allows PHI about one patient to be disclosed to aid in treating a different patient (for example, a later patient exhibiting similar symptoms or genotype).[44] The minimum necessary standard does not apply, so the disclosure could be quite detailed or even identifiable.[45] Notably, however, these disclosures can only be made to health care providers, who are HIPAA-regulated and subject to additional duties of confidentiality. State licensing statutes for physicians, nurses, and health care facilities place them under strong confidentiality norms enforceable through disciplinary sanctions or loss of license.[46] Professional ethics standards, like those of the American Medical Association, add soft-law norms on top of the legally enforceable confidentiality requirements.[47] Further, a health care provider breaching patient confidentiality can lose eligibility to receive Medicare payments or violate hospital accreditation standards—sanctions that pose existential threats to the provider's commercial viability.[48] Thus, Norm 6 disclosures can only go to parties that are

42. *See* 45 C.F.R. § 164.512(f)(1)-(6).

43. *See id.* § 164.512(e)(ii)-(vi).

44. *See id.* § 164.502(a)(1)(ii); *see also* HHS, FAQ 512, *supra* note 35 (clarifying that, except for psychotherapy notes, a HIPAA-covered doctor may disclose a patient's information to another doctor without individual authorization for use in treating "another patient," not necessarily a family member of the person whose data is disclosed).

45. *See* 45 C.F.R. § 164.502(b)(2)(i).

46. *See* BARRY R. FURROW ET AL., HEALTH LAW 117 (8th ed., 2018) (listing fiduciary duties of licensed health care professionals, including "a duty to hold in confidence information learned about a patient in a treatment relationship," "a duty to provide care in a manner consistent with the standard of care," "a duty to obtain a patient's informed consent prior to treatment," and "a duty not to abandon a patient with whom a treatment relationship has been formed").

47. *See, e.g.*, *Confidentiality, Opinion 3.2.1*, AMA CODE MED. ETHICS, https://www.ama-assn.org/delivering-care/ethics/confidentiality [https://perma.cc/773Y-AVUL] (last visited Sept. 23, 2023) ("Physicians in turn have an ethical obligation to preserve the confidentiality of information gathered in association with the care of the patient.").

48. *See HIPAA Violations & Enforcement*, AM. MED. ASS'N, https://www.ama-assn.org/practice-management/hipaa/hipaa-violations-enforcement [https://perma.cc/49QE-3G74] (last visited Sept. 23, 2023) (noting that a provider's breach of patient privacy can

"information fiduciaries," to use Jack Balkin's phrase.[49] His concept of an information fiduciary involves more limited duties than those of doctors and lawyers, who have duties of loyalty and duties of care that go far beyond a mere duty to handle information carefully.[50] The recipients of Norm 6 treatment disclosures, as health care professionals, clearly qualify as information fiduciaries.

Other norms direct disclosures to federal agencies governed by the Privacy Act of 1974, to state agencies governed by similar state privacy frameworks, and to courts and other bodies subject to strong norms of privacy protection for data they receive.[51] By authorizing disclosures to a narrow set of trustworthy actors, the Privacy Rule offers a degree of comfort that data disclosed without consent will remain protected in the hands of the recipients.

Unfortunately, the same cannot be said of all of the disclosures the Privacy Rule allows. Some of its norms allow unconsented data disclosures to parties—such as commercial researchers, external software service providers, and private-sector drug and device manufacturers—that are neither HIPAA-covered nor subject to other laws on confidentiality.[52] The prospect of unconsented sharing of data, potentially in identifiable form, naturally alarms people unless the Privacy Rule sets

---

result in exclusion from the Medicare program, among other sanctions); *see also Medical Record—Security*, JOINT COMM'N, https://www.jointcommission.org/standards/standard-faqs/hospital-and-hospital-clinics/information-management-im/000001462/ [https://perma.cc/W7FW-XX7P] (last visited Sept. 23, 2023) (noting, as part of the standards for accreditation by The Joint Commission, an influential private accreditation body for hospitals, clinical laboratories, and other health care providers, the need for strict privacy and data security compliance).

49. *See* Jack M. Balkin, *Information Fiduciaries and the First Amendment*, 49 U.C. DAVIS L. REV. 1183, 1205-09 (2016) [hereinafter Balkin, *Information Fiduciaries*] (introducing and developing the concept of an information fiduciary); *see also* Jack M. Balkin, *The Three Laws of Robotics in the Age of Big Data*, 78 OHIO ST. L.J. 1217, 1227, 1229 (2017) [hereinafter Balkin, *Three Laws*] (calling for controllers of AI/ML algorithms to act as information fiduciaries of their clients, customers, and end-users to whose data they have access).

50. *See* FURROW ET AL., *supra* note 46, at 117 (itemizing the various fiduciary duties of physicians); Balkin, *Three Laws*, *supra* note 49, at 1229 (explaining the lesser requirements of an information fiduciary and noting, for example, that monetizing personal data is central to the business model of many entities operating AI/ML in social media contexts, where harnessing data to recoup expenses or make a profit does not by itself violate the social media operator's information fiduciary duties to its customers and concluding that social media providers have lower duties of care to their users, and the users repose less trust in the operators, than what is expected of health care professionals, who have fiduciary duties to warn and to act in patient's best interests); *see also* Balkin, *Information Fiduciaries*, *supra* note 49, at 1183 (discussing the concept of an information fiduciary).

51. *See* 5 U.S.C. § 552a(d); *see, e.g.*, *supra* Table 1, Norm 7 (disclosures to the U.S. Department of Health & Human Services); *id.* Norm 15 (disclosures to public health authorities); *id.* Norm 17 (disclosures to health oversight agencies); *see also id.* Norms 9-12 (allowing disclosures to various agencies charged with detecting abuse and neglect, courts and administrative agencies, law enforcement bodies, and others as required by law, subject to various procedural protections such as verification procedures or obtaining a subpoena).

52. *See, e.g.*, *supra* Table 1, Norm 13 (allowing unconsented disclosure with waiver approved by an IRB/privacy board); *id.* Norm 16 (allowing unconsented disclosure to FDA-regulated entities for activities that the FDA *requires* them to do).

meaningful limits on how recipients can use the data after they receive it. HHS felt it lacked jurisdiction to regulate downstream uses of PHI by non-HIPAA-covered data recipients.[53] HHS seriously considered "limiting the type or scope of the disclosures permitted" but felt forced to allow wide data sharing to promote "key public goals such as research, public health, and law enforcement."[54] HHS wanted recipients of PHI to follow "safeguards, including restrictions on re-disclosure, to ensure that individual subjects are not harmed"[55] but felt it lacked authority to impose the safeguards.

HHS's position was disingenuous. HHS has ample power to require *covered entities* to impose information fiduciary duties, by means of contract, on parties to whom they disclose PHI. HHS employed this approach in one of the Privacy Rule's norms. Norm 18—allowing unconsented disclosure of almost fully de-identified limited data sets—requires contractual use limitations in the form of a Data Use Agreement (DUA) restricting the use, redisclosure, or re-identification of the data and requiring various other privacy protections.[56] HHS could have used this same approach in HIPAA's other informational norms allowing nonconsensual disclosure of PHI but chose not to do so. This was an unforced error: HHS had the necessary jurisdiction but chose not to exercise it, leaving lingering gaps in data protection under the Privacy Rule.[57]

Through the lens of modern control-over-information theory, the Privacy Rule strikes many observers as gravely flawed—a flagrant display of the "health-care data protection exceptionalism" of medical privacy law.[58] Medical privacy law is a common shorthand for a large body of state and federal laws affecting the use and sharing of data in traditional clinical health care settings, such as doctors' offices, clinics, hospitals, rehabilitation and nursing facilities, and clinical laboratories.[59] Professor Terry rightly observes that calling it "privacy"

53. *See* Standards for Privacy of Individually Identifiable Health Information, 64 Fed. Reg. 59918, 59923 (proposed Nov. 3, 1999) (to be codified at 45 C.F.R. pts. 160-164) ("[T]he proposed regulation does not directly cover many of the persons who obtain identifiable health information from the covered entities . . . . [W]e are, therefore, faced with creating new regulatory permissions for covered entities to disclose health information, but cannot directly put in place appropriate restrictions on how many likely recipients of such information may use and re-disclose such information.").

54. *Id.*

55. *See HHS, 1997 Recommendations*, *supra* note 38; *see also* Standards for Privacy of Individually Identifiable Health Information, 64 Fed. Reg. at 59968.

56. *See* 45 C.F.R. § 164.514(e)(3)(i), (e)(4) (2022).

57. *See infra* Section IV.B.

58. *See* Terry, *supra* note 7, at 143.

59. *See* Evans, *Rules for Robots*, *supra* note 18, at 12-15 (summarizing the web of state and federal privacy laws and other general health laws that, taken together, establish the medical privacy framework that governs data privacy in clinical health care settings).

law is a misnomer: for example, it does not restrict data collection.[60] Instead, it "employs a downstream data protection model ('confidentiality')" after data collection already has occurred.[61] Yet to call the Privacy Rule a confidentiality law is also a misnomer. The Privacy Rule does not impose new duties of confidentiality within its own text—neither on the medical data handlers it regulates nor on the long list of outside parties to which it lets regulated entities disclose data without patient permission.[62] The privacy protection (if any) comes not from the Privacy Rule itself but from other laws placing covered entities—and many but not all of the data recipients with whom they can share data—under information fiduciary duties.[63]

If the Privacy Rule is not really a privacy law or a confidentiality law, then what is it? It is best described as what Helen Nissenbaum calls a contextual privacy scheme. As summarized in Table 1, it lists "informational norms"—a set of data flows considered appropriate and necessary in and around one specific context (in this case, clinical health care).[64] Information sharing that is appropriate in one context might be highly inappropriate in other contexts. Thus, inquiring about people's annual income is appropriate when they apply for a home loan but not when meeting them for the first time at a cocktail party. The Privacy Rule deems 27 information flows to be permissible in and around the traditional clinical health care setting, and it provides instructions to follow for each such information flow. To many observers, that is not a privacy regulation at all, but a data sharing regulation, designed for a highly unusual context where many data handlers have at least some preexisting legal duties to treat data confidentially. The next Section traces how the Privacy Rule came to be the way it is.

C. The Privacy Rule's Policy Rationale

HIPAA was an insurance statute, not a privacy statute, but its framers saw health insurance claims shifting to electronic formats that posed novel risks to medical privacy. Legislators inserted a few lines deep in the statute to show they were aware of the privacy

60. *See* Terry, *supra* note 7, at 165 (noting that "the culture of medicine has seemed to favor collecting *everything*" that might remotely advance its mission of treating the sick).

61. *Id.* at 164.

62. *See infra* Section III.B (recommending reforms to address this gap in the Privacy Rule's protections).

63. *See supra* notes 46-48 and accompanying text (describing laws placing data recipients under independent duties of confidentiality).

64. *See supra* Table 1. *See generally* HELEN NISSENBAUM, PRIVACY IN CONTEXT: TECHNOLOGY, POLICY, AND THE INTEGRITY OF SOCIAL LIFE 129-57 (2010) (describing contextual privacy schemes); Adam Barth et al., *Privacy and Contextual Integrity: Framework and Applications*, *in* PROCEEDINGS OF THE IEEE SYMPOSIUM ON SECURITY AND PRIVACY 184, 184-98 (2006) (same); Helen Nissenbaum, *Privacy as Conceptual Integrity*, 79 WASH. L. REV. 119 (2004) (same).

concerns.[65] Those lines called for Congress to address medical privacy in separate legislation based on recommendations HIPAA ordered the U.S. Department of Health and Human Services (HHS) to provide by 1997.[66] If Congress failed to enact privacy legislation by August 21, 1999, HIPAA authorized HHS to promulgate federal health privacy regulations.[67]

HHS delivered its privacy recommendations to Congress on time in 1997.[68] After viewing them, Congress chose not to legislate. It is a fool's errand to read too much meaning into Congress's actions, let alone its failures to act. By 1997, data privacy was already divisive. "[D]elegations can be particularly useful to Congress with respect to divisive issues" best punted to unelected agency officials.[69] Whether dithering or purposeful, Congress left HHS to formalize its 1997 recommendations in regulations, which came to be known as the HIPAA Privacy Rule.[70]

HHS proceeded with pained reluctance, protesting that HIPAA does not confer the full set of authorities it takes to do a good job regulating medical privacy.[71] HHS exhorted Congress to pass new legislation.[72] Not getting a response, HHS soldiered up and conducted one of the most meticulously researched, vigorous, and

65. *See* U.S. DEP'T OF HEALTH & HUM. SERVS., SUMMARY OF THE HIPAA PRIVACY RULE 1-2 (2003), https://www.hhs.gov/sites/default/files/privacysummary.pdf?language=es [https://perma.cc/7WT7-VXF8] (discussing the Administrative Simplification provisions at §§ 261-264 of HIPAA, authorizing the Secretary of HHS to set standards for electronic exchange, security, and privacy of health information).

66. Health Insurance Portability and Accountability Act of 1996, Pub. L. No. 104-191, § 264(a)-(c), 110 Stat. 1936, 2033-34 (codified as amended at 42 U.S.C. 1320d-2).

67. *Id.* § 264(c); *see* U.S. DEP'T OF HEALTH & HUM. SERVS., *supra* note 65, at 1-2.

68. *See HHS, 1997 Recommendations*, *supra* note 38.

69. Margaret H. Lemos, *The Consequences of Congress's Choice of Delegate: Judicial and Agency Interpretations of Title VII*, 63 VAND. L. REV. 363, 369-70 (2010).

70. *See* 45 C.F.R. pts. 160, 164 (2022); *see also* Barbara J. Evans, *The Interplay of Privacy and Transparency in Health Care*, *in* TRANSPARENCY IN HEALTH AND HEALTH CARE IN THE UNITED STATES 30, 33-38 (Holly Fernandez Lynch, I. Glenn Cohen, Carmel Shachar & Barbara J. Evans eds., 2019) (tracing how the Privacy Rule incorporated the major principles outlined in *HHS 1997 Recommendations, supra* note 38).

71. *See* Standards for Privacy of Individually Identifiable Health Information, 64 Fed. Reg. 59918, 59923 (proposed Nov. 3, 1999) (to be codified at 45 C.F.R. pts. 160-164) ("We believed then, and still believe, that there is an urgent need for legislation to establish comprehensive privacy standards for all those who pay and provide for health care, and those who receive information from them. This proposed rule implements many of the policies set forth in the [HHS 1997] Recommendations. However, the HIPAA legislative authority is more limited in scope than the federal statute we recommend, and does not always permit us to propose the policies that we believe are optimal. Our major concerns with the scope of the HIPAA authority include the limited number of entities to whom the proposed rule would be applicable, and the absence of strong enforcement provisions and a private right of action for individuals whose privacy rights are violated.").

72. *See id.*

participatory rulemaking processes in U.S. history, addressing over 52,000 public comments[73] in mammoth rulemaking documents[74] and reopening its December 2000 final rule for a second round of comments[75] and amendments.[76]

The 1997 recommendations HIPAA required HHS to prepare later served as the template for the Privacy Rule and are the key to understanding its rationale.[77] HHS framed the challenge as being to preserve longstanding medical privacy norms on an altered landscape of massively decentralized health care mediated by electronic information flows. At the outset, HHS examined the informational norms that traditionally existed in clinical health care.[78] HHS indicated it had "carefully examined the many uses that the health professions, related industries, and the government make of health information."[79] Medical privacy was historically a state-law concern.[80] HHS ascribed the need for federal intervention to recent changes in health care delivery and payments: "There was a time when our health care privacy was protected by our family doctors—who kept hand-written records about us sealed away in big file cabinets," but the single-provider delivery model of the past gave way to vast "networks of insurers and health care professionals" linked by electronic transfers of "secrets . . . from doctors to hospitals to insurance companies."[81] The new federal medical privacy law must protect patients' privacy while still preserving essential information flows—including many unconsented data flows—on which clinical health care has always depended.

---

73. *See* Ko, *supra* note 4, at 500.

74. *See, e.g.*, Standards for Privacy of Individually Identifiable Health Information, 65 Fed. Reg. 82462, 82462-829 (Dec. 28, 2000) (to be codified at 45 C.F.R. pts. 160, 164) (finalizing the HIPAA Privacy Rule in December 2000); Standards for Privacy of Individually Identifiable Health Information, 64 Fed. Reg. at 59918-60065.

75. *See* Standards for Privacy of Individually Identifiable Health Information, 66 Fed. Reg. 12738 (proposed Feb. 28, 2001) (to be codified at 45 C.F.R. pts. 160, 164).

76. *See* Standards for Privacy of Individually Identifiable Health Information, 67 Fed. Reg. 53182 (Aug. 14, 2002) (to be codified at 45 C.F.R. pts. 160, 164) (amending various provisions of the 2000 final Privacy Rule).

77. *See HHS, 1997 Recommandations*, *supra* note 38.

78. *See infra* text accompanying notes 80-86 (discussing traditional information norms of clinical health care).

79. *See HHS, 1997 Recommendations*, *supra* note 38, § I.

80. *See generally* P. JON WHITE ET AL., PRIVACY AND SECURITY SOLUTIONS FOR INTEROPERABLE HEALTH INFORMATION EXCHANGE: REPORT ON STATE MEDICAL RECORD ACCESS LAWS (2009), https://www.healthit.gov/sites/default/files/290-05-0015-state-law-access-report-1.pdf [https://perma.cc/5ARR-ZKCT] (providing a multistate survey of state laws governing clinical health care records); *see also supra* notes 46-48 and accompanying text (discussing state requirements for licensed health care professionals, including fiduciary duties of confidentiality).

81. *See HHS, 1997 Recommendations*, *supra* note 38, § I.

Informational norms of traditional health care. The informational norms of traditional medical privacy law reflect the context for which it was designed. That context is the clinical health care encounter, a transaction in which patients disclose highly personal information about themselves to a health care professional, who—quite apart from privacy laws—is already under legally enforceable fiduciary duties to protect the confidentiality of that information and to act in the patient's best interests.[82] As part of the transaction, the health care professional combines the patient's personal information with other sources of medical knowledge to draw expert inferences about the patient's health. These inferences are then shared with the patient, either verbally or through actions the professional takes (such as writing a prescription or ordering a course of treatment) during and after the clinical encounter.

Information about the individual patient is not the only, or necessarily the most important, input to clinical health care. If that were true, clinical inferences would resemble the old joke about the management consultant who borrows your watch and tells you what time it is.[83] In the joke, the "watch" is detailed company-specific data that the client company laboriously develops and feeds to the consultant. In clinical health care, the "watch" is detailed data about themselves that patients supply to the clinician directly or by letting the clinician examine them. In reality, clinical health care incorporates many additional information flows less visible to the patient but equally crucial to the patient's health care.

Clinical inferences draw on a base of general medical knowledge, which is an accretion of information drawn from *other people's* health experiences. Only for the simplest maladies (such as when a patient's finger has a splinter in it that obviously should not be there) can health care providers treat one patient in informational isolation from others. Rare disease sufferers can attest that clinical inference often fails for " 'n-of-1' single-instance medical mysteries" when a patient presents with a novel condition for which the clinician lacks comparators: it is difficult for doctors to infer what might be wrong if they never saw a similar malady before.[84]

Traditionally, the sharing of data about similar and contrasting cases took place inside the family doctor's head, without patients' consenting to have their past medical experiences considered or their

82. *See supra* notes 46-48 and accompanying text.

83. *See* Katie Hope, *What Does a Management Consultant Do Anyway?*, BBC NEWS (Jan. 5, 2016), https://www.bbc.com/news/business-35220061 [https://perma.cc/J82Y-J3B5].

84. *See* Edward Hancock, *Matt Might: Genetic Testing and 'Crowdscreening' Enabling Precision Medicine, Faster Research*, SEVEN BRIDGES (Dec. 4, 2015), https://www.sevenbridges.com/matt-might-genetic-testing-crowdscreening-phenotypes-changing-medical-research-better/ [https://perma.cc/QHS7-UP3L] (discussing Matt Might's quest for diagnosis and treatment of his son who had the first recorded case of a novel N-glycanase (NGLY1) deficiency, a congenital disorder of glycosylation).

files consulted as their doctor treated other patients.[85] Patients have never had a GDPR-style "right to be forgotten" while their doctor is treating other patients.[86] Their doctor was an information fiduciary, holding their data in the privacy of the doctor's intuition.[87] The resulting intuition, however, was freely shared with other patients.

Modern health care distills many sources of general knowledge—all incorporating data about other people—such as case reports describing other physicians' experiences while treating their patients, larger observational studies recording which treatments worked and did not work for patients with various characteristics, results of systematic clinical research studies, and public health surveillance reports of diseases circulating in the patient's location.[88] No individual is an island in clinical health care: health care providers use information about other people past and present when treating the patient in front of them. Then, after treatment takes place, people's data are shared with insurers and other payors to facilitate billing; with regulators, courts, and oversight bodies charged with maintaining safety standards; and with various other actors involved in operating the health care system.[89] Each treatment encounter depends on all these information flows. The patient's personal information is just one small part of the essential informational ecosystem supporting the day-to-day provision of health care.

These flows serve what Faden et al. have characterized as a "norm of common purpose" that is "similar to what John Rawls calls the principle of the common good, a principle presiding over matters that affect the interests of everyone."[90] None of us can receive optimal health care in an informational vacuum where our physicians' advice is uninformed by the treatment experiences of prior patients, whether shared through explicit data flows or implicitly in our care providers' memories and intuitions. In turn, each patient's treatment experiences

85. *See HHS, 1997 Recommendations*, *supra* note 38, § I (describing the traditional informational ecosystem of health care).

86. *Cf.* Regulation 2016/679, *supra* note 40, art. 17 (granting individuals a "right to erasure ('right to be forgotten')").

87. *See supra* notes 46-48 and accompanying text.

88. *See* U.S. PREVENTIVE SERVS. TASK FORCE, GUIDE TO CLINICAL PREVENTIVE SERVICES (2d ed. 1996) (comparing various types of study design based on their perceived evidentiary weight).

89. *See HHS, 1997 Recommendations*, *supra* note 38, § I.I (enumerating national priority activities for which unimpeded data access is deemed to provide important social benefits).

90. Ruth R. Faden et al., *An Ethics Framework for a Learning Health Care System: A Departure from Traditional Research Ethics and Clinical Ethics*, 43 HASTINGS CTR. REP. S16, S23 (2013) (emphasis omitted).

contribute to the general knowledge base for others. There is "reciprocity of advantage" because all patients stand to benefit from the burdens placed on their own and other people's privacy.[91]

Translating traditional norms to the modern digital health care environment. To replicate this same reciprocity in a decentralized, networked modern health care industry, HHS recommended in 1997 that providers should be "permitted to disclose health information without patient authorization to provide health care to *any* patient."[92] HHS noted that such disclosures were routine in traditional health care practice and are allowed by various state and federal medical confidentiality laws, which often direct such disclosures to another health care professional—in other words, to someone already under a strong fiduciary duty of confidentiality not just to their own patient but to other people whose data makes its way into the patient's medical records.[93]

The HIPAA Privacy Rule's so-called "treatment exception" implements this longstanding informational norm.[94] It is the broadest and least restricted pathway for sharing patients' PHI in the entire regulation.[95] HHS confirmed in later guidance that this exception permits broad disclosure of patients' information, without individual authorization, to other health care providers for treatment of "another patient" who might (or might not) be a family member of the first patient.[96] Your data can be used, without your consent, to treat a total stranger.[97] The HIPAA Privacy Rule did not create this norm, which replicates information flows long existing inside the brains of isolated family physicians and modernizes them for today's dense, decentralized provider networks.

At first glance, the European Union's (EU) General Data Protection Regulation (GDPR) seems to lack a similar provision allowing one

91. *See* Eric R. Claeys, *Takings, Regulations, and Natural Property Rights*, 88 CORNELL L. REV. 1549, 1587-89, 1619-21 (2003) (tracing the "reciprocity of advantage" or "common benefit of all" concepts in nineteenth and early twentieth-century state and federal cases that justified incursions on one person's rights in contexts where the individual gains the benefit of similar burdens placed on others).

92. *See HHS, 1997 Recommendations*, *supra* note 38, § II.E (emphasis added).

93. *See id.*; *see also supra* notes 46-48 and accompanying text (describing various laws imposing duties of confidentiality on health care providers).

94. *See supra* Table 1, Norm 6.

95. *See* Letter from William W. Stead, Chair, Nat'l Comm. on Vital & Health Stat., to Honorable Sylvia M. Burwell, Secretary, U.S. Dep't of Health & Hum. Servs. app. A (Nov. 9, 2016), https://www.ncvhs.hhs.gov/wp-content/uploads/2013/12/2016-Ltr-Privacy-Minimum-Necessary-formatted-on-ltrhead-Nov-9-FINAL-w-sig.pdf [https://perma.cc/3SLU-PH34] (listing various Privacy Rule norms allowing unconsented disclosure and use of data); *see also* Barbara J. Evans & Gail P. Jarvik, *Impact of HIPAA's Minimum Necessary Standard on Genomic Data Sharing*, 20 GENETICS MED. 531, 532-34 (2018) (discussing nonconsensual data sharing under the Privacy Rule).

96. *See* HHS, FAQ 512, *supra* note 35 (clarifying that the treatment exception allows unconsented data disclosures for the benefit of "another patient").

97. *See id.*

patient's data to be used in treating others.[98] The reality is more nuanced. In the clinical care context that the Privacy Rule regulates, it is in some respects *stronger* than the GDPR. The Privacy Rule sets a federal floor of medical privacy protections; states are free to set higher standards but cannot go lower.[99] In contrast, the GDPR grants the twenty-seven EU Member States leeway to go higher *or lower* than the GDPR's baseline consent standard when establishing their own medical privacy laws.[100] A 2021 report for the European Commission describes many instances where Member States allow unconsented flows of clinical health data, often resembling the Privacy Rule's informational norms in Table 1.[101] EU Member States can and do provide exceptions allowing health care providers to obtain and process patients' sensitive health information "for others than the data subject . . . if a significant medical interest prevails."[102] The Privacy Rule's treatment exception is consistent with data sharing practices in these EU Member States.[103]

98. Regulation 2016/679, *supra* note 40; *see also* CONSUMERS, HEALTH, AGRIC. & FOOD EXEC. AGENCY, EUROPEAN COMM'N, ASSESSMENT OF THE EU MEMBER STATES' RULES ON HEALTH DATA IN THE LIGHT OF GDPR 9, 23 (2021) [hereinafter EU MEMBER STATES' RULES] (defining the "primary purpose" of clinical health data as treatment of the person the data describe, and treating everything else as a "secondary purpose," seemingly relegating treatment of other patients to a lesser status along with along with research, public health uses, and management of the health care system).

99. *See* 45 C.F.R. §§ 160.202-.203 (2022) (Privacy Rule preemption provisions).

100. *See* Regulation 2016/679, *supra* note 40, art. 6 (requiring consent for processing of personal data in § 1(a) but allowing unconsented processing for various purposes such as legal compliance, "to protect the vital interests of the data subject or of another natural person," for tasks "carried out in the public interest" in § 1(b)-(f) and allowing Member States to specify provisions "to adapt the application of the rules" in some of these circumstances); *see also id.* art. 9 (addressing the processing of "special categories of personal data" which include health data and requiring consent in § 2(a) but allowing Member States to establish different conditions and safeguards for data used in "preventive or occupational medicine, for the assessment of the working capacity of the employee, medical diagnosis, the provision of health or social care or treatment or the management of health or social care systems and services" in § 2(h); in public health in § 2(i); and for public interest purposes including scientific research in § 2(j)); *id.* art. 89 (allowing Member State law to derogate from various rights provided by the GDPR when those "rights are likely to render impossible or seriously impair the achievement" of various public-interest goals including scientific research).

101. *See generally* EU MEMBER STATES' RULES, *supra* note 98 (surveying flows of clinical data under the laws of EU member states).

102. *Id.* at 26 (citing French and Netherlands law as examples).

103. The apparent difference between the Privacy Rule and the GDPR in this regard may be semantic, tracing to the fact that EU citizens generally have universal access to health care services, which blurs the line between using a person's data to treat others *versus* a public health use of the data. Access to treatment is part of Europe's commitment to public health. In contrast, the United States sharply distinguishes the concepts of "public health" and "treatment of others" because it conceives public health activities as benefiting entire populations, whereas treatment activities benefit the fortunate few who have access to health care.

Beyond treatment uses of data, HHS identified other essential information flows in and around the clinical health care context.[104] These support values that most people regard as non-frivolous and important: enabling organ transplants; detecting child abuse; tracking epidemics; validating workers' compensation claims; facilitating dignified burial of the deceased; providing accurate medical evidence in judicial proceedings where justice depends on it; enabling regulatory oversight of health care; detecting and addressing medical misadventures when treatment fails to serve its hoped-for value of improving health; discovering new ways to treat illness; and other social values to which clinical medicine traditionally contributes.[105] The HIPAA Privacy Rule allows unconsented data flows for these high priority uses but subjects them to various alternative privacy protections.[106]

When crafting the norms identified in Table 1, HHS was "aware of the concerns of privacy and consumer advocates" about controlling access to their data, but HHS determined that "[t]he allowable disclosures and corresponding restrictions we recommend reflect a balancing of privacy and other social values."[107] The Privacy Rule was never all about your individual autonomy; it is about the contextual ends and values of clinical health care.[108] It is about making health care work and, as Part II argues, about making health care work more equitably—something that, to date, American health care has struggled to do.

Faden et al. observe that modern control-over-information theory, with its emphasis on autonomy, "never emphasized obligations of patients to contribute" to the general knowledge base by sharing their data.[109] They call for these "traditional presumptions . . . to change."[110] The Privacy Rule changes those presumptions, not by crafting a new set of informational norms, but by reaffirming the older set of norms that had long existed in medical privacy law. The question is which approach—control-over-information theory or the duty-based approach of medical privacy law—is better tailored for the ethical challenges that lie ahead in the age of AI-enabled health care. The next Part explores what those challenges are.

---

104. *See supra* Table 1 (listing various purposes for which the Privacy Rule allows disclosures).

105. *Id.*

106. *See id.* (summarizing which norms apply alternative privacy protections).

107. *See HHS, 1997 Recommendations*, *supra* note 38, § I.I (calling for "limited disclosures of health information without patient consent for specifically identified national priority activities").

108. *See supra* notes 92-97 and accompanying text.

109. *See* Faden et al., *supra* note 90, at S23.

110. *Id.*

II. THE ETHICAL CHALLENGE OF AI-ENABLED HEALTH CARE

This discussion explores the ethical challenges of AI-enabled health care, focusing on AI/ML CDS tools, which are a large and important category of medical AI at the core of the AI-driven transformation of traditional health care delivery.[111] CDS tools offer recommendations to health care professionals on how to diagnose, predict the course of, or treat a patient's disease.[112] They work by comparing data about the particular patient to a source of general medical knowledge and inferring what might be wrong with the patient and which treatments might work best.[113] Some CDS tools are simple and do not incorporate AI/ML: for example, simple CDS tools might rely on published literature, clinical practice guidelines, or information from FDA-approved drug labeling as their source of general medical knowledge.[114] For AI/ML CDS tools, however, the source of general medical knowledge can include insights that an AI algorithm gleaned by processing data reflecting a very large number of other people's health care experiences—for example, which treatments typically worked best for past patients with symptoms and characteristics similar to the patient now being treated?[115]

CDS tools are designed to assist, rather than to take the place of, health care professionals working in clinical care settings such as hospitals, clinics, clinical laboratories, and nursing homes.[116] This

111. *See Clinical Decision Support*, HEALTHIT.GOV (Apr. 10, 2018), https://www.healthit.gov/topic/safety/clinical-decision-support [https://perma.cc/BAC4-ZVM8] (describing a range of CDS tools providing decisional support to health care professionals).

112. *See* 21st Century Cures Act, Pub. L. No. 114-255, § 3060(a), 130 Stat. 1033 (2016) (codified at 21 U.S.C. § 360j(o)) (defining five categories of medical software and delimiting the extent of the FDA's authority to regulate each of them); *see also id.* § 360j(o)(1)(E) (discussing the category of software commonly known as clinical decision support software without using that name, and including as one of its attributes that it functions "for the purpose of 'supporting or providing recommendations to a health care professional about prevention, diagnosis, or treatment of a disease or condition.' "); U.S. FOOD & DRUG ADMIN., CLINICAL DECISION SUPPORT SOFTWARE: GUIDANCE FOR INDUSTRY AND FOOD AND DRUG ADMINISTRATION STAFF (2022), https://www.fda.gov/regulatory-information/search-fda-guidance-documents/clinical-decision-support-software [https://perma.cc/QY8N-B28P] (using the term "Clinical Decision Support Software" to refer to the category of software described at 21 U.S.C. § 360j(o)(1)(E) of the 21st Century Cures Act).

113. *See* Julia Adler-Milstein et al., *Meeting the Moment: Addressing Barriers and Facilitating Clinical Adoption of Artificial Intelligence in Medical Diagnosis* 15 (National Academy of Medicine Discussion Paper, September 29, 2022) (noting that Diagnostic Decision Support (DDS) tools are one type of CDS tool and that "CDS tools combine general medical 'knowledge' with patient-specific information to produce recommended diagnoses. With AI-DDS systems, that knowledge can include inferences generated internally by an AI/ML algorithm").

114. *See id.*; *see also Clinical Decision Support*, *supra* note 111 (providing various examples).

115. *See* Adler-Milstein et al., *supra* note 113, at 15.

116. *See supra* notes 112-14.

distinguishes CDS tools from consumer-facing health apps and at-home monitoring tools, for which the user is often a medically untrained layperson without skills to challenge the software's recommendations.[117] For CDS tools, the user is a trained professional who, at least in theory, could query, challenge, or even reject the software's recommendations.[118] CDS tools also are distinct from AI algorithms imbedded within hardware medical devices, such as software processing dental X-ray images to highlight suspected cavities, where the software is optimizing device performance rather than overtly offering recommendations to a health care professional.[119]

AI/ML CDS tools are already in use at many health care facilities, recommending diagnoses and treatments and affecting patient care, often without patients knowing that AI tools are in use.[120] A recent report by the U.S. Government Accountability Office and National Academies of Science, Engineering, and Medicine identified six ethical and practical challenges with AI/ML CDS tools.[121] The three ethical challenges are to protect privacy, to access high-quality data for use in training the AI, and to reduce potential biases in the training data, which can cause the tools to perform poorly for patients who are unlike the people reflected in the training data.[122] Bias is thus a source of

117. *See generally* David A. Simon et al., *Skating the Line Between General Wellness Products and Regulated Devices: Strategies and Implications*, J.L. & BIOSCIENCES, July-Dec. 2022 (discussing differences between consumer-facing medical products and FDA-regulated clinical decision support tools).

118. *See, e.g.*, 21st Century Cures Act, Pub. L. No. 114-255, § 3060(a), 130 Stat. 1033, 1130-33 (2016) (codified as amended at 21 U.S.C. § 360j(o)(1)(E)) (establishing the FDA's jurisdiction to regulate some categories of clinical decision support software and focusing the FDA's oversight on software that is not intended to enable the "health care professional to independently review the basis for such recommendations that such software presents" so that there is an intent that the "health care professional rely primarily on any of such recommendations to make a clinical diagnosis or treatment decision regarding an individual patient").

119. *See* 21 U.S.C. § 360j(o)(1)(E) (confirming FDA's jurisdiction to regulate software whose "function is intended to acquire, process, or analyze a medical image or a signal from an in vitro diagnostic device or a pattern or signal from a signal acquisition system"); *see also* Bradley Merrill Thompson, *Learning from Experience: FDA's Treatment of Machine Learning*, MOBIHEALTHNEWS (Aug. 23, 2017), http://www.mobihealthnews.com/content/learning-experience-fda%E2%80%99s-treatment-machine-learning [https://perma.cc/2U3J-XTTZ] (discussing AI algorithms imbedded in diagnostic imaging devices).

120. Rebecca Robbins & Erin Brodwin, *An Invisible Hand: Patients Aren't Being Told About the AI Systems Advising Their Care*, STAT NEWS (July 15, 2020), https://www.statnews.com/2020/07/15/artificial-intelligence-patient-consent-hospitals/ [https://perma.cc/52GC-B3Y2].

121. *See* U.S. GOV'T ACCOUNTABILITY OFF., GAO-21-7SP, ARTIFICIAL INTELLIGENCE IN HEALTH CARE: BENEFITS AND CHALLENGES OF TECHNOLOGIES TO AUGMENT PATIENT CARE 21 (2020).

122. *Id.* (listing these three ethical challenges and also noting various operational challenges including "difficulties in scaling" software systems to serve large populations in diverse health care settings, "limited transparency of AI tools," and "uncertainty about liability").

future health care inequities, if CDS tools offer inaccurate recommendations for patients who were not well represented in the original training data used when developing the tools.

Section II.A recognizes health equity as the central ethical challenge for AI/ML CDS tools and examines the role privacy policy plays in abetting inequity and social injustice in twenty-first-century health care. It explores the potential for control-over-information privacy policies to contribute to harmful biases in AI/ML medical tools. Section II.B explores how this happened, critiquing aspects of twentieth-century bioethics that tend to obscure the fact of human diversity when, in reality, solving the problem of bias in AI software requires that we acknowledge and embrace the richness of human diversity.

A. The Critique That Consent Norms Contribute to Health Care Inequity

A growing body of literature finds that AI/ML CDS tools perform more reliably for white male patients than for women, transgender patients, people of color, and persons experiencing economic disadvantages.[123] A celebrated achievement of twenty-first-century medicine—AI/ML CDS software—threatens to become a new source of invidious discrimination in health care.

The problem of bias is widely attributed—and with a large measure of truth—to structural and systemic inequities in U.S. health care. That may not be the entire story, though, and other contributing factors also need to be explored. In particular, the ethical and legal norms with which a society governs the acquisition of training data for AI/ML CDS software influence how reliable—and how equitably reliable—its recommendations will be. This discussion critiques a widely shared bioethical norm: that secondary uses of clinical health data should require de-identification or informed consent. Could biomedicine's widespread adherence to this DOGC norm be contributing to the problem of non-inclusive training data, fueling health inequities and social injustice?

Bias in AI/ML training data is a multi-headed Hydra with many different causes.[124] Only some of these causes can be traced to the ethical rules governing acquisition of data for scientific use.[125] Other

123. *See infra* notes 126-31 and accompanying text.

124. *See Hydra: Greek Mythology*, BRITANNICA, https://www.britannica.com/topic/Hydra-Greek-mythology [https://perma.cc/TTF4-HH2F] (last visited Sept. 23, 2023) (defining "Hydra" as a "gigantic water-snake-like monster with nine heads (the number varies), one of which was immortal," as described in Hesiod's *Theogony*).

125. *See* U.S. FOOD & DRUG ADMIN., USE OF REAL-WORLD EVIDENCE TO SUPPORT REGULATORY DECISION-MAKING FOR MEDICAL DEVICES 21 (Aug. 31, 2017), https://www.fda.gov/media/99447/download [https://perma.cc/M3VP-X3US] (defining "Bias"

sources of bias reflect broad policy failures of the health care system and the biomedical research environments where medical software operates. For example, people denied access to health care services leave no trails of clinical health data that can be used to train future AI/ML CDS tools. This problem reflects a systemic failure of financing and compassion. It has nothing to do with the ethical rules for acquiring data for scientific use. Neither a notice-and-consent privacy scheme, nor any alternative regime that might be imagined, can mobilize data that simply do not exist. AI/ML CDS tools are only as equitable as the inequitable health care system from which they draw their real-world training data.

Law considers discrimination "invidious" when people are treated in damaging ways because of race, gender, or class without a rational reason to do so (for example, there could be a good reason to disadvantage a historically privileged group to correct past injustices).[126] AI/ML CDS tools show great promise for improving health care, but recent empirical studies reveal their ominous potential to become instruments of invidious health care discrimination.[127] Training data for AI/ML CDS tools tend to overrepresent men of European ancestry while underrepresenting members of other racial and ethnic groups and women.[128] Empirical studies of how CDS tools perform for transgender patients do not even appear to exist, but it is known that these patients face special health risks (such as elevated incidence of aortic aneurysms in transgender women).[129] Those risks can

---

as "any systematic error in the design, conduct, analysis, interpretation, publication, or review of a study and its data that results in a mistaken estimate of a treatment's effect on disease" and noting that such errors can result "from flaws in the method of selecting study participants, in the procedures for gathering data, and in the decision of how and whether to publish the results").

126. *See Invidious Discrimination Law and Legal Definition*, USLEGAL, https://definitions.uslegal.com/i/invidious-discrimination/ [https://perma.cc/QN43-LFKA] (last visited Sept. 23, 2023); *Invidious Discrimination*, LEGAL INFO. INST., https://www.law.cornell.edu/wex/invidious_discrimination#:~:text=Treating%20a%20class%20of%20persons,criminal%20law [https://perma.cc/Y4ML-SK9F] (last visited Sept. 23, 2023).

127. *See, e.g.*, U.S. GOV'T ACCOUNTABILITY OFF., *supra* note 121, at 24.

128. For concerns with racial biases in AI/ML data and algorithms, see, e.g., Adewole S. Adamson & Avery Smith, *Machine Learning and Health Care Disparities in Dermatology*, 154 JAMA DERMATOLOGY 1247, 1247 (2018); Ruha Benjamin, *Assessing Risk, Automating Racism*, 366 SCIENCE 421, 421 (2019); Ziad Obermeyer et al., *Dissecting Racial Bias in an Algorithm Used to Manage the Health of Populations*, 366 SCIENCE 447, 450 (2019); Alice B. Popejoy et al., *The Clinical Imperative for Inclusivity: Race, Ethnicity, and Ancestry (REA) in Genomics*, 39 HUM. MUTATION 1713, 1714, 1717-18 (2018). For concerns with gender-based disparities, see, e.g., CAROLINE CRIADO PEREZ, INVISIBLE WOMEN: DATA BIAS IN A WORLD DESIGNED FOR MEN 90-91, 93 (2019).

129. *See* Zaria Gorvett, *Why Transgender People Are Ignored by Modern Medicine*, BBC (Aug. 16, 2020), https://www.bbc.com/future/article/20200814-why-our-medical-systems-are-ignoring-transgender-people [https://perma.cc/5EQJ-9D9T] (noting elevated risk of aortic aneurism in transgender females); *see also* National Cancer Institute, *Keynotes: Dr. Deven McGraw and Dr. Karl Surkan on Personal Control of Genomic Data for Research*, YOUTUBE (Feb. 20, 2020), https://youtu.be/-sSJK_LEW6U [https://perma.cc/R7ZT-7Y7A] (discussing health experiences of BRCA-positive transgender male patients at 24:53).

be obscured if AI/ML training data sets force-fit patients into gender-binary categories without further nuancing to highlight special medical needs within those categories.[130] AI/ML tools that perform well in high-resource, well-staffed academic medical centers often underperform at lower-resourced community hospitals where much of the American population receives care.[131]

Accepting that much of this problem is systemic and hard for individuals to change, there is an aspect of it that lawyers and bioethicists can influence: the ethical norms for acquiring health data for secondary uses, such as for research or for inclusion in AI/ML training data. Ethical norms can themselves become a source of invidious discrimination in AI/ML CDS software, if the norms have disparate impacts and produce training data sets that fail to reflect the full diversity of patients who ultimately will be treated using the software. For excluded or underrepresented groups, the software is under-informed and might produce unreliable—even dangerous—recommendations. This discussion seeks to distinguish the various contributors to bias with a view to identifying those that might be responsive to shifts in the informational norms surrounding acquisition of training data.

Systemic bias in health care data. The most obvious contributor to systemic bias, already noted, is that patients denied access to health care generate little or no real-world clinical data to include in AI/ML training data. AI/ML algorithms cannot learn from data that do not exist.

---

130. *See, e.g.*, Curtis S. Tenney et al., *A Crisis of Erasure: Transgender and Gender-Nonconforming Populations Navigating Breast Cancer Health Information*, 5 INT'L J. INFO. DIVERSITY & INCLUSION 132, 132, 134 (2021) (noting that many health information norms and research practices fail to recognize gender-nonconforming categories); Christine Labuski & Colton Keo-Meier, *The (Mis)Measure of Trans*, 2 TSQ: TRANSGENDER STUD. Q. 13, 13 (2015) (noting "transgender's instability as a research variable" and calling for "more precise methodological orientations in trans research, particularly regarding gender and sexual orientation"); *see also* Sari L. Reisner et al., *"Counting" Transgender and Gender-Nonconforming Adults in Health Research: Recommendations from the Gender Identity in US Surveillance Group*, 2 TSQ: TRANSGENDER STUD. Q. 34, 34-56 (2015) (providing recommendations for including gender minority adults in health research); T. Benjamin Singer, *The Profusion of Things: The "Transgender Matrix" and Demographic Imaginaries in US Public Health*, 2 TSQ: TRANSGENDER STUD. Q. 58, 58 (2015) (cautioning that "[d]emographic categories are double-edged swords in that they are necessary for the redirection of resources toward socially marginalized people; at the same time, they often constitute the conditions of containment of these same people").

131. *See* W. Nicholson Price II, *Medical AI and Contextual Bias*, 33 HARV. J.L. & TECH. 88, 91, 95, 113 (2019) (discussing narrow validity of AI systems developed in resource-rich contexts when implemented in lower-resource settings). *See generally* ERIC TOPOL, DEEP MEDICINE: HOW ARTIFICIAL INTELLIGENCE CAN MAKE HEALTHCARE HUMAN AGAIN (2019) (discussing lack of reproducibility, narrow validity, nontransparency of data, and overblown claims about the benefits of medical software); Matthew Zook et al., *Ten Simple Rules for Responsible Big Data Research*, PLOS COMPUTATIONAL BIOLOGY, Mar. 2017 (identifying similar limits); Danah Boyd & Kate Crawford, *Critical Questions for Big Data: Provocations for a Cultural, Technological, and Scholarly Phenomenon*, 15 INFO. COMMC'N & SOC'Y 662 (2012) (same).

Yet even when training data are inclusive and represent all members of society, algorithmic biases can mischaracterize what the data are saying.[132] Using hospital admission as proof of a heart attack erases the heart attacks of women told to go home and get some rest because doctors dismissed their symptoms as stress-related.[133] The algorithm will fail to glean useful insights from their experiences and provide less accurate results for women than for men. "Men's . . . health needs largely define insurance coverage,"[134] and "[u]ntil recently, medical research generally calibrated 'normal' on a trim white male."[135] The systemic problem is biased doctoring, accurately reflected in the data. Nevertheless, a good algorithm should be able to detect biased doctoring when that is what the data reveals.

Another example of algorithmic bias is the finding by Obermeyer et al. of racial bias in a commercial risk-prediction tool used by many large health systems and payers in the United States.[136] The algorithm viewed the amount of money spent on a person's health care as a proxy for how sick the person was.[137] It erroneously inferred that Black patients were healthier than they actually were, because their health care expenditures were low.[138] In reality, they were just not receiving adequate care for the illnesses they did have.[139] Algorithmic bias is not systemic in the sense of being beyond human control; it is well within human control to write a good algorithm that interprets inequity as inequity instead of interpreting it as "women and Black people are super-healthy." Bioethicists and patient safety regulators may not be the ones to write better algorithms themselves, but their voices can help bring about a system in which software developers are required to do so.

Another kind of bias, selection bias, occurs when the group of people included in (selected for) a study are not representative of the entire

---

132. Algorithmic bias, as its name suggests, is introduced not by defects in the data being analyzed, but by the algorithm that processes the data—for example, due to errors in the calculations the algorithm performs or faulty assumptions incorporated as the algorithm was designed.

133. *See generally* ELINOR CLEGHORN, UNWELL WOMEN: MISDIAGNOSIS AND MYTH IN A MAN-MADE WORLD (2021) (discussing male normativity of the U.S. health care system).

134. CATHARINE A. MACKINNON, TOWARD A FEMINIST THEORY OF THE STATE 224 (1989).

135. *See* Janice P. Nimura, *Why 'Unwell Women' Have Gone Misdiagnosed for Centuries*, N.Y. TIMES (June 8, 2021), https://www.nytimes.com/2021/06/08/books/review/unwell-women-elinor-cleghorn.html [https://perma.cc/7XL4-CG4F] (reviewing CLEGHORN, *supra* note 133).

136. *See* Obermeyer et al., *supra* note 128, at 447-49; *see also* Salman Ahmed et al., *Examining the Potential Impact of Race Multiplier Utilization in Estimated Glomerular Filtration Rate Calculation on African-American Care Outcomes*, 36 J. GEN. INTERNAL MED. 464, 466 (2021) (examining how a numerical adjustment used in calculating severity of kidney disease understated the severity of disease in African Americans).

137. *See* Obermeyer et al., *supra* note 128, at 447-49.

138. *Id.*

139. *Id.*

population that ultimately will rely on results from that study.[140] "Bias in data used to develop AI tools can reduce their safety and effectiveness for patients who differ—whether genetically or in socioeconomic status, general health status, or other characteristics—from the population whose data were used to develop the tool."[141] Consent bias is one form of selection bias. It occurs when people who consent (self-select) to include their data in a study differ from the population at large.[142]

Consent is the product of individual choices, yet consent bias is partly systemic. For example, people cannot consent to be in AI/ML training data if they are never asked to do so. Academic medical centers, as early adopters of electronic health record systems, have been leaders in research to develop AI/ML CDS tools.[143] Patients privileged to receive care at such sites stand a better chance of being represented—of having people like them[144]—in training data.[145] In contrast, "less than 1% of federal [research and development] expenditures went to historically Black colleges and universities (HBCUs) in 2019."[146] Just 7.4% and 6.6% of the National Science Foundation (NSF) and the National Institutes of Health (NIH) grant awards flow to Black and Latino or Hispanic innovators—"far below those groups' share of the population."[147]

---

140. SABRINA NOUR & GILLES PLOURDE, *Pharmacoepidemiology in the Prevention of Adverse Drug Reactions*, *in* PHARMACOEPIDEMIOLOGY AND PHARMACOVIGILANCE: SYNERGISTIC TOOLS TO BETTER INVESTIGATE DRUG SAFETY 25 (2019) ("Selection bias occurs when the study population is not representative of the target population so that the measure of risks/benefits does not accurately represent the target population to which conclusions are being extended.").

141. *See* U.S. GOV'T ACCOUNTABILITY OFF., *supra* note 121, at 24.

142. *See* Evans, *Data Ownership*, *supra* note 41, at 95-96 (citing and discussing various studies of consent bias); *see also* Brian Buckley et al., *Selection Bias Resulting from the Requirement for Prior Consent in Observational Research: A Community Cohort of People with Ischaemic Heart Disease*, 93 HEART 1116 (2007) (defining "consent bias" as "a term coined to describe the selection bias resulting from the loss of non-consenters to any cohort").

143. *See* John D Halamka et al., *Early Experiences with Personal Health Records*, 15 J. AM. MED. INFO. ASS'N. 1, 1-2 (2008) (describing examples of electronic records systems installed at several high-resource academic medical centers); John D. Halamka, *Early Experiences with Big Data at an Academic Medical Center*, 33 HEALTH AFFAIRS 1132, 1132-38 (2014) (describing the resources required to transform raw data into a useful knowledge resource). *But see* Vindell Washington et al., *The HITECH Era and the Path Forward*, 377 NEW ENG. J. MED. 904, 905 (2017) (noting that only about 10% of hospital systems used EHRs in 2008).

144. *See* LAWRENCE LESSIG, CODE AND OTHER LAWS OF CYBERSPACE 152 (1999) ("[N]o one spends money collecting these data to actually learn anything about you. They want to learn about people like you." (emphasis omitted)).

145. *See* Price, *supra* note 131, at 67, 79-80, 91-93.

146. Mark Muro et al., *Congress Needs to Prioritize Inclusion in Our Slumping Innovation System*, BROOKINGS INST. (Aug. 11, 2021), https://www.brookings.edu/blog/the-avenue/2021/08/11/congress-needs-to-prioritize-inclusion-in-our-slumping-innovation-system/ [https://perma.cc/G77D-97LQ].

147. *See id.*; *cf.* Francis S. Collins et al., *Affirming NIH's Commitment to Addressing Structural Racism in the Biomedical Research Enterprise*, 184 CELL 3075, 3075 (2021) (discussing NIH's commitment to promote greater equity in the innovation system).

There are systemic, racial, and gender barriers to entry and professional success at all levels of the U.S. innovation system, reflecting, for example, community norms and stereotyping of roles, disparities in completing STEMM[148] education, and in obtaining funds for research and commercialization of discoveries.[149] Even if a woman of color manages to succeed in research, "all else being equal, patent applications with women as lead inventors are rejected more often than those with men as lead inventors" by the U.S. Patent and Trademark Office.[150] Assuming diverse populations are even asked to consent to have their data used to train AI/ML CDS software, they might hesitate to say "yes" if the researcher asking for their data does not look like them. Why arm your oppressors with detailed personal information about you?

Refusing to consent is an individual choice and thus hard to characterize as systemic. Yet these individual choices are responsive to systemic factors reflecting a nation's broad policies on how and where to allocate research funding and how to bestow the privilege (by wealth, insurance status, location, color, and gender) of receiving treatment at leading research centers. In this sense, then, consent bias is partly systemic. There are no easy bioethical solutions to structural and systemic inequities, which usually require fiscal resources rather than ethical nostrums to cure.

The disparate impact of informed consent norms. It does not discount the extent of systemic bias to recognize that some sources of bias are non-systemic and could be reduced without overhauling the entire, inequitable health care system. Omitting people from AI/ML training data (or granting them a right to stay out) ostensibly respects their autonomy and protects them from privacy harms, but the resulting AI/ML tool can subject non-consenters to real medical harm if they

148. STEMM stands for Science, Technology, Engineering, Mathematics, and Medicine.

149. *See* Lisa D. Cook & Janet Gerson, *The Implications of U.S. Gender and Racial Disparities in Income and Wealth Inequality at Each Stage of the Innovation Process*, WASH. CTR. FOR EQUITABLE GROWTH (July 24, 2019), https://equitablegrowth.org/the-implications-of-u-s-gender-and-racial-disparities-in-income-and-wealth-inequality-at-each-stage-of-the-innovation-process/ [https://perma.cc/2GDC-5L7R] (reviewing literature exploring systemic factors causing racial and gender gaps in participation in STEMM fields); *see also* Maria Klawe, *How Can We Encourage More Women to Study Computer Science?*, CRA-WP (2015), https://cra.org/cra-wp/629/ [https://perma.cc/B6NV-TM4N] (noting a 40% decline over the past twenty years in women's completion of degrees in computer science). *But see* Marcia McNutt & Laura Castillo-Page, *Promoting Diversity and Inclusion in STEMM Starts at the Top*, 27 NATURE MED. 1864, 1864-65 (2021) (expressing commitment to increase diverse participation in STEMM fields).

150. *See* Cook & Gerson, *supra* note 149, at 8 (citing Kyle Jensen et al., *Gender Differences in Obtaining and Maintaining Patent Rights*, 36 NATURE BIOTECHNOLOGY 307, 307 (2018)).

seek care from providers who rely on it.[151] People have the potential to reduce these harms by contributing their data for use in training medical AI tools. Why don't they?

Studies show that consenters differ medically from the population at large.[152] For example, people with stigmatizing health conditions might hesitate to share their data for research lest a data breach "out" embarrassing medical facts about them.[153] There is a tradeoff between producing high-quality, generalizable scientific results *versus* showing respect for autonomy by seeking consent. Consent bias exists because of that tradeoff.

Some bioethicists question whether consent bias has a material impact on scientific results.[154] Others argue that even if consent bias is material, it is ethically irrelevant, as in Franklin Miller's statement that even if a data use has high social value, if consent is logistically difficult or impossible to obtain, and if requiring consent may undercut the scientific validity of results, these facts "do not in themselves constitute valid ethical reasons for waiving a requirement of informed consent."[155] Would it be a "valid ethical reason" to question consent-based data acquisition schemes if it turns out they are promoting racial, gender, and socioeconomic inequities in AI-enabled health care? Alternatively, is the ethical norm of informed consent so sacred that racial, gender, and socioeconomic disparities in health care are a small price to pay for the benefits AI/ML tools confer on white males?

---

151. *See supra* notes 120-23 (discussing harms from biased AI/ML medical tools).

152. *See generally* Buckley et al., *supra* note 142; IOM, PRIVACY REPORT, *supra* note 5, at 209-14 (surveying studies of consent and selection bias); Khaled El Emam et al., *A Globally Optimal k-Anonymity Method for the De-Identification of Health Data*, 16 J. AM. MED. INFO. ASS'N 670, 670 (2009); Steven J. Jacobsen et al., *Potential Effect of Authorization Bias on Medical Record Research*, 74 MAYO CLINIC PROC. 330 (1999); Jack V. Tu et al., *Impracticability of Informed Consent in the Registry of the Canadian Stroke Network*, 350 NEW ENG. J. MED. 1414 (2004); Steven H. Woolf et al., *Selection Bias from Requiring Patients to Give Consent to Examine Data for Health Services Research*, 9 ARCHIVES FAM. MED. 1111 (2000).

153. *See* Evans, *Data Ownership*, *supra* note 41, at 95 (noting that the underlying reasons for consent bias are not fully understood). *But see* Buckley et al., *supra* note 142, at 1116-17 (citing studies that found that people who decline to consent were more likely to be female and younger than 60 years old; that persons with "sensitive" diagnoses like reproductive disorders, mental disorders, or infectious diseases were less willing to consent; that willingness to consent varies by age group; and that people who consent are less likely to live in economically deprived areas than non-consenters are).

154. *See, e.g.*, Mark A. Rothstein & Abigail B. Shoben, *Does Consent Bias Research?*, 13 AM J. BIOETHICS 27, 27 (2013) (arguing "claims about the degree of consent bias are overstated").

155. Franklin G. Miller, *Research on Medical Records Without Informed Consent*, 36 J.L. MED. & ETHICS 560, 560 (2008) (describing but not necessarily endorsing this view, which Miller describes as: "The facts that historically much valuable population-based observational research has been conducted without informed consent, that obtaining consent would often make such research impossible to conduct, and that selection biases associated with soliciting consent may compromise its scientific validity, do not in themselves constitute valid ethical reasons for waiving a requirement of informed consent").

Professor Tauber remarks that foundational American bioethical works after 1970 stressed individual autonomy and consent norms without explaining how the principle of autonomy "compete[s] with other moral tenets" such as the principles of beneficence and justice.[156] Professor Snead adds that bioethics also "elevates the principles of autonomy and self-determination" above equality.[157] It may be time to ask whether twentieth-century bioethics—in which "autonomy usually trumps other contenders" in the hierarchy of moral values—is still appropriate in a twenty-first-century, AI-enabled health care system serving an ever more diverse patient population.[158]

There is mounting evidence that consenters differ not just medically but also demographically from those who do not consent.[159] Conditioning inclusion in AI/ML training data on individual consent invites disparate underrepresentation of population subgroups that feel reluctant to consent because of unpleasant past research experiences.[160] "[T]he research community is on notice . . . that there are important differences in preferences by race and ethnicity" in terms of whether people view research as important and how comfortable they feel having their data used in research and under what conditions.[161] Empirical work by Jagsi et al. found "racial and ethnic minorities may be particularly concerned about consent to any participation in research" and pointed out that this result is consistent with findings from other studies.[162]

---

156. *See* TAUBER, *supra* note 19, at 16; *see also* Alfred I. Tauber, *Sick Autonomy*, 46 PERSPECTIVES BIOLOGY & MED. 484, 488 (2003) (noting that "although autonomy, beneficence, justice, and non-malfeasance each claim consideration, autonomy usually trumps other contenders" and citing Beauchamp and Childress, *supra* note 16, as having made a similar observation); *id.* (citing Paul Root-Wolpe, *supra* note 16, for the statement that autonomy "indisputably . . . become the central and most powerful principle in ethical decision making in American medicine"); SNEAD, *supra* note 16 (expressing similar views).

157. *See* SNEAD, *supra* note 16, at 71.

158. *See* Tauber, *supra* note 156, at 488.

159. *See* Buckley et al., *supra* note 142, at 1117 (citing a study that, as early as 2003, identified an association suggesting that people living in "deprived areas" were less likely to consent to research than more privileged people were).

160. *See, e.g.*, Spector-Bagdady, *supra* note 8, at 1, 2-3, 6-8 (recounting past research abuses including the Tuskegee Syphilis Study that left Black males untreated for syphilis to observe how the disease progresses when untreated, research that purposefully infected vulnerable Guatemalans with sexually transmitted diseases, and the case of Henrietta Lacks); Nanibaa' A. Garrison, *Genomic Justice for Native Americans: Impact of the Havasupai Case on Genetic Research*, 38 SCI. TECH.& HUM. VALUES 201, 201-04 (2013) (describing a case involving unconsented secondary use of blood specimens collected from the Havasupai Tribe); H.K. Beecher, *Ethics and Clinical Research*, 274 NEW ENGL. J. MED. 1354, 1368-70 (1966) (detailing a variety of research abuses against elderly, vulnerable, and cognitively impaired persons).

161. *See* Spector-Bagdady, *supra* note 8, at 2-3 (citing Reshma Jagsi et al., *Perspectives of Patients with Cancer on the Ethics of Rapid-Learning Health Systems*, 35 J. CLIN. ONCOL. 2315 (2017)).

162. Jagsi et al., *supra* note 161, at 2321 (citing Raymond Gene De Vries et al., *Understanding the Public's Reservations about Broad Consent and Study-by-Study Consent for Donations to a Biobank: Results of a National Survey*, PLOS ONE, July 2016).

Moreover, de-identification requirements can diminish data utility. For example, stripping away identifiers, such as the patient's zip code, might make it hard to detect non-inclusiveness and bias in AI/ML training data.[163] De-identification also contributes to the erasure of transgender patients in AI/ML training data: sorting data into binary "male" and "female" buckets affords anonymity while obscuring details about individual life trajectories (e.g., "assigned male at birth, female since age 23")[164] that have medical significance (e.g., "possible heightened risk for aortic aneurysm"). Inclusive AI/ML data are good, but inclusive, nuanced AI/ML data are even better. De-identification can diminish nuancing.[165]

One solution ethicists propose is to conduct better outreach to underrepresented groups and to promote better practices for obtaining their consent as ways to improve data inclusivity while preserving valuable nuancing. Programs like the NIH-funded All of Us and Bridge2AI research programs already are working to develop more inclusive data resources for future biomedical and AI research.[166]

But what if, at the end of all persuasion, underrepresented populations still choose not to consent to have their data included in AI/ML training data? At that point, is the ethically appropriate response to say, in an echo of Miller's statement,[167] "We respect your autonomy. You had your chance to participate and declined, and we respect that" and then forge ahead designing medical software tools that might prove deadly to the non-consenters? Is there ever a point where it is ethically appropriate to question the unwavering fealty to informed consent itself? Where is the coherence of bioethical norms that exalt individual autonomy to the point where it trumps concerns about racism, gender bias, transgender exclusion, and poverty, which diminish the autonomy of millions of individuals? The foregrounding of autonomy/consent in modern bioethics and control-over-information

163. *See* U.S. GOV'T ACCOUNTABILITY OFF., *supra* note 121, at 24.

164. *See supra* notes 121-22 and accompanying text.

165. *See, e.g.*, Off. Director, Nat'l Inst. Health, Request for Information on Proposed Updates and Long-Term Considerations for the NIH Genomic Data Sharing Policy 2 (NOT-OD-22-029, Nov. 30, 2021) (noting, in discussing de-identification requirements of the current N.I.H. Genomic Data Sharing Policy, that "[c]ertain data elements considered potentially identifiable, such as date ranges shorter than a year, may have scientific utility, especially when studying disease progression (e.g., with COVID-19) or higher resolution location data than the regulatory standard (e.g., full ZIP codes or mobile location data), which may be valuable for studying the social determinants of health or environmental risk").

166. *See, e.g.*, *About*, NAT'L INSTITUTES HEALTH, https://allofus.nih.gov/about [https://perma.cc/X455-38X6] (last visited Sept. 23, 2023) (describing the NIH All of Us Research Program); *Bridge to Artificial Intelligence*, NAT'L INSTITUTES HEALTH, https://commonfund.nih.gov/bridge2ai [https://perma.cc/8HWK-8RQN] (last visited Sept. 23, 2023) (describing the NIH Bridge2AI Program).

167. *See* Miller, *supra* note 155, at 560.

theory is especially problematic in the era of AI-enabled health care where consent bias poses medical safety concerns for population subgroups that are underrepresented in AI training data.

B. The Critique That Consent Norms Are Rooted in a Denial of Diversity

This Section argues that the nub of the problem with twentieth-century bioethics and its control-over-information privacy theory is its failure to acknowledge human diversity. This failure takes on greater importance as the twenty-first-century health system treats an ever more diverse patient population.

Critical legal theorists have challenged traditional "legal liberalism" for its "assumption that all persons share certain 'samenesses,' such as rationality or autonomy."[168] Twentieth-century bioethics incorporated a similar assumption: specifically, that all persons share sameness as bearers of rationality, autonomy, and, it might be added, representativeness of the population at large.[169] This is seen, for example, in one of its central dogmas: that people who consent to biomedical research are "altruistic."[170]

This dogma calls for people who consent to research to receive clear notice that they, personally, are unlikely to benefit from their research participation, and the benefits, if any, will probably flow to others.[171] When consenters falsely expect to gain a personal benefit from research, this even has a name: "therapeutic misconception," in which consenters wrongly believe research activities will cure them.[172] This altruism narrative portrays consenters as bestowing positive externalities on others. As in Miller's statement, consenters have no duty to rescue others: they are free of any ethical or moral imperative to do

168. *See* Cynthia V. Ward, *On Difference and Equality*, 3 LEGAL THEORY 65, 65 nn.1-2 (1997) (citing literature since 1989).

169. *See supra* notes 154-55 and accompanying text (discussing the tendency in twentieth-century bioethics to treat bias dismissively); *infra* notes 194-95 and accompanying text (discussing twentieth-century medical science's presumption that small clinical trials, populated heavily by white men, were representative of the entire population).

170. *See, e.g.*, Lynn A. Jansen, *The Ethics of Altruism in Clinical Research*, 39 HASTINGS CTR. REP. 26, 26-28 (2009).

171. *See id.* at 26 ("Altruistic motives can explain why rational people with full understanding would agree to participate in trials that offer little or no direct therapeutic benefit and expose them to significant risks of harm."); *id.* at 27 (noting that this is relevant to assessing valid informed consent); *see also* THE PRIV. PROT. STUDY COMM'N, *supra* note 21, at 567 (noting that "research and statistical activities generally do not lead to an immediate or direct benefit for the individual subject as such. The researcher asks for the individual's participation or for information about him, but society as a whole, rather than the individual, is the ultimate beneficiary").

172. *See* Paul S. Appelbaum et al., *False Hopes and Best Data: Consent to Research and the Therapeutic Misconception*, 17 HASTINGS CTR. REP. 20, 20 (1987).

so.[173] Hence, they are altruistic. If letting your data be used in research is framed as altruistic, then consent norms bestow a right against forced altruism.

For this altruism narrative to work, however, twentieth-century bioethics implicitly presumed consenters are the same as (representative of) non-consenters: one consenter is as informative as any other, for purposes of scientific study.[174] Unless this were true, a consenter's participation would not benefit others and would not be altruistic. The altruism narrative fails if there is human diversity.

Likewise, twentieth-century bioethics presumed all people share sameness as bearers of individual autonomy. This presumption has the unintended effect of oppressing subpopulations for whom the exercise of autonomy (e.g., consenting to have their data in AI/ML training data sets) poses group-specific risks (e.g., women's fear that data breaches might leak their gynecological records to prosecutors in the world after *Dobbs v. Jackson Women's Health Organization*,[175] or the fear among undocumented residents that leaking data to Immigrations and Customs Enforcement might lead to deportation[176]). Conditioning data acquisition on notice and consent invites consent bias, which fuels future health disparities as people are left out (because they choose to be left out). Opt-out consent helps somewhat but does not fully eliminate consent bias.[177] Groups facing unusually dire risks from inclusion in a database might be motivated to expend the extra effort it takes to opt out.

The autonomy-based consent norms of late twentieth-century bioethics incorporate "a highly specific model of personhood that was constructed . . . for a white male elite"[178] who, as it happens, were the only people who had much autonomy back when Locke and Kant extolled the virtues of autonomy in terms that profoundly influenced

173. *See* Jansen, *supra* note 170, at 27 (characterizing participation in clinical research as "morally optional—not a moral duty" and noting this is "the common view" with "a long history"); *see also* Miller, *supra* note 155, at 560.

174. *See supra* notes 154-55 and accompanying text (dismissing the materiality and ethical importance of consent bias, views that seem to rest on an implicit presumption of human sameness).

175. *See* Dobbs v. Jackson Women's Health Org., 142 S. Ct. 2228 (2022) (finding no constitutional right to abortion and opening the door for some states to enact laws imposing various civil and criminal sanctions on persons providing, assisting, or receiving abortions).

176. *See Immigration and Customs Enforcement*, DEP'T HOMELAND SEC. (Feb. 26, 2023), https://www.dhs.gov/topics/immigration-and-customs-enforcement [https://perma.cc/G3FJ-GF8C] (describing the U.S. federal agency that arouses fear of deportation among undocumented residents).

177. *See* Evans, *Data Ownership*, *supra* note 41, at 96 (discussing opt-out consent, which is a scheme where people's data are included by default unless they take affirmative steps to remove their data, and contrasting it with opt-in consent, where data are not included without a person's affirmative consents).

178. *See* Ward, *supra* note 168, at 65.

the development of American bioethics.[179] Both philosophers' notion of "autonomy" coexisted with the subservience of women and the colonization or enslavement of entire populations in their day, apparently with approval and even the cooperation of Locke and Kant: "Locke had investments in the African slave trade, and wrote, or at least had a major hand in writing, the Carolina Constitution, which enshrines hereditary African slavery."[180] Kant argued that the "merits that are proper to her sex" for a woman consist of pleasing men.[181] While it seems somewhat baffling today, these philosophers became foundational thinkers of twentieth-century American bioethics. Both appear to have conceived autonomy as the autonomy of cis-gendered European males, which suggests a need to treat with caution those aspects of American bioethics that rest heavily on their work. Rigid norms requiring consent for the scientific use of data are one such aspect.

From its inception in the 1960s and 1970s, the field of bioethics has been—not entirely, but overwhelmingly—a domain of white scholars, and the field allegedly propagates white normativity.[182] Consent-based data acquisition protects the interests of a privileged elite that enjoys the level of autonomy that bioethics presumes, while ignoring real differences that can lead to other people's data being left out of data sets that, today, increasingly affect the quality of health care. The resulting health disparities promote social injustice. Recent empirical studies document invidious bias in AI/ML CDS tools developed to date,

179. *See generally* Tauber, *supra* note 156 (discussing how the works of Locke and Kant influenced bioethicists' embrace of an atomistic conception of autonomy after 1970); SNEAD, *supra* note 16, at 71-72 (tracing the influence of Locke and Kant as well as other thinkers on the field of bioethics).

180. Charles W. Milles, *Locke on Slavery*, *in* THE LOCKEAN MIND 487 (Jessica Gordon-Roth & Shelley Weinberg eds., 2021) (adding that " 'there is a clear contradiction' between his theory and his practice," which various commentators over the years have tried to resolve).

181. *See* Mari Mikkola, *Kant on Moral Agency and Women's Nature*, 16 KANTIAN REV. 89, 89 (2011) (noting that Kant "describes women as coquettes and writes that they have 'a beautiful understanding' due to which '[l]aborious learning or painful pondering, even if a woman should greatly succeed in it, destroy the merits that are proper to her sex' ") (internal citations and emphasis omitted); *see also* Sally Sedgwick, *Can Kant's Ethics Survive the Feminist Critique?*, *in* FEMINIST INTERPRETATIONS OF IMMANUEL KANT 77, 78 (Robin May Schott ed., 1997) (noting that one major critique "concerns Kant's notion of moral autonomy" which "is said to reflect features more of male than of female identity"); Robin May Schott, *Introduction*, *in* CAN KANT'S ETHICS SURVIVE FEMINIST CRITIQUE? 1, 5 (1990) ("Given Kant's explicit endorsement of the subordination of wives to their husbands, and to the exclusion of women from intellectual or political rights, it is no surprise that many feminists consider Kant to be an exemplar of philosophical sexism.").

182. *See generally* Catherine Myser, *White Normativity in U.S. Bioethics: A Call and Method for More Pluralist and Democratic Standards and Policies*, *in* DEFINING VALUES AND OBLIGATIONS 241 (2007).

during a period when DOGC data acquisition norms were widely followed in the United States, at least in medicine and biomedical research.[183]

An interesting question is why this defect of bioethics—namely, its implicit assumption of human sameness—is only now becoming so apparent. Diversity has always existed; people were never all the same. This truth was well understood at least as far back as 1892, when Sir William Osler famously said of the practice of medicine, "If it were not for the great variability among individuals, medicine might as well be a science and not an art."[184] In light of this truth, how did twentieth-century bioethics function as well as it has done for so long?

Biomedicine's rapid advances after 1940, ironically, came by ignoring that truth and embracing a study methodology—the randomized, controlled clinical trial (RCT)—that masks human variability by presuming sameness.[185] The first modern, multicenter RCT was reported in 1948.[186] RCTs gained a central role in biomedicine after the 1962 Drug Amendments[187] required them as the basis for the FDA to approve new drugs.[188] In 2009, Janet Woodcock, then head of the FDA's drug division, remarked, "Over the past half century, biomedical science has developed randomized, controlled clinical-trial methods that can distinguish treatment effects [whether a drug works for the

183. *See* THE PRIV. PROT. STUDY COMM'N, *supra* note 21, at 280, 283 (describing the Privacy Protection Study Commission's recommendation for consent in informational research and its subsequent implementation in regulations such as the Common Rule).

184. *See* Lawrence J. Lesko & Janet Woodcock, *Translation of Pharmacogenomics and Pharmacogenetics: A Regulatory Perspective*, 3 NATURE REVIEWS: DRUG DISCOVERY 763, 764 (2004) (quoting Osler).

185. *See* LAWRENCE M. FRIEDMAN ET AL., FUNDAMENTALS OF CLINICAL TRIALS 2 (3d ed. 1998) (noting that the term "clinical trial" is variously defined but generally includes the following elements: the study is prospective, following study subjects forward in time from a defined baseline point, which need not be a calendar date but could be a stage of illness); *see also* BENGT D. FURBERG & CURT D. FURBERG, EVALUATING CLINICAL RESEARCH: ALL THAT GLITTERS IS NOT GOLD 11 (2d ed. 2007) (making similar points and noting that concurrent groups of study subjects receive either an intervention (one or more treatments that are under study) or a control (either a placebo or an alternative treatment with which the intervention is being compared)). If a clinical trial is randomized, subjects are assigned randomly to receive either the intervention or the control. *See id.* at 11-12.

186. *See* FRIEDMAN ET AL., *supra* note 185, at 1 (citing Streptomycin in Tuberculosis Trials Comm'n., Med. Research Council, *Streptomycin Treatment of Pulmonary Tuberculosis*, 2 BRIT. MED. J. 768, 769-78 (1948)); *see also* Robert M. Califf, *Clinical Trials Bureaucracy: Unintended Consequences of Well-Intentioned Policy*, 3 CLINICAL TRIALS 496, 496 (2006) (same). *But see* FRIEDMAN ET AL., *supra* note 185, at 1 (noting the use of randomization via coin toss in a 1931 clinical trial, J.B. Amberson et al., *A Clinical Trial of Sanocrysin in Pulmonary Tuberculosis*, 25 AM. REV. TUBERCULOSIS 401 (1931), which also was the first reported study to use blinding).

187. Drug Amendments of 1962, Pub. L. No. 87-781, 76 Stat. 780 (codified as amended in scattered sections of 21 U.S.C.).

188. *See* 21 C.F.R. § 314.50(d)(5)(ii) (2022) (calling for data from controlled clinical studies in new drug applications).

average person in the RCT] from the noise of human variability."[189] There you have it: Did you ever volunteer for a clinical trial? If so, good for you. If not, the FDA views you as medical "noise."

RCTs cut through the noise of human difference to ferret out our presumed sameness (whether or not we really are the same). This approach was highly useful for several decades. It let scientists, undistracted by human variability, discover some useful facts that are true on average. It deferred science's reckoning with human difference into the future—that is, until now.

A respected 1996 report honored RCTs as the highest-quality form of medical evidence.[190] This reverence for RCTs can be seen, in part, as a response to the primitive state of information technology in the twentieth century.[191] Science lacked the computational tools to study large samples of the human population and had to content itself with running little 600- to 3,000-person RCTs and hoping the results represented everybody.[192] Biomedical science either had to assume a high degree of human sameness or else own up to its own narrow validity.

Problems with the 1962 drug approval framework, with its reliance on RCTs, already were evident by the mid-1970s, when a Joint Commission on Prescription Drug Use confirmed that safe and effective FDA-approved drugs were neither safe nor effective for many Americans.[193] Drugs approved on data from predominately male clinical trial populations injure women, for example.[194] This and other failures of evidence from RCTs rebutted the presumption of

189. Janet Woodcock & Lawrence J. Lesko, *Pharmacogenetics—Tailoring Treatment for the Outliers*, 360 NEW. ENGL. J. MED. 811, 811 (2009).

190. *See generally* U.S. PREVENTIVE SERVS. TASK FORCE*, supra* note 88.

191. *See* Barbara J. Evans, *Seven Pillars of a New Evidentiary Paradigm: The Food, Drug, and Cosmetic Act Enters the Genomic Era*, 85 NOTRE DAME L. REV. 419, 438 (2010) [hereinafter Evans, *Seven Pillars*].

192. *See* COMM. ON THE ASSESSMENT OF THE U.S. DRUG SAFETY SYS., INST. OF MED., THE FUTURE OF DRUG SAFETY 36 (Alina Baciu et al. eds., 2007), http://books.nap.edu/openbook.php?record_id=11750 [https://perma.cc/NT5B-GVNN] (discussing typical size and duration of clinical drug trials and limits to their generalizability).

193. *See generally* JOINT COMM'N ON PRESCRIPTION DRUG USE, FINAL REPORT (1980) (assessing the clinical safety of drugs FDA had previously approved using clinical trial data).

194. *See, e.g.*, Niti R. Aggarwal et al., *Sex Differences in Ischemic Heart Disease*, 11 CIRCULATION: CARDIOVASC. QUAL. & OUTCOMES 1, 3, 10 (2018) (discussing gender differences in the effectiveness of medicines tested on predominantly male trial populations); INST. OF MED., WOMEN'S HEALTH RESEARCH: PROGRESS, PITFALLS, AND PROMISE (2010), https://nap.nationalacademies.org/catalog/12908/womens-health-research-progress-pitfalls-and-promise [https://perma.cc/7M5N-85AD] (same). *See generally* G. Tripepi et al., *Bias in Clinical Research*, 73 KIDNEY INT. 148 (2008) (discussing various biases in results from clinical trials).

human sameness, amid mounting evidence of disparities in health outcomes.[195] The same treatments do not produce the same results for everyone.[196]

As the twentieth century closed, the pretense of human sameness was increasingly untenable. Improved information technology ushered in "an era of large volumes of data on platforms conducive to analyses."[197] The shift to informational research—large-scale data-driven discovery using people's health data and biospecimens—was underway.[198] AI/ML medical software is one outgrowth of that shift. In an AI-enabled health care system, the presumption of sameness in twentieth-century bioethics is no longer workable.

Virtually all bioethicists support justice and equity in health care but rarely question whether consent-based data acquisition norms might be feeding injustice. This is reminiscent of the way most white people support racial justice until conversation turns to the matter of reparations. Lockean assertions about individual data ownership and control are as hard to shake as the Lockean theory that one's property is the product of one's own meritorious labor, and good luck and privilege had nothing to do with it.[199] It remains hard for many Americans even to contemplate whether control-over-information theory and hallowed consent norms might be contributing to health care inequities.

195. *See, e.g.*, INST. OF MED., GUIDANCE FOR THE NATIONAL HEALTHCARE DISPARITIES REPORT (Elaine K. Swift ed., 2002), https://www.nap.edu/catalog/10512/guidance-for-the-national-healthcare-disparities-report [https://perma.cc/3X5P-NNEE] (discussing mounting evidence of health care disparities).

196. *See* Fred D. Brenneman et al., *Outcomes Research in Surgery*, 23 WORLD J. SURGERY 1220 (1999).

197. *See* Irony, *supra* note 9, at 93, 95.

198. *See* Spector-Bagdady, *supra* note 8, at 4 (discussing the shift from human subjects research that studies people's bodies to informational "research with all the stuff [such as data and biospecimens] derived from them").

199. *See* Carol M. Rose, *Possession as the Origin of Property*, 52 U. CHI. L. REV. 73, 73 (1985) (describing Locke's theory that individuals establish ownership of a thing by commingling their labor with it, and citing LAWRENCE BECKER, PROPERTY RIGHTS: PHILOSOPHIC FOUNDATIONS 49 (1977) for the notion that a psychological sense of desert is a driving force behind the appeal of this labor theory).

C. The Critique That Consent Norms Fail to Protect Privacy

A growing body of scholarly work questions whether consent provides effective privacy protection.[200] The problem with control-over-information theory is that the control it provides is largely an illusion.[201]

In modern, large-scale, highly interconnected information environments, privacy is interdependent: each individual's privacy is "affected by the decisions of others, and could be out of their own control."[202] This is seen in social networks, where a friend's decision to post a photo of a party you both attended can expose embarrassing information about you.[203] It is also true in large-scale informational research environments. The Common Rule and HIPAA Privacy Rule both define "research" as a systematic investigation that aims to produce *generalizable* knowledge.[204] Generalizable studies reveal facts about everybody to whom the results are generalizable, and this is true whether or not they consented for their data to be used in the study.

Suppose, for example, that Tommie is a thirty-year-old woman who enjoys white wine and is asked to contribute her health records and lifestyle information to a study of early-onset Alzheimer's disease. Tommie refuses to consent, fearing that the researchers might suffer a data breach and leak her sensitive personal input data, resulting in stigmatization, discrimination, or embarrassment. The study proceeds without Tommie's data and discovers a strong statistical association between white wine consumption and early-onset cognitive decline. Even though Tommie declined to share her data with the researchers,

---

200. *See, e.g.*, THE IEEE GLOB. INITIATIVE ON ETHICS OF AUTONOMOUS AND INTELLIGENT SYS., ETHICALLY ALIGNED DESIGN: A VISION FOR PRIORITIZING HUMAN WELL-BEING WITH AUTONOMOUS AND INTELLIGENT SYSTEMS, VERSION 2, 102-05 (2017), https://standards.ieee.org/wp-content/uploads/import/documents/other/ead_v2.pdf [https://perma.cc/3M9E-EVEY]; *see also* WOODROW HARTZOG, PRIVACY'S BLUEPRINT: THE BATTLE TO CONTROL DESIGN OF NEW TECHNOLOGIES (2018); Cate, *supra* note 5, at 1797; Khalid El Emam et al., *A Systematic Review of Re-Identification Attacks on Health Data*, PLOS ONE, Dec. 2011; Georgios A. Kaissis et al., *Secure Privacy-Preserving and Federated Machine Learning in Medical Imaging*, 2 NATURE MACH. INTEL. 305 (2020); Paul Ohm, *Broken Promises of Privacy: Responding to the Surprising Failure of Anonymization*, 57 UCLA L. REV. 1701 (2010).

201. *See, e.g.*, Ellen W. Clayton et al., *The Law of Genetic Privacy: Applications, Implications, and Limitations*, 6 J.L. & BIOSCIENCES 1, 36 (2019) ("The first step to meaningful protection of genetic privacy may be the societal recognition that health privacy, including genetic privacy, is now largely a mirage.").

202. Gergely Biczók & Pern Hui Chia, *Interdependent Privacy: Let Me Share Your Data*, *in* PROCEEDINGS OF THE 17TH INTERNATIONAL CONFERENCE ON FINANCIAL CRYPTOGRAPHY AND DATA SECURITY (2013).

203. *See* Mathias Humbert, *When Others Impinge Upon Your Privacy: Interdependent Risks and Protection in a Connected World* 66 (Mar. 13, 2015) (Ph.D. Thesis, École Polytechnique Fédérale de Lausanne), https://infoscience.epfl.ch/record/205089/files/EPFL_TH6515.pdf [https://perma.cc/PC4M-HVYW].

204. *See* 45 C.F.R. § 46.102(l) (2022) (Common Rule); *see also* 45 C.F.R. § 164.501 (2022) (HIPAA Privacy Rule).

and hackers could not possibly find any data about her in the study data set, the results of the study nevertheless could stigmatize Tommie. Anybody who sees Tommie sipping a glass of white wine at a party might infer that she is at risk for early-onset cognitive issues. Because the study findings are generalizable, they affect her whether or not she participated in the study. Control-over-information theory empowers people to block the use of their data as study *inputs*, yet Tommie's stigmatization flows from the study *outputs* and the adverse inferences about her that can be drawn from them.

If you define "privacy" as the individual's ability to control uses of personal information, then controlling personal information gives you "privacy."[205] But this is a circular, tautological privacy protection. Controlling flows of our personal *input* data may not protect what we really care about in the way of privacy protection. For example, it will not protect us from having unwanted inferences drawn about us based on data, including our own and data from other people, that already are circulating in the world.

Notice-and-consent data privacy schemes resemble the "dummy thermostats" in modern office buildings: simulated thermostats (some of which even make hissing sounds when manipulated) that give people the illusion they can control their office temperature when, in reality, the thermostats are not connected to the heating and air conditioning system.[206] There is some evidence that people's perception of thermal comfort has a placebo effect, so that dummy thermostats promote workplace harmony and worker satisfaction.[207] Pretending that notice, consent, and de-identification protect people's data privacy has no beneficial placebo effects: believing your privacy is protected does not protect it. Consent norms may foster societal harmony by making people feel they can control their privacy risks. Yet, to a large degree,

205. *See supra* note 23 (citing sources that have defined privacy as an individual right to control personal information).

206. *See* Jared Sandberg, *Employees Only Think They Control Thermostat*, WALL ST. J. (Jan. 15, 2003, 12:01 AM), https://www.wsj.com/articles/SB1042577628591401304 [https://perma.cc/HMW5-7SEZ] (reporting that "[a] lot of office thermostats are completely fake—meant to dupe you into thinking you've altered the office weather conditions" and citing an HVAC specialist for the estimate that "90% of office thermostats are dummies" but noting variation in estimates).

207. *See* Daven Hiskey, *Most Thermostat Controls in Large Office Buildings Don't Do Anything*, TODAY I FOUND OUT (July 18, 2012), http://www.todayifoundout.com/index.php/2012/07/most-thermostat-controls-in-large-office-buildings-dont-usually-do-anything/ [https://perma.cc/7J6A-9957] (providing examples of installations of controls that hiss to give employees the impression that the fake thermostat works); *see also* Barbara A. Checket-Hanks, *Placebo Stats*, AIR CONDITIONING HEATING REFRIGERATION NEWS (March 27, 2003), https://www.achrnews.com/articles/92414-placebo-stats [https://perma.cc/ERY8-49GK] (noting that, "[i]n many cases, the placebo effect of the nonfunctional thermostats seems to be successful").

privacy loss is interdependent and systemic; it is baked into the fabric of the data-driven informational economy we, as a society, have chosen to inhabit.[208] Individuals have only a limited degree of control.

This, then, is the most damning critique of control-over-information theory: it does not provide strong privacy protection.[209] There are more effective ways to protect privacy in the large data sets used to train and operate artificial intelligence/machine learning (AI/ML) software. Computational methods of privacy protection (privacy by design or "PBD") can provide effective, measurable protection.[210] For example, federated learning can improve data security and privacy when combined with other approaches such as differential privacy, homomorphic encryption, secure multiparty computation, and secure hardware implementation.[211] In large-scale, complex information systems—such as AI/ML CDS tools—protecting privacy requires engineering and computer science skills. We are no longer in the era, just fifty years ago, when privacy was thought to be protected by discussing autonomy as conceived by Locke and Kant.

The HIPAA Privacy Rule is designed in a way that could support greater reliance on PBD, because it recognizes two alternative methods for de-identifying data. The first, the so-called "safe harbor" method, involves stripping away specific data elements (e.g., names, zip codes, hospital admission dates, among others) that are thought to allow data to be re-identified; it is in wide use but is widely criticized as ineffective.[212] The second method, less widely used, is

208. *See* Cynthia Dwork et al., *Calibrating Noise to Sensitivity in Private Data Analysis*, *in* PROCEEDINGS OF THE THIRD THEORY OF CRYPTOGRAPHY CONFERENCE 265 (2006) (conceiving, in a groundbreaking work in the field of computer science, that the amount of privacy we lose by consenting to include our personal information in a data set is merely a "differential" over and above the baseline of privacy loss we already suffer because of the inferences that can be drawn about us even when our data are not included); *see also* Alexandra Wood et al., *Differential Privacy: A Primer for a Non-Technical Audience*, 21 VAND. J. ENT. & TECH. L. 209 (2018) (providing a highly accessible description of Dwork et al.'s theory of differential privacy).

209. *See supra* note 200 (citing various sources questioning the effectiveness of notice and consent as a privacy protection and suggesting alternative approaches); *see also* Evans, *Rules for Robots*, *supra* note 18, at 8-10 (discussing the interdependent and systemic nature of privacy risk in large-scale, generalizable data analysis, such that opting out of contributing one's data for analysis no longer fully protects against having adverse or stigmatizing inferences drawn about oneself).

210. *See generally* Kaissis et al., *supra* note 200, at 307-09 (summarizing these PBD techniques). *See also* Ittai Dayan et al., *Federated Learning for Predicting Clinical Outcomes in Patients with COVID-19*, 27 NATURE MED. 1735 (2021) (applying federated learning to train AI models with data from multiple sources while protecting privacy); Kobbi Nissim et al., *Bridging the Gap Between Computer Science and Legal Approaches to Privacy*, 31 HARV. J.L. & TECH. 687, 714-33 (2018) (discussing differential privacy).

211. *See* Kaissis et al., *supra* note 200, at 308-39.

212. *See* 45 C.F.R. § 164.514(b)(2)(i) (2022) (allowing the so-called "safe harbor" de-identification method, which strips away eighteen specific types of data elements that might

statistical de-identification, which correctly frames data privacy as a probabilistic phenomenon subject to effective and measurable protection, just as PBD frames it.[213] Through PBD, meaningful privacy protection can be built into the design of data systems and algorithms that store and process people's data. PBD is not a privacy panacea, however. It is computationally expensive, requiring extra grinding by computers to protect privacy while still producing the desired outputs.[214] Some PBD approaches deliberately blur outputs to make re-identification more difficult, which forces tough tradeoffs between privacy and accuracy.[215] Other PBD approaches are costly to set up and administer.[216] Still, PBD can deliver effective privacy protection that is quantifiable, which is more than de-identification and consent can boast.

Moving past the false belief that safe harbor de-identification and consent protect privacy could open the door to new data acquisition strategies aimed at reducing the racial, gender, and class-related biases that infect the current generation of AI/ML CDS tools while offering superior privacy protection to what we have today. In a crucial respect, the Privacy Rule was ahead of its time, enabling modern, computational methods of privacy protection.[217] These methods are capable of outperforming the feeble de-identification and consent-based privacy protections that have been in vogue since the 1970s, but only if there is societal consensus that it is time to take advantage of them.

III. LEGAL PATHWAYS TO DIVERSE, INCLUSIVE AI/ML TRAINING DATA

This Part explores legal pathways available under the Privacy Rule that offer opportunities to reduce consent bias in AI/ML training data and promote greater equity in AI-enabled health care. These promising pathways exist against the backdrop of an inconvenient fact:

---

serve to identify a person); *see, e.g.*, Ohm, *supra* note 200, at 1736-38, 1740-41 (discussing the risks of re-identification of data protected using the HIPAA Privacy Rule's safe harbor de-identification standard).

213. *See* 45 C.F.R. § 164.514(b)(1) (allowing certification by "[a] person with appropriate knowledge of and experience with generally accepted statistical and scientific principles" that "the risk is very small" that the information could be re-identified). *See generally* Cynthia Dwork & Aaron Roth, *Algorithmic Foundations of Differential Privacy*, 9 FOUNDATIONS & TRENDS THEORETICAL COMPUT. SCI. 211 (2014) (explaining differential privacy and its probabilistic foundations).

214. *See* Kaissis et al., *supra* note 200, at 305.

215. *See, e.g.*, Zhiyu Wan et al., *Sociotechnical Safeguards for Genomic Data Privacy*, 23 NATURE REVIEWS GENETICS 429, 430-37 (2022) (comparing various privacy technical safeguards used in PBD and noting that some of them lose too much accuracy or pose other practical limits to be suitable for genomic data).

216. *See generally* Dayan et al., *supra* note 210, at 1735-37 (discussing federated learning, which can entail significant coordination efforts among distributed data sites).

217. *See e.g.*, *supra* Table 1, Norm 2b.

during the twenty-five years since Congress enacted HIPAA, problems with health equity have scarcely been eliminated and may even have grown worse. How, then, can the Privacy Rule be touted as an engine of future health equity if, twenty-five years past its inception, inequity is still rampant? The short answer is that the Privacy Rule's equity-serving data acquisition pathways have, till now, been underutilized, and Part IV will explore why that happened. For now, in Part III, the goal is simply to identify several legal pathways for accessing data under the Privacy Rule that hold promise, if the goal is to make AI/ML CDS tools safe, effective, and more equitable.

The five data access pathways discussed below invoke eight of the Privacy Rule's 27 informational norms summarized in Table I. Prominent among these is the treatment exception which facilitates flows of clinical data for use in treating patients (Norm 6 in Table 1). Because AI/ML CDS tools are, by definition, intended for use in clinical health care, using data to train and operate these systems potentially constitutes a treatment purpose under that norm. Another relevant norm allows disclosure of data for payment and health care operations (Norm 14). Two norms allow access to data for biomedical research, through waivers of individual authorization (Norm 13) or as almost fully de-identified limited data sets pursuant to a Data Use Agreement (DUA) (Norm 18). For FDA-regulated CDS software,[218] software developers may be able to access data to support compliance with FDA regulatory requirements (Norm 16). Finally, the Privacy Rule offers pathways for sharing data with public health authorities and their contractors (Norm 15) and with health oversight agencies (Norm 17)—pathways that could help the FDA and state health oversight agencies, including medical practice regulators, develop diverse, inclusive data resources to ensure that AI/ML CDS tools are safe and effective for all patients who, in actual clinical practice, will be relying on them. These pathways are described in turn below.

A. AI/ML CDS Software as a Treatment Use of Data

HIPAA's treatment exception allows covered entities to use and disclose PHI without individual authorization to a health care provider for treatment purposes, including for treatment of people other than the patient the data describe.[219] There is no "minimum necessary" restriction on such disclosures.[220] This informational norm replicates information flows that traditionally took place within a family doctor's

218. *See* 21st Century Cures Act, Pub. L. No. 114-255, § 3060(a), 130 Stat. 1033 (2016) (codified at 21 U.S.C. § 360j(o)(1)(E)) (defining the scope of the FDA's jurisdiction to regulate CDS software).

219. *See* 45 C.F.R. § 164.502(a)(1)(ii) (2022); *see also* HHS, FAQ 512, *supra* note 35.

220. *See* 45 C.F.R. § 164.502(b)(2)(i).

brain as the doctor compared and contrasted the current patient with past observations about other patients. The HIPAA Privacy Rule provides an important alternative privacy protection: such disclosures can only be made to health care providers already under strong state fiduciary duties of confidentiality.[221]

CDS tools, by definition, are designed for use by clinicians to treat patients in clinical health care settings.[222] Data used when developing CDS tools, as well as data incorporated into the final software products, arguably fit within HIPAA's treatment exception. The data are informing a health care provider's clinical intuition, but in a modern way where clinical intuition is informed by recommendations from CDS tools.

The precise scope of what counts as a "treatment" use of data is the crucial question in determining which uses qualify for access under the Privacy Rule's norm for treatment-related disclosures. Is training software that is eventually destined for use in clinical settings "treatment"? Or does the use of data in CDS software only count as a treatment use after the software has been delivered into the hands of a specific clinician who is using it to treat a specific patient during a specific treatment encounter? There are, as yet, no precedents to guide this line-drawing.

The HIPAA statute and Privacy Rule do not define what counts as "treatment." Their silence accords with other major federal health care statutes, such as the Clinical Laboratory Improvement Amendments of 1988[223] and the Food, Drug, and Cosmetic Act,[224] which use terms like "diagnosis," "prevention," and "treatment" without defining their precise scope. Their silence respects federalism. Federal laws affect the practice of medicine, but Congress and federal agencies try to respect the states' primacy in the sphere of medical practice regulation.[225]

The states, through their medical practice acts, other health statutes, and common law, define the scope of medical practice and clinical care, including which activities constitute medical treatment.[226] Thus,

---

221. *See supra* notes 46-48 and accompanying text.

222. *See supra* notes 111-13 (defining CDS tools).

223. Pub. L. No. 100-578, 102 Stat. 2903 (1988) (codified as amended at 42 U.S.C. § 263a).

224. Pub. L. No. 75-717, 52 Stat. 1040 (1938) (codified as amended at 21 U.S.C. §§ 1-2252).

225. *See, e.g.*, David G. Adams, *The Food and Drug Administration's Regulation of Health Care Professionals*, *in* FUNDAMENTALS OF LAW AND REGULATION 423, 423 (David G. Adams et al. eds., 1999) (noting that the FDA, as a matter of policy, "has traditionally taken the position that it does not regulate the practice of medicine or pharmacy and has generally avoided regulatory actions that would directly restrict or interfere with professional service to patients").

226. *See* 46 AM. JUR. PROOF OF FACTS 2D *Existence of Physician and Patient Relationship* §§ 3, 5, 6, 9, Westlaw (database updated July 2023) (providing an overview of the states' roles

the states ultimately determine the scope of the Privacy Rule's treatment exception. Whether a given use of data in CDS software has a treatment purpose is for the states to decide. This grants the states a significant, as-yet-unutilized power to affect how freely data can flow to AI/ML medical software to reduce its racial, gender, and socioeconomic biases and, ultimately, to promote more equitable health care. If states pass laws defining the development, training, and utilization of CDS software as a treatment use of data, then such activities would be eligible to acquire data using the Privacy Rule's treatment exception.

A critique is that this access pathway has the potential to feed a form of systemic bias that Nicholson Price has described.[227] Namely, the treatment exception creates differential access to data by allowing data to move to clinician-researchers (for example, by medical academics marshaling data from their own health care institutions to develop CDS software), while disfavoring access by non-physician software developers.[228] This could concentrate CDS software development activities at academic medical centers staffed by researchers who are practicing clinicians, resulting in CDS tools that work best for privileged patient populations treated at those sites. Independent, non-physician software developers, who might design tools to serve a more diverse population, cannot take advantage of this access pathway.

This differential access reflects valid privacy concerns: non-physician software developers are not subject to the fiduciary duties states impose on physicians and thus would evade the alternative privacy protection on which the Privacy Rule's treatment exception relies. This suggests a difficult tradeoff between protecting privacy and addressing systemic bias. Fortunately, there is a way to avoid that tradeoff. If states enact laws defining CDS software as a treatment use of data, the states could simultaneously place CDS software developers under the same information fiduciary duties that apply to clinicians.[229] Then, treatment disclosures—whether made to a clinician or to a non-physician software developer—would be on equal footing, with the same level of data access and the same alternative privacy protections. Software developers unwilling to embrace these state information fiduciary duties could still gain access to data through other pathways but would be ineligible for access under the Privacy Rule's treatment exception.

---

in these areas); *see also* Patrick D. Blake, Note, *Redefining Physicians' Duties: An Argument for Eliminating the Physician-Patient Relationship Requirement in Actions for Medical Malpractice*, 40 GA. L. REV. 573, 601 (2006) (same).

227. *See* Price, *supra* note 131, at 66-67, 91-94.

228. *See* HHS, FAQ 512, *supra* note 35.

229. *See supra* notes 46-48 and accompanying text (summarizing these fiduciary duties).

B. Diversion of Health Data from Health Care Operational Uses

This pathway merits discussion because it is a major pathway through which clinical health data leak into a variety of research and commercial uses, which could include use in AI/ML CDS software. Because this pathway is already widely used, the aim here is not to recommend it but to critique it and propose steps for addressing gaps in the privacy protections it provides.

The Privacy Rule allows unconsented disclosure or use of PHI for health care operations.[230] This allows HIPAA-covered providers to use or disclose PHI without consent for a broad array of business operational purposes which include—among other things—quality improvement studies to explore ways to provide better care to their patient populations.[231] A 2021 proposed rulemaking, if finalized, would expand the scope of permitted disclosures under this norm by removing some business operational data flows from even the light protection of HIPAA's minimum necessary standard.[232]

When health care providers conduct quality improvement studies in-house, the data remain protected by the Privacy Rule and fiduciary duties the provider already has under state law, so the use poses little incremental privacy risk. In practice, however, many providers outsource their business operational computing to external information service providers, and the Privacy Rule's health care operations provision lets them transfer PHI without consent to the service provider for these operational activities.[233] After receiving the PHI, external information service providers can divert the data to other uses (such as to train AI/ML tools the information service provider happens to be developing) by de-identifying the data. Once de-identified, information is no longer PHI and is no longer protected by the Privacy Rule.[234] De-identifying data, of course, limits its utility and

230. *See supra* Table 1, Norm 14.

231. *See* 45 C.F.R. §§ 164.502(a)(1)(ii), .506 (2022).

232. *See* Proposed Modifications to the HIPAA Privacy Rule To Support, and Remove Barriers to, Coordinated Care and Individual Engagement, 86 Fed. Reg. 6446, 6533 (proposed Jan. 21, 2021) (proposing to add a new exception at 45 C.F.R. § 164.502(b)(2)(vii) of the Privacy Rule exempting certain disclosures to health plans from HIPAA's minimum necessary standard); *see also* Ellen W. Clayton et al., Comments on Proposed Modifications to the HIPAA Privacy Rule To Support, and Remove Barriers to, Coordinated Care and Individual Engagement (May 5, 2021), https://www.regulations.gov/comment/HHS-OCR-2021-0006-1116 [https://perma.cc/LU5D-BCX9] (noting that disclosures to health plans are a payment or operational purpose, traditionally subject to minimum necessary restrictions, and noting that the proposed revision fundamentally alters the Privacy Rule and "opens a conduit for patients' whole medical records to flow into a wide variety of 'big data' applications without their consent").

233. *See supra* note 231.

234. *See* Jim Hawkins et al., *Non-Transparency in Electronic Health Record Systems*, *in* TRANSPARENCY IN HEALTH AND HEALTH CARE IN THE UNITED STATES 273, 280 (Holly Fernandez Lynch et al. eds, 2019).

can make it harder to detect biases.[235] The fact that this data access pathway is widely used could be a contributing factor in the biases observed in currently available CDS tools.

Aware of the potential for operational data to be diverted for other purposes, HHS amended the Privacy Rule in 2013 so that it now applies to a covered entity's business associates—parties that receive PHI from a covered entity in order to perform services for the covered entity.[236] Information service providers performing business operational computing for a HIPAA-covered health care provider clearly are its business associates and now are HIPAA-covered entities themselves. This, unfortunately, falls short of making them information fiduciaries, because the Privacy Rule looks to external sources of law to impose fiduciary duties.[237] Moreover, as HIPAA-covered entities, business associates can exercise the Privacy Rule's full panoply of disclosure-friendly informational norms which, for example, would allow them to create their own Institutional Review Boards (IRBs) to approve research uses of the PHI without individual consent.[238] This is a serious gap in the Privacy Rule's protections.

The Privacy Rule does require a written Business Associate Agreement (BAA) between the covered entity and business associate.[239] In theory, the BAA could impose contractual duties for a business associate to act as an information fiduciary with respect to the PHI it receives. Unfortunately, HHS's model BAA terms do not include any limitations on re-identification, redisclosure, or reuse of data.[240] Jim Hawkins et al. characterize it as "fanciful" to assume covered entities will insert meaningful privacy protections in BAAs, absent legal requirements to do so.[241] Under modern contract theory, each contract term has a price.[242] When a covered entity is purchasing software services from a business associate, demanding contract terms that impose information fiduciary duties on the business associate

235. U.S. GOV'T ACCOUNTABILITY OFF., *supra* note 121, at 24.

236. *See* 45 C.F.R. § 160.103 (2022) (defining business associate); *see also Direct Liability of Business Associates*, U.S. DEP'T HEALTH & HUM. SERVICES (July 16, 2021), https://www.hhs.gov/hipaa/for-professionals/privacy/guidance/business-associates/fact-sheet/index.html [https://perma.cc/9M5Z-HG5F] (discussing the 2013 revisions).

237. *See supra* notes 53-54 and accompanying text (discussing HHS's determination that it lacked jurisdiction to impose fiduciary duties on downstream data recipients under the Privacy Rule).

238. *See supra* Table 1, Norm 13.

239. *See Business Associate Contracts: Sample Business Associate Agreement Provisions*, U.S. DEP'T HEALTH & HUM. SERVICES (Jan. 25, 2013), https://www.hhs.gov/hipaa/for-professionals/covered-entities/sample-business-associate-agreement-provisions/index.html [https://perma.cc/PJG3-DEX9].

240. *See* Hawkins et al., *supra* note 234, at 281 (discussing deficiencies of HHS's sample BAA).

241. *Id.* at 275.

242. *See* Michael I. Meyerson, *The Efficient Consumer Form Contract: Law and Economics Meets the Real World*, 24 GA. L. REV. 583, 589-90 (1990).

would make the negotiated price of those software services go up. Unless HHS requires covered entities to require meaningful privacy protections in BAAs, they seemingly have no incentive to do it.

C. *Two Pathways of Access for AI/ML Research*

Covered entities can supply data without consent for use in research in two ways: as an almost fully de-identified limited data set or by having an IRB or special-purpose HIPAA privacy board (together, "IRB") waive the requirement for individual authorization.[243] Access to limited data sets requires a Data Use Agreement (DUA) restricting the recipient's reuse, re-identification, or redistribution of data.[244] No DUA is required for access to data under a waiver, which relies on IRB oversight as its method of ensuring privacy risks have been minimized, although the waiver provision does require IRBs to require "[a]dequate written assurances" that the data recipient will not reuse or disclose the data inappropriately.[245]

The Privacy Rule's waiver provision is subject to the minimum necessary standard, but this standard can still allow access to identifiable data if the researcher justifies why identifiers are needed.[246] Researchers receiving data under HIPAA's waiver provision sometimes are HIPAA-covered entities, for example if they are employed by a HIPAA-covered academic medical center.[247] More generally, however, this pathway can move data, potentially in identifiable form, to researchers who are neither information fiduciaries nor HIPAA-covered.

Concerned by this lack of downstream privacy protections once data are shared with researchers, many IRBs hesitate to grant waivers for identifiable data.[248] Their caution is understandable from a privacy

243. *See supra* Table 1, Norm 18 (limited data set); *id.* Norm 13 (waiver provision).

244. *See* 45 C.F.R. § 164.514(e)(3)(i), (e)(4) (2022).

245. *See id.* § 164.512(i); *see id.* § 164.512(i)(2)(ii)(A)(3) (requiring "[a]dequate written assurances" concerning the recipient's reuse and redisclosure of the data received under a waiver).

246. *See* 45 C.F.R. § 164.512(i)(2)(ii)(A)(1)-(2) (requiring an IRB, when approving a waiver for release of data with identifiers, to require a plan to protect the identifiers from improper use and disclosure and a plan for destroying the identifiers at the earliest opportunity consistent with conduct of the research); *see also* Letter from William W. Stead, *supra* note 95, at 9 (discussing operation of the minimum necessary standard).

247. *See* 45 C.F.R. §§ 160.102-.103 (2022) (clarifying who is a covered entity).

248. Barbara J. Evans & Harlan M. Krumholz, *People-Powered Data Collaboratives: Fueling Data Science with the Health-Related Experiences of Individuals*, 26 J. AM. MED. INFO. ASS'N 159, 160 (2018) ("Twenty-first century science often requires multidimensional data assembly—not only from many individuals but also across many different data holders that store portions of each person's data—to capture people's complete experiences over time. De-identifying data can thwart necessary linkages, and there is growing awareness that large, linked datasets are inherently re-identifiable, which makes many IRBs reluctant to approve waivers.").

standpoint, but it can impede efforts to reduce bias in AI/ML CDS software. Removing identifiers can diminish data utility—for example, making it harder to monitor the inclusivity of AI/ML training data, rendering patients who may have special medical needs (e.g., transgender patients) invisible, or making it difficult (or impossible) to link HIPAA-covered clinical data with external sources of information bearing on social determinants of health (e.g., data on fitness, nutrition, or education).[249] IRBs may feel more comfortable sharing limited data sets for use in research, but the near-complete de-identification that occurs when creating a limited data set can create some of these same issues with data utility.[250]

It is time to face a troubling fact: anonymity masks the invidious reality that some patient subgroups are routinely underserved, erased, or left out of AI/ML training data. Tackling the nation's ongoing health disparities and developing equitable medical AI may require improved access to identifiable health data for use in AI/ML research. For such access to be remotely acceptable, from a privacy standpoint, recipients of the data need to be placed under strong information fiduciary duties. The Privacy Rule's waiver provision imposes no such duties.[251] That is one major gap.

The waiver provision has a second crucial gap. Many scholars agree that the "central ethical issue" in research that uses people's health information is to ensure that the research offers sufficient public benefit to justify the burdens it places on individual privacy.[252] Even as bioethicists advocated for consent norms in the 1970s, they acknowledged that some unconsented research uses of data are ethically justified. A National Commission formed under the National Research Act of 1974 published influential recommendations in 1978.[253] For research that uses existing health data, the Commission

249. *See id.*; *see also* U.S. GOV'T ACCOUNTABILITY OFF., *supra* note 121, at 24.

250. *See supra* Table 1, Norm 18; 45 C.F.R. § 164.514(e)(3)(i), (e)(4).

251. *See supra* notes 53-54 and accompanying text.

252. *See* Antoine C. El Khoury et al., *Bioethical Issues in Pharmacoepidemiologic Research*, *in* PHARMACOEPIDEMIOLOGY 623, 637 (Brian L. Strom et al. eds., 5th ed. 2012) ("The central ethical issue in pharmacoepidemiologic research is deciding what kinds of projects will generate generalizable knowledge that is widely available and highly valued, and do this in a manner that protects individuals' right to privacy and confidentiality."); *see also* NAT'L BIOETHICS ADVISORY COMM'N, ETHICAL AND POLICY ISSUES IN RESEARCH INVOLVING HUMAN PARTICIPANTS xviii, 103 (2001), http://bioethics.georgetown.edu/nbac/human/overvol1.pdf [https://perma.cc/Q9LY-JV5V] (recognizing the need for nonconsensual data use in some circumstances and including, as a necessary criterion, that an IRB determine that "the benefits from the knowledge to be gained from the research study outweigh any dignitary harm associated with not seeking informed consent"); Peter D. Jacobson, *Medical Records and HIPAA: Is It Too Late to Protect Privacy?*, 86 MINN. L. REV. 1497, 1497-99 (2002) (arguing that the most important issue to resolve is which public health objectives are sufficiently important to override the individual's interest in nondisclosure).

253. *See* National Research Act of 1974, Pub. L. No. 93-348, 88 Stat. 342, 348-51 (codified as amended in scattered sections of 42 U.S.C.) (creating the National Commission);

concluded that "[i]f the subjects are not identified or identifiable, the research need not be considered to involve human subjects" and should not require consent.[254] Moreover, even "where the subjects are identified, informed consent may be deemed unnecessary" if certain conditions are met.[255] Those conditions included a public benefit requirement: an IRB must determine that "the importance of the research justifies such invasion of the subjects' privacy."[256]

When drafting the Privacy Rule, HHS was mindful of this recommendation. HHS's original proposal would have required IRBs to determine, before approving a waiver, that "the research is of sufficient importance so as to outweigh the intrusion of the privacy of the individual whose information is subject to the disclosure."[257] Unfortunately, this proposal drew "a large number" of negative comments—many of them from IRBs—expressing doubt that IRBs are competent to balance public and privacy interests or to apply a public benefit criterion consistently.[258] HHS ultimately backed down and dropped the burdensome public benefit requirement from the criteria for IRB approval of waivers of consent for research uses of data.[259]

As things stand, people whose data are used in research without consent under the Privacy Rule's waiver provision have no assurance that their sacrifice serves any socially beneficial purpose at all. This reality feeds concerns about "data colonialism" and fuels popular mistrust that clinical health data are being harnessed for software

---

Protection of Human Subjects: Institutional Review Board; Report and Recommendations of the National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research, 43 Fed. Reg. 56174, 56174-98 (Nov. 30, 1978) (publishing recommendations as required by the National Research Act of 1974).

254. *See* Protection of Human Subjects: Institutional Review Board; Report and Recommendations of the National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research, 43 Fed. Reg. at 56181.

255. *Id.*

256. *Id.* at 56179; *see also id.* at 56181 (reporting findings of the Privacy Protection Study Commission, which elaborated this balancing requirement more specifically: "[M]edical records can legitimately be used for biomedical or epidemiological research, without the individual's explicit authorization," provided that the medical care provider (who in all likelihood would have been the data holder in that era of paper records) determines "that the importance of the research or statistical purpose for which any use of disclosure is to be made is such as to warrant the risk to the individual from additional exposure of the record or information contained therein," and provided that an IRB ensures this condition has been met).

257. *See* Standards for Privacy of Individually Identifiable Health Information, 65 Fed. Reg. 82462, 82698 (Dec. 28, 2000) (to be codified at 45 C.F.R. pts. 160, 164).

258. *Id.*

259. *See id.* (revising the balancing requirement in the December 2000 version of the Privacy Rule); *see also* Standards for Privacy of Individually Identifiable Health Information, 67 Fed. Reg. 53182, 53270 (Aug. 14, 2002) (to be codified at 45 C.F.R. pts. 160, 164) (dropping the balancing requirement altogether in the currently effective version of HIPAA's waiver provisions at 45 C.F.R. § 164.512(i)).

developers' private, commercial gain instead of for socially beneficial purposes such as reducing health disparities and ensuring equitable medical AI.[260]

Failure to implement the National Commission's 1978 recommendation was a gross breach of trust with the American people. As long as this breach continues, the data acquisition strategies now needed to achieve equitable medical AI will always have a whiff of illegitimacy.

D. Access to Data by FDA-Regulated Software Developers

The Privacy Rule allows covered entities to provide data to private-sector drug and medical device manufacturers to enable them to perform various tasks that the FDA requires with respect to products already on the market.[261] This provision does not let covered entities supply PHI to help manufacturers develop *new* medical products and prove they are safe and effective, but it helps them comply with regulatory requirements in the postmarketing period after a product is already in clinical use.[262] The postmarketing period is precisely when it is crucial to monitor AI/ML tools closely to make sure that they are providing equitable results in actual clinical use.

A little-known feature of the Privacy Rule is this interaction with the FDA's regulatory requirements. The FDA can, by requiring product manufacturers to provide answers, trigger access to the data they need to provide those answers.[263] Suppose, for example, that the FDA promulgated a regulation requiring sellers of FDA-regulated AI/ML CDS tools to report whether their software is providing biased, unreliable recommendations for patients in certain demographic groups. By imposing this requirement, the FDA would make software developers eligible to receive HIPAA-protected PHI under the Privacy Rule's norm for disclosure of data to FDA-regulated entities.[264]

---

260. *See Glossary: Data Colonialism*, PURDUE UNIV.: CRITICAL DATA STUD., https://purdue.edu/critical-data-studies/collaborative-glossary/data-colonialism.php [https://perma.cc/6SRX-EWQ6] (last visited Sept. 23, 2023) (defining "data colonialism" as "the process by which governments, non-governmental organizations and corporations claim ownership of and privatize the data that is produced by their users and citizens").

261. *See* 45 C.F.R. § 164.512(b)(iii) (2022) (allowing such disclosures "[t]o collect or report adverse events," "[t]o track FDA-regulated products," "[t]o enable product recalls, repairs, or replacement, or lookback (including locating and notifying individuals who have received products that have been recalled, withdrawn, or are the subject of lookback)," or "[t]o conduct post marketing surveillance").

262. *Id.*

263. *Id.*

264. *Id.*

This norm is permissive in the sense of allowing, but not requiring, covered entities to disclose PHI.[265] Thus, the FDA cannot *force* covered entities to disclose data to FDA-regulated manufacturers; it merely has the power to make it lawful under the Privacy Rule for covered entities to do so. The FDA exercises this power by imposing mandatory requirements on the drug and device manufacturers it regulates.[266] This power is an important one, because covered entities possess large stores of information (such as data on patient adverse events and on the clinical performance of devices in use at health care facilities) bearing on the safety of FDA-regulated products. Subject to HIPAA's minimum necessary standard, covered entities regularly share data with manufacturers under this norm to aid their efforts to improve patient safety.

To date, the FDA has relied heavily on voluntary oversight methods for AI/ML CDS tools.[267] Voluntary approaches, unfortunately, do not impose legal compliance requirements that trigger data access under this norm. For several years, the FDA has been "reimagining its approach to digital health medical devices" and its policies are still a work in progress.[268] In 2016, § 3060 of the 21st Century Cures Act placed some (but not all) CDS software within the definition of a medical device that the FDA can regulate.[269] The statute draws a line between CDS tools that are and are not within the FDA's regulatory jurisdiction.[270] The FDA's September 2022 Guidance on Clinical

265. *Cf. supra* Table 1, Norm 3 (HIPAA's right for individuals to inspect and receive copies of their own medical data, which is the only one of the Privacy Rule's 27 informational norms that is mandatory in the sense of requiring covered entities to disclose data).

266. This statement is supported by the plain text of 45 C.F.R. § 164.512(b)(iii), which allows a covered entity to disclose PHI to "[a] person subject to the jurisdiction of the Food and Drug Administration (FDA) with respect to an FDA-regulated product or activity for which that person has responsibility, for the purpose . . . (A) [t]o collect or report adverse events" and "(D) [t]o conduct post marketing surveillance." Persons subject to the FDA's jurisdiction include, for example, FDA-regulated drug and device manufacturers and developers and vendors of FDA-regulated medical software. The FDA has the power, by promulgating regulations, to place the companies it regulates under legally binding responsibilities to "collect or report adverse events" or "conduct post marketing surveillance" for their products. Once the FDA has imposed such a responsibility on an FDA-regulated company, the Privacy Rule allows HIPAA-covered entities to disclose PHI to the company to facilitate its compliance with the FDA's requirement. *Id.* § 164.512(b)(iii).

267. *See infra* notes 268-73 and accompanying text (discussing FDA's reliance on nonmandatory guidance and voluntary programs for regulation of CDS tools).

268. U.S. FOOD & DRUG ADMIN., DIGITAL HEALTH INNOVATION ACTION PLAN 5 (2017), https://www.fda.gov/downloads/MedicalDevices/DigitalHealth/UCM568735.pdf [https://perma.cc/VEX2-6H2Z].

269. *See* 21st Century Cures Act, Pub. L. No. 114-255, § 3060(a), 130 Stat. 1033, 1131-32 (2016) (codified at 21 U.S.C. § 360j(o)).

270. *See* 21 U.S.C. § 360j(o)(1)(E) (providing that CDS software is not subject to FDA regulation if it meets three statutory criteria, one of which addresses whether the health care provider using the software would be able to independently review the basis for the software's recommendations; *see id.* § 360j(o)(1)(E)(iii)); *see also* Adler-Milstein et al., *supra* note 113, at 16-17 (discussing how these statutory criteria affect the FDA's jurisdiction to regulate CDS tools).

Decision Support Software sought to clarify this line.[271] However, the Guidance has been widely criticized as deviating from the line Congress drew in the 21st Century Cures Act.[272] Thus, six years after enactment of the statute, there are continuing questions about the scope of the FDA's authority to regulate CDS tools. Moreover, guidance documents by their nature are voluntary and do not impose legally enforceable requirements.[273] To trigger access to data under HIPAA's exception at 45 C.F.R. § 164.512(b)(iii), the FDA seemingly would need to promulgate regulations as opposed to relying on voluntary guidance documents as a regulatory tool.[274]

As the FDA gains experience and know-how to guide its future policies for AI/ML CDS tools, the software industry no doubt appreciates the agency's measured approach to imposing binding regulatory requirements. Unfortunately, the agency's reluctance to act through mandatory regulations has an unintended consequence: it hinders access under the Privacy Rule to clinical data developers need in their efforts to make CDS tools safe, effective, and more equitable.

E. Creating Common Data Infrastructure for Equitable AI/ML Medical Software

The data acquisition pathways discussed up till now move data directly into the hands of parties engaged in researching, developing, selling, or using AI/ML CDS software.[275] Ultimately, however, the United States needs a shared data infrastructure to promote inclusivity and justice in medical AI.

Nations with single-payer national health systems often have comprehensive health data resources in standard formats for all their citizens, but the fragmented American health care system presents a tougher challenge.[276] Acquiring raw clinical health care data is only

---

271. *See* U.S. FOOD & DRUG ADMIN., *supra* note 112.

272. *See* Barbara J. Evans, *FDA Regulation of Physicians' Professional Speech*, 3 J. FREE SPEECH L. (forthcoming 2023) (manuscript at 19-20, 19 nn.131-32), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4501746 [https://perma.cc/VP3F-2Y4Q] (citing and providing a meta-analysis of fourteen client alerts and podcasts published by major law firms active in the FDA regulatory area in response to the CDS Guidance and finding a surprising level of consensus that the Guidance is out of accord with Section 3060 of the Cures Act).

273. *See* Mark Seidenfeld, *Substituting Substantive for Procedural Review of Guidance Documents*, 90 TEX. L. REV. 331, 347 (2011) (noting that agency guidance documents have no binding legal force independent of the regulations they interpret or implement); *see also* Nina A. Mendelson, *Regulatory Beneficiaries and Informal Agency Policymaking*, 92 CORNELL L. REV. 397, 400 n.17 (2007) (noting that it is common practice for guidance documents to disclose their non-binding nature to make clear that they are non-binding/non-legislative rules versus binding/legislative); U.S. FOOD & DRUG ADMIN, *supra* note 112 (disclosing its non-binding nature).

274. *See supra* notes 271-73 and accompanying text.

275. *See supra* Sections III.A-D.

276. *See* Barbara J. Evans, *Congress' New Infrastructural Model of Medical Privacy*, 84 NOTRE DAME L. REV. 585, 594-95 (2009) (citing discussion at an FDA Public Workshop on

the first step. Clinics, hospitals, insurers, and other sources of clinical health data do not employ standardized record formats and they describe the same medical conditions differently.[277] The President's Council of Advisors on Science and Technology (P-CAST) is pessimistic that a standard medical record format can ever emerge in the United States.[278] Determining whether a patient had a specific health event, such as a heart attack, requires knowledge of the recordkeeping conventions at the source site.

Even when the Privacy Rule lets AI/ML software developers acquire data, large investments of skilled labor and capital are required to verify the incoming information, detect duplicative entries, and convert data into a common format for analysis.[279] Supplying data separately to multiple software developers requires them to make duplicative efforts and investment; moreover, it increases patients' exposure to privacy risks to have their data transferred to multiple recipients. As with other networked infrastructure (e.g., pipelines, telecommunication systems, and power grids), the efficient solution is to develop a common data infrastructure to reduce duplicative capital investment and promote consistent oversight and safeguards (in this case, for data privacy).[280]

Public health authorities (such as the FDA) and health oversight agencies (such as state agencies responsible for clinical safety and licensure of health care facilities and providers) are positioned to play an important role in the creation of shared data infrastructure for medical AI because the Privacy Rule allows covered entities to share PHI with them (and with contractors they might engage to develop and operate the shared infrastructure).[281] This might sound far-fetched, but this approach has already been employed in a different

---

postmarketing drug safety studies indicating that the United States, because of its fragmented structure for delivering and paying for health care, lacks infrastructure for creating the longitudinal health records such studies require, such that the first evidence of emerging drug safety problems often comes from other nations with national health systems that provide greater coordination of care and thus more complete health records tracking patients over time).

277. *See* PRESIDENT'S COUNCIL OF ADVISORS ON SCI. & TECH., EXEC. OFF. OF THE PRESIDENT, REPORT TO THE PRESIDENT: REALIZING THE FULL POTENTIAL OF HEALTH INFORMATION TECHNOLOGY TO IMPROVE HEALTHCARE FOR AMERICANS: THE PATH FORWARD 39 (2010).

278. *See id.* ("[A]ny attempt to create a national health IT ecosystem based on standardized record formats is doomed to failure . . . . With so many vested interests behind each historical system of recording health data, achieving a natural consolidation around one record format . . . would be difficult, if not impossible.").

279. *See* Evans, *Data Ownership*, *supra* note 41, at 90-92 (discussing steps required to transform raw data from treatment encounters into useful data resources for analysis).

280. *See* CHARLES F. PHILLIPS, JR., THE REGULATION OF PUBLIC UTILITIES 51-54 (3d ed. 1993) (discussing natural monopoly characteristics of networked infrastructure facilities).

281. *See supra* Table 1, Norm 15 (allowing unconsented disclosures to public health authorities and their contractors); *id.* Norm 17 (allowing unconsented data disclosures to health oversight agencies).

context. The Food and Drug Administration Amendments Act of 2007 (FDAAA) authorized the FDA to a nationwide data network, known as the Sentinel System, for use in postmarketing assessments and studies of the clinical safety of FDA-approved drugs.[282]

Congress called for the system to include data for twenty-five million patients by July 2010 and 100 million by July 2012, although both targets were quickly surpassed and 230.2 million patients had data in the system at least transiently during 2000-2021.[283] To enhance privacy protections, the FDA chose a federated (distributed) data architecture, which means people's PHI is not actually transferred to a centralized database and instead remains at the source locations (often a HIPAA-covered health system or insurer), with the various sites cooperating to convert data into a common format and perform localized studies to answer queries about drug safety.[284] Because Congress authorized the FDA to develop this system as a public health activity, unconsented uses and disclosures of data fall under the Privacy Rule's norm on disclosures to public health authorities and their contractors.[285]

Congress authorized the FDA to enter Collaborative Data Use Agreements (CDUAs) with private-sector entities, such as companies and academic researchers, allowing access to the Sentinel Data infrastructure to study a wide variety of drug safety topics, subject to statutorily required privacy protections.[286] To date, the FDA has not aggressively used these powers, perhaps reflecting a lack of staff resources to administer this role, which tasks the FDA with regulating

---

282. *See* 21 U.S.C. § 355(k)(3)-(4); *FDA Sentinel System's Coronavirus (COVID-19) Activities*, SENTINEL, https://www.sentinelinitiative.org/featured-topics/coronavirus-covid-19 [https://perma.cc/JM84-GRHY] (last visited Sept. 23, 2023) (describing the Sentinel data infrastructure, its uses, and its privacy policies). *See generally* FDA Amendments Act of 2007, Pub. L. No. 110-85, 121 Stat. 823 (codified as amended in scattered sections of 21 U.S.C.).

283. *See* 21 U.S.C. § 355(k)(3)(B)(ii) (setting the Congressional targets); *see also Sentinel Distributed Database (SDD) Statistics Summary: 2000-2021*, SENTINEL (July 5, 2022), https://www.sentinelinitiative.org/about/key-database-statistics#sentinel-distributed-database-sdd-statistics-summary-2000-2021 [https://perma.cc/6UV6-28Q5] (reporting that 365.1 million unique patient identifiers have had data in the system during the period 2000-2022 and noting that one person can have more than one unique patient identifiers if the person changes health plans).

284. *See How Sentinel Works: The Sentinel Common Data Model and Sentinel Distributed Database*, SENTINEL, https://www.sentinelinitiative.org/about/how-sentinel-gets-its-data#how-sentinel-works-the-sentinel-common-data-model-and-sentinel-distributed-database [https://perma.cc/XUM2-NCCG] (last visited Sept. 23, 2023).

285. *See supra* Table 1, Norm 15.

286. *See* RICHARD PLATT ET AL., MINI-SENTINEL AND CLINICAL TRIALS TRANSFORMATION INITIATIVE: DEVELOPING APPROACHES TO CONDUCTING RANDOMIZED TRIALS USING THE MINI-SENTINEL DISTRIBUTED DATABASE 36-41 (2014), https://www.sentinelinitiative.org/sites/default/files/Methods/Mini-Sentinel_Methods_CTTI_Developing-Approaches-to-Conducting-Randomized-Trials-Usi.pdf [https://perma.cc/P9LX-SFN7] (describing the powers Congress granted to FDA under 21 U.S.C. § 355(k)(4) to enter Collaborative Data Use Agreements and discussing the scope of studies Congress allows and the required privacy protections).

access to a national data infrastructure. The fact that the Sentinel System relies heavily on administrative data (e.g., insurance claims data) could have reduced the demand for access: administrative data suffer various data quality issues that limit their utility in scientific research and, moreover, they offer a biased sample of a fairly privileged subpopulation: people who have health insurance.[287]

The Sentinel System does not itself answer the need for a national data infrastructure for AI/ML CDS software. The administrative data (insurance claims data) it assembles are not an ideal resource for AI/ML research and, further, Congress authorized its use for drug-safety studies, whereas CDS tools are regulated as medical devices.[288] Nevertheless, it provides a model of what a national data infrastructure for AI/ML CDS software might look like and of the legislative provisions for creating it and triggering access to data under the Privacy Rule.

States seeking to gain a leadership position in AI-enabled health care and to ensure that their citizens are well-represented in AI/ML training data could follow the legislative model set out in FDAAA, appointing a state health oversight agency to act as the state's AI data infrastructure regulator. As such, the responsible state agency could receive data under the Privacy Rule's norm allowing covered entities to disclose data to health oversight agencies.[289]

IV. ACHIEVING STATE-OF-THE-ART PRIVACY PROTECTION IN MEDICAL AI

The Privacy Rule offers multiple legal pathways for assembling more inclusive, representative data sets for socially beneficial purposes, such as to make AI/ML CDS tools perform more equitably across the full range of demographic subgroups receiving treatment at American health care facilities. Why, then, is the current generation of AI/ML CDS tools performing so poorly in this regard? This Part starts by examining factors that have made data controllers reluctant to embrace the socially beneficial data sharing practices that the Privacy Rule sought to promote. It then identifies specific reforms that might help overcome this reluctance.

287. *See* Evans, *Seven Pillars*, *supra* note 191, at 483-85 (discussing advantages and limitations of Sentinel's reliance on administrative data).

288. *See* 21 U.S.C. § 355 (k)(3)-(4) (authorizing the Sentinel System for drug-related purposes); *see also supra* note 269 and accompanying text (explaining that Congress treats FDA-regulated CDS software as a medical device).

289. *See supra* Table 1, Norm 17.

A. Why the Privacy Rule Has Underperformed Its Original Promise

IRBs and other clinical data gatekeepers display ongoing reluctance to allow non-consensual data access and privacy by design (PBD), clinging to safe harbor de-identification and consent as their "go-to" privacy protections even though the Privacy Rule and other medical privacy laws do not require them to do so.[290] Even when policymakers strive to create vibrant scientific data commons, as with the National Institute of Health's Genomic Data Sharing Policy, they sometimes cling to safe harbor de-identification methods that can compromise data utility and hinder uptake of state-of-the-art computational privacy safeguards (PBD).[291]

IRBs—a creation of mid-twentieth-century bioethics—sometimes impede socially justified data access. A disturbing example occurred when IRBs blocked non-consensual data access under the Privacy Rule for an FDA drug safety activity, when the requested use was for a congressionally authorized public health purpose that clearly was not regulated by the Common Rule (which envisions IRB involvement).[292] Access was sought under the Privacy Rule's norm for disclosures to public health authorities, which does not involve IRBs as decisionmakers.[293] Yet private IRBs, without a clear legal basis, blocked the data use, which Congress determined would offer public benefit in the form of improved drug safety.

IRBs, as currently composed, generally lack information science and statistical qualifications required to oversee privacy protection plans that harness PBD.[294] The HHS Office for Human Research Protections (OHRP) administers the system of IRB ethics review

290. *See* Kaissis et al., *supra* note 200 (noting an ongoing, heavy reliance on "[a]nonymization (the removal of private data from a record) and pseudonymization (replacement of sensitive entries with artificially generated ones while still allowing re-attribution using a look-up table) . . . [which] are currently the most widely used privacy preservation techniques for medical datasets").

291. *See, e.g.*, Off. Director, Nat'l Inst. Health, *supra* note 165, at 2-3 (explaining that the NIH Genomic Data Protection Policy requires the use of safe harbor de-identification under the Privacy Rule, acknowledging that doing so can reduce data utility and provide sub-optimal privacy protection, and seeking comment on whether the policy should be amended to allow statistical de-identification under the Privacy Rule).

292. Sarah L. Cutrona et al., *Validation of Acute Myocardial Infarction in the Food and Drug Administration's Mini-Sentinel Program*, 22 PHARMACOEPIDEMIOL DRUG SAFETY 40, 44 (2013) (noting that hospital IRBs in seven cases required a patient signature to release charts even though the data request indicated that the data were for a congressionally approved FDA public health study that did not fall under the purview of the Common Rule and for which HIPAA allows unconsented release of data).

293. *See supra* Table 1, Norm 15; 45 C.F.R. § 164.512(b)(1) (2022).

294. *See* 45 C.F.R. § 46.107 (2022) (stating the qualifications for IRB members, without requiring any member of an IRB to have knowledge of privacy law, computer science, informational privacy, or probability).

under the Common Rule.[295] In 2011, OHRP noted that "questions have been raised about the extent and quality of the protections afforded by current informed consent requirements and practices."[296] OHRP questioned the wisdom of using IRBs to manage "informational risks" (i.e., privacy risks) that depend on "the nature of the information and the degree of identifiability of the information."[297] OHRP conceded that "[i]t is not clear that [IRB] members have appropriate expertise regarding data protections"[298] and concluded "[s]tandardized data protections, rather than IRB review, may be a more effective way to minimize informational risks."[299] This vote of no confidence in IRBs came from the regulator that created and administers the IRB system.

To reduce inappropriate privacy oversight by IRBs, OHRP revised the Common Rule in 2017 so that secondary uses of identifiable private information or identifiable biospecimens are now exempt from the Common Rule, if the uses are subject to the Privacy Rule's regulation of research, health care operations, or public health uses of data.[300] Thus, activities that use PHI are exempt from the Common Rule if the

295. *See* 45 C.F.R. § 46.103(a) (requiring entities regulated by the Common Rule to provide assurances of their compliance with the regulation to the "Office for Human Research Protections, HHS, or any successor office"); *see also id.* §§ 46.107-.109 (describing the requirements for IRB membership, IRB functions and operations, and IRB review of research under the Common Rule).

296. Human Subjects Research Protections: Enhancing Protections for Research Subjects and Reducing Burden, Delay, and Ambiguity for Investigators, 76 Fed. Reg. 44512, 44513 (proposed July 26, 2011) (to be codified at 45 C.F.R. pts. 46, 160, & 164).

297. *Id.* at 44516.

298. *Id.*

299. *Id.*

300. *See* 45 C.F.R. § 46.104(d)(4) (providing that "consent is not required" for secondary research use of "identifiable private information or identifiable biospecimens," and exempting such data uses from the Common Rule if "[t]he research involves only information collection and analysis involving the investigator's use of identifiable health information when that use is regulated under 45 CFR parts 160 and 164, subparts A and E [i.e., the HIPAA Privacy Rule], for the purposes of 'health care operations' or 'research' . . . [or] 'public health activities and purposes' "); *see also* Federal Policy for the Protection of Human Subjects, 82 Fed. Reg. 7149, 7192, 7194 (Jan. 19, 2017) (revising the Common Rule to create this exemption and explaining that there is "no requirement that the information and biospecimens must be pre-existing at the time that the investigator begins a particular research study" and that the study "could include specimens that are added . . . during the course of the study," and that this exemption "will allow investigators to see identifiable private information, and also allow them to retain and record that information (including the identifiers) as part of their research records"); *id.* at 7186, 7192-93 (explaining that, under the revised Common Rule, " 'exempt' does not always mean exempt from all of the requirements of the Common Rule" and weighing whether the scope of "protections required under HIPAA . . . were sufficient" to address research ethical concerns with informational research that uses identifiable private information). *But see id.* at 7194 (concluding, after considering public comments, that the 45 C.F.R. § 46.104(d)(4)(iii) exemption "introduces a clearer distinction between when the Common Rule and the HIPAA Privacy Rule apply to research in order to avoid duplication of regulatory burden. We believe that the HIPAA protections are adequate for this type of research, and that it is unduly burdensome and confusing to require applying the protections of both HIPAA and an additional set of protections").

data user is a HIPAA-covered entity, but this exemption does not apply if the data user is not HIPAA-covered.[301] This change took effect in January 2019.[302] The Common Rule no longer requires consent or IRB gatekeeping when data flows are destined for a HIPAA-regulated use.[303]

The Privacy Rule does not *require* IRB review before a data-holding covered entity shares data with another HIPAA-covered entity for research. The Privacy Rule does grant data-holding covered entities the *option* of using their IRBs, in lieu of creating a special-purpose HIPAA privacy board, to approve waivers of authorization for research at other HIPAA-covered facilities.[304] Other than that, IRBs have no role in administering the Privacy Rule.

Still, neither regulation expressly prevents a data-holding institution from involving its IRB, as a matter of institutional policy, in matters for which IRB oversight is not legally required.[305] IRBs regularly interfere and block data flows that the Privacy Rule sought to enable.[306] Doing so is questionable if it rests on the fallacy that IRBs somehow enhance privacy protection—a view OHRP rejected in 2011.[307] When Common Rule IRBs intrude into HIPAA privacy oversight, they venture outside their expertise, often requiring DOCG privacy schemes in lieu of state-of-the-art privacy protection.[308]

---

301. *See* Federal Policy for the Protection of Human Subjects, 82 Fed. Reg. at 7192 (stating that "[n]ot all investigators are part of a covered entity and thus some investigators are not required to comply with [the HIPAA Privacy and Security] rules" and thus would not be entitled to this exemption from the Common Rule).

302. 45 C.F.R. § 46.101(l)(2) (noting that the general compliance date for the revised Common Rule was January 21, 2019).

303. The HIPAA Privacy Rule does not use IRBs, although it does give covered entities the option of using their existing Common Rule IRB (instead of creating a special-purpose HIPAA privacy board) for one task only: approving waivers of authorization for research uses of data. *See supra* Table 1, Norm 13; 45 C.F.R. § 164.512(i) (2022); *see also supra* note 282 (discussing and citing the provision of the revised Common Rule removing HIPAA-regulated activities from Common Rule oversight).

304. *See supra* Table 1, Norm 13.

305. *See* 45 C.F.R. § 46.112 (allowing research institutions to impose stricter ethics review procedures than the Common Rule requires). The Privacy Rule achieves similar effect in its 26 informational norms that allow (but do not require) data disclosures, leaving covered entities free to add additional conditions at their discretion before agreeing to disclose data. *See supra* Table 1, Norms 1-2, 4-27.

306. *See, e.g.*, Cutrona et al., *supra* note 292; *see also* IOM, PRIVACY REPORT, *supra* note 5, at 70-71 (drawing a conclusion, after studying researchers' access to data, that "[t]here is a great deal of variation[,] . . . with many covered entities, IRBs, and Privacy Boards interpreting the HIPAA Privacy Rule very conservatively" and noting that "[t]hese interpretations impede some important research activities, and can also limit the validity and generalizability of some research results").

307. *See supra* notes 295-99 and accompanying text.

308. *See* 45 C.F.R. § 46.107 (stating undemanding qualifications for membership on an IRB). Because the HIPAA Privacy Rule largely rejects the use of IRBs, it states no qualifications for IRBs.

Data-holding institutions whose institutional policies voluntarily insert IRBs in privacy oversight have a duty to ensure the IRBs are suitably skilled for that task. If institutions shirk this duty and continue staffing their IRBs with volunteers lacking state-of-the-art skills for privacy oversight, their states should intervene. The Common Rule and Privacy Rule are both "floor" regulations stating minimal federal standards and allowing states to set stricter standards of protection.[309] States should require any IRB that oversees informational research or data disclosures under the Privacy Rule to have majorities composed of data security and privacy experts knowledgeable of computer science and engineering, PBD, and probability and statistics (including statistical bias).

The Common Rule and HIPAA preemption provisions would overlay these stricter state IRB staffing requirements onto the weak federal standards.[310] People with the needed skills are in demand and may be unlikely to volunteer. This could increase data-holding institutions' costs to staff their IRBs, possibly deterring the institutions from involving IRBs in privacy oversight for which they lack expertise. That would not necessarily be a bad outcome, unless one prefers "dummy thermostats" to real privacy protections.[311]

In fairness to IRBs, on the whole they are staffed by well-intentioned people volunteering their time out of a sincere desire to protect the rights people whose data are held in biomedical information systems. Their reluctance to employ the Privacy Rule's norms for unconsented data access is driven, at least in part, by justified concerns about gaps in its alternative privacy protections.[312] The next Section explores reforms to help strengthen privacy protections for data used as AI/ML training data pursuant to the nonconsensual access pathways Part III identified.

B. Addressing the Privacy Rule's Lingering Privacy Gaps

The rulemaking process that created the Privacy Rule was attentive to federalism concerns and included extensive consultations with the States under Executive Orders No. 12,612 and No. 13,132, both

309. *See* Common Rule, 45 C.F.R. § 46.101(f) ("This policy does not affect any state or local laws or regulations . . . that provide additional protections for human subjects."); Privacy Rule, 45 C.F.R. § 160.203(b) (2022) (preventing the preemption of state laws that provide more-stringent privacy protections than the HIPAA Privacy Rule does). *See generally* William W. Buzbee, *Asymmetrical Regulation: Risk, Preemption, and the Floor/Ceiling Distinction*, 82 NYU L. REV. 1547, 1551-52, 1554 (2007) (discussing federal regulations that establish floors from those that set regulatory ceilings).

310. *See* 45 C.F.R. §§ 160.202-.203 (Privacy Rule preemption provisions); *see also* 45 C.F.R. § 46.101(f) (deferring to state and local laws and regulations that provide stricter protection than the Common Rule does).

311. *See* Sandberg, *supra* note 206 (discussing non-functioning "dummy" thermostats).

312. *See supra* Section I.B (summarizing the Privacy Rule's alternative protections).

addressing federalism.[313] States were adamant about retaining the power to implement stronger medical privacy protections than the Privacy Rule provides.[314] This allocation of power appears in the Privacy Rule's preemption provisions, which defer to more-stringent privacy requirements of state law.[315] This puts states in a position to address gaps in the alternative privacy protections the Privacy Rule provides when data are shared without individual authorization. The states should consider the following measures:

Strengthen information fiduciary duties for controllers of AI/ML CDS tools. Parties who control AI/ML CDS tools—whether they are researchers, software developers/vendors, or health care providers that use the software—should be subject to strong information fiduciary duties in their handling of personal information used in these systems. For data used in AI/ML software without individual authorization under the Privacy Rule's informational norms, these duties should include (at a minimum) restrictions on reuse, re-identification, and redisclosure of the data.

States could implement these requirements in various ways. One approach would be through contractual privacy protections establishing information fiduciary duties: (1) as state-mandated terms to be included in covered entities' BAAs (for disclosures of PHI to business associates); (2) through state requirements for covered entities to use DUAs for all unconsented data disclosures, even when the Privacy Rule does not require a DUA; and (3) as state-required terms to be included in those DUAs—for example, limitations on reuse, re-identification, and redisclosure of the data and restrictions on the diversion of PHI to other uses by de-identifying it using the Privacy Rule's lax safe harbor method.

An additional option to consider would be for states to treat data subjects (the people whose PHI is used in AI/ML CDS tools) as third-party beneficiaries of BAAs, DUAs, and software vendor contracts. This would promote accountability by granting data subjects a private right of action. HHS considered a private right of action to be essential for promoting accountability but felt it lacked authority to include one in the Privacy Rule, leaving the matter for states to resolve.[316]

313. *See* Evans, *Institutional Competence*, *supra* note 4, at 1212-13 (discussing HHS's consultations with the states, in compliance with these Executive Orders, during the rule-making process for the Privacy Rule).

314. *See id.* at 1213.

315. *See* 45 C.F.R. §§ 160.202-.203 (preemption provisions).

316. *See supra* notes 53, 71 (quoting statements by HHS, in the preamble to the proposed HIPAA Privacy Rule, complaining about the insufficiency of HHS's jurisdiction under HIPAA to regulate medical privacy and calling on Congress to pass legislation creating a private right of action for data subjects in order to ensure covered entities will be accountable for privacy protection—something that HHS felt it could not do via regulation).

An alternative approach would avoid reliance on contract law altogether. Under this approach, states could amend their medical records acts and common law, so that *all* parties receiving PHI from HIPAA-covered entities are subject to the same informational fiduciary duties that apply to health care providers.

Nudging the medical software industry toward modern computational privacy protections. The Privacy Rule allows safe harbor or statistical de-identification. As already noted, safe harbor de-identification continues in wide use because it is straightforward and easy to understand, despite concerns about its effectiveness.[317] States should impose more-stringent requirements that encourage covered entities and recipients of PHI to transition away from the weak safe harbor method and toward state-of-the-art computational methods for protecting privacy.

Because safe harbor de-identification is widely used at present, these state requirements should provide flexibility and a reasonable transition period. For example, states might continue to allow safe harbor de-identification but impose a more-stringent requirement for covered entities and data users that use it to perform a privacy impact assessment that includes a numerical estimate of re-identification risk. Such a requirement would allow safe harbor de-identification but nudge those who rely on it to start thinking in statistical terms about privacy risks. Over time, states could implement progressive requirements for users who receive or use PHI in AI/ML CDS tools to transition to statistical de-identification/privacy by design.

As computational techniques for privacy protection continue to evolve, states should avoid prescribing specific techniques (e.g., federated learning, differential privacy, etc.). Instead, they should state broad, general requirements for AI/ML CDS software to implement privacy protections that provide measurable guarantees against a broader range of attacks than are addressed by the Privacy Rule (which, like many other regulations, conceives re-identification attacks as the principal threat to privacy).[318]

---

317. *See* Kaissis et al., *supra* note 200, at 307-08 (noting an ongoing, heavy reliance on safe harbor de-identification).

318. *See* Nissim & Wood, *supra* note 24, at 11 ("Privacy regulations and related guidance contemplate a limited set of specific attacks and privacy failure modes. As one example, many regulations make an implicit assumption that re-identification via record linkage—i.e. the re-identification of one or more records in a deidentified dataset by uniquely linking these records with identified records in a publicly available dataset—is the primary or sole privacy failure mode. Other central concepts appearing in privacy regulations, including personally identifiable information, (de-)identification, linkage and inference, are often defined from this point of view. For example, many privacy regulations require data providers to protect information that can be linked to an individual in order to safeguard against record linkage. As a result, these requirements are often interpreted as requiring the protection of information one can foresee being used in a record linkage attack. However, in the last two decades, researchers have identified new attacks and privacy failure modes."). *See generally*

These state requirements should recognize that there are tradeoffs between privacy protection, on the one hand, and data utility, computational efficiency, and accuracy, on the other hand.[319] Therefore, states should provide a procedure and a set of criteria for review and approval of privacy plans involving such tradeoffs.

Staffing requirements for IRBs involved in data access decisions. When HIPAA-covered institutions choose to involve their IRBs in administering Privacy Rule access to data for use in AI/ML medical software, states should require the IRBs to be staffed with majorities of qualified information privacy specialists. Institutions whose IRBs are only involved in the oversight of traditional clinical research would not be subject to these requirements.

Truthful disclosure of the limits of consent and risks of nonparticipation. Five decades of exposure to twentieth-century bioethics and its control-over-information privacy theory have convinced many Americans—including people who serve on IRBs and individual members of the public—that safe harbor de-identification and consent/authorization provide effective privacy protections.[320] IRBs continue to require, and members of the public continue to desire, DOGC privacy norms. We, as a society, perform a bioethical belief that informed consent is an ethical imperative that protects privacy and promotes public trust. This belief persists despite mounting evidence that consent does not protect privacy and that it introduces consent biases, including invidious ones.[321]

This situation points to the need for education and re-messaging about informed consent, at least for data used in AI/ML CDS software. The consent-as-altruism narrative of twentieth-century bioethics breaks down for AI/ML CDS tools, where medical benefits are largely internalized to consenters.[322] The better narrative is that non-consent is an act of self-harm: being left out of training data can cause medical software to deliver ill-informed recommendations for you and for people medically and demographically similar to you.[323] Nonconsensual data use may be ethically justified if it promotes inclusivity in AI/ML training data, reduces future health disparities, and promotes health equity and social justice. Bioethics has long highlighted the risks of research participation during the informed consent

Wan et al., *supra* note 215, at 431 fig.1 (providing an overview of various types of privacy attacks); Kaissis et al., *supra* note 200, at 306 tbl.1 (same).

319. *See* Wan et al., *supra* note 215, at 436-37 (noting that some technical safeguards for privacy cause losses of accuracy and efficiency).

320. *See* IOM, PRIVACY REPORT, *supra* note 5, at 66 (summarizing surveys of public attitudes about de-identification and informed consent for data uses).

321. *See supra* note 200 (citing various studies that have questioned whether consent is effective as a privacy protection); *see also supra* Section II.A (discussing disparate impacts of consent norms).

322. *See supra* Section II.B.

323. *Id.*

process for research uses of health information. When seeking consent to use data to train AI/ML CDS tools, there are potential risks of *not* participating that seemingly ought to be disclosed.

Disclosing the impacts of consent bias would not violate bioethical norms against coercion[324] because those norms (like all bioethical norms) are directed at humans—in this case, the humans conducting AI/ML research. According to these norms, *researchers* must not create circumstances that pressure people to consent to research.[325] Consent bias is not something researchers create (even if some other forms of bias, such as algorithmic biases and systemic biases, are human creations to which researchers sometimes do contribute). Consent bias is a statistical reality caused by *research participants'* individual decisions to consent or not consent. Those individual decisions have consequences for people's future health. Disclosing a risk in order to help people make well-informed decisions is not the same thing as intentionally presenting a harm "in order to obtain compliance" (i.e., coercion).[326] This, admittedly, is a fine line deserving careful bioethical guidance on how IRBs and the researchers they oversee can best navigate it. Agencies that fund or regulate research, institutions that hold data or conduct AI research, bioethicists, and IRBs should develop policies on disclosure of the risks of non-participation in data sets used in training AI/ML CDS tools.

Public benefit criteria for unconsented uses of data in AI/ML research. States should require public benefit criteria for AI/ML software that relies on the Privacy Rule's waiver provision as a means of acquiring PHI to use as training data. The public benefit criteria would be a more-stringent state-level add-on to the criteria the Privacy Rule already requires for approval of waivers.[327] The goal is to implement the National Commission's 1978 recommendation that non-consensual research uses of identifiable (or re-identifiable) health data are ethically justified only if "the importance of the research or statistical purpose for which any use of disclosure is to be made is such as to warrant the risk to the individual from additional exposure of the

324. *See, e.g.*, 45 C.F.R. § 46.116 (2022) (requiring that "[a]n investigator shall seek informed consent only under circumstances that provide the prospective subject or the legally authorized representative sufficient opportunity to discuss and consider whether or not to participate and that minimize the possibility of coercion or undue influence").

325. *See Informed Consent FAQs*, U.S. DEP'T HEALTH & HUM. SERVICES, https://www.hhs.gov/ohrp/regulations-and-policy/guidance/faq/informed-consent/index.html [https://perma.cc/S32Y-DLM2] (last visited Sept. 23, 2023) (explaining that "[c]oercion occurs when an overt or implicit threat of harm is intentionally presented by one person to another in order to obtain compliance").

326. *Id.*

327. *See supra* Table 1, Norm 13; *see also* 45 C.F.R. § 164.512(i)(2)(ii) (2022) (listing criteria that an IRB or privacy board must determine have been met before it can approve a waiver or alteration of the Privacy Rule's authorization requirements).

record or information contained therein."[328] The procedures for making this determination—and what, substantively, it means for research to be "socially beneficial" so as to confer a public benefit—will be challenging to define.

Deciding which lines of research offer enough social benefit to warrant the use of waivers requires a global science policy perspective and is not a suitable question for local IRBs. One approach would be for legislatures (or a publicly accountable oversight body) to identify general categories of AI/ML medical research that offer sufficient public benefit to justify unconsented access to data. For example, the use of waivers might be allowed only in projects that aim to achieve inclusivity across demographic, gender, racial, and/or socioeconomic lines. Research that appears unlikely to produce broadly generalizable results would not be eligible to acquire data using waivers and instead would have to rely on the Privacy Rule's other informational norms (such as obtaining individual authorizations for the research).

Another public benefit criterion might be to insist that any CDS tools developed by the research comply with certain standards of business transparency and accountability. Such standards could require transparent business practices (for example, no gag clauses in software vendor contracts, so that health care providers can freely air problems they encounter while using the software[329]); reasonable pricing; requirements to validate the software's performance on diverse patient populations and to clearly disclose if the software is not fit for purpose in all subgroups; and standards of algorithmic transparency/explainability to help health care providers understand and challenge the software's recommendations.

A public benefit standard for unconsented research uses of data is conceptually similar to the "public use" requirement in takings law.[330] In takings, the state has eminent domain power to take property for "public use" without the owner's consent, subject to payment of just compensation.[331] The waiver provisions of the Common Rule and Privacy Rule resemble a scheme of private eminent domain, in which people's data can be "taken" at the discretion of private-sector bodies

328. *See* Protection of Human Subjects: Institutional Review Board; Report and Recommendations of the National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research, 43 Fed. Reg. 56174, 56181 (Nov. 30, 1978) (quoting the Privacy Protection Study Commission); *see also* BELMONT REPORT, *supra* note 15, pt. C.2.

329. *See* Hawkins et al., *supra* note 234, at 273, 275 (discussing gag clauses and other nontransparent business practices in electronic health record vendor contracts).

330. *See* Robin Paul Malloy & James Charles Smith, *Private Property, Community Development, and Eminent Domain*, *in* PRIVATE PROPERTY, COMMUNITY DEVELOPMENT, AND EMINENT DOMAIN 1, 8 (Robin Paul Malloy ed., 2008) (discussing the public use requirement).

331. U.S. CONST. amend. V.

(IRBs and HIPAA privacy boards).[332] Patients subjectively sense a "taking" when their data are used in research pursuant to waivers, and some express a desire to share in proceeds that ultimately may flow from the research.[333] Unfortunately, the nature of "big data" health research (and most AI/ML software) is that it gleans insights by processing very large data sets incorporating many people's data, such that the contribution of any individual would probably be valued in pennies. Moreover, takings law compensates fair market value, not subjective values such as the dignitary insult of having one's consent right waived by an IRB.[334] What people want—and deserve—when their data are used without consent is an assurance that their sacrifice serves a socially beneficial purpose.

In discussing takings of property, Professor Merrill suggests an analytical approach that identifies traits that signal when a taking *does not* serve a public purpose.[335] A similar approach might be helpful as states develop public benefit criteria for unconsented data uses. It can be easier to define what a publicly beneficial use of data *is not* than to define what it *is*. For example, researchers unwilling to disclose their research objectives, methods, and results might be presumed to have private, commercial aims, as opposed to advancing socially beneficial medical knowledge. Data uses that advance private aims should be allowed, but they should not be eligible to acquire people's data through waivers. Persons desiring to use the public's PHI without consent must be open about what they plan to do with it. States could develop a list of red flags that weaken the presumption that a proposed research use of data offers sufficient public benefit to justify the use of waivers. There may be sound ethical justifications for unconsented uses of people's PHI, but such uses need to have guardrails that deter frivolous, scientifically unjustified uses and private rent-seeking behaviors.

Legal accountability for inequitable AI/ML medical software. In its 1997 recommendations to Congress on health data privacy, HHS stressed that an accountable system of privacy protections must include a private right of action enabling individuals to bring lawsuits

332. *See* Evans, *Data Ownership*, *supra* note 41, at 77-82 (exploring this analogy and concluding that even if patients had property rights to "own" their data, the data still would be subject to unconsented use under eminent domain principles).

333. *See, e.g.*, Moore v. Regents of the Univ. of Cal., 793 P.2d 479, 480 (Cal. 1990).

334. *See* Evans, *Data Ownership*, *supra* note 41, at 81-82 (concluding that, if people owned their data, the "just compensation" for an unconsented data use would likely be zero under current takings doctrine, which does not compensate subjective value, such as an owner's emotional attachment to the property).

335. *See* Thomas W. Merrill, *The Economics of Public Use*, 72 CORNELL L. REV. 61, 90-92 (1986).

to enforce their civil rights.[336] The same is true more generally to protect people's safety and civil rights in an age of AI-enabled health care. The subject of liability is a hated one in medical circles, but the fact remains that tort lawsuits are an important legal pathway for incentivizing safety, equity, and privacy if a system is serious about promoting accountability. A recent GAO report noted the liability landscape surrounding medical AI is still developing, and it is hard to say how liability may ultimately be apportioned between software developers and vendors, on the one hand, and health care facilities and professionals that use software systems, on the other hand.[337]

The FDA's involvement in regulating CDS software as a medical device has the effect of classifying it as a medical product, rather than as an information service. This opens the door to product liability suits when medical software contributes to patient injuries.[338] Design defect suits under state law offer a possible mechanism for promoting inclusivity and equity in AI/ML medical tools: it is a design defect when software purports to be for general clinical use but was developed with training data that fails to reflect all the patients who will be relying on it.[339] States should consider recognizing product liability causes of action for design defects and also for failure-to-warn in situations where patients are injured by AI/ML CDS tools that used non-inclusive training data or that failed to disclose that the software is not fit for purpose in underrepresented population subgroups.[340] Not all AI/ML CDS tools will be subject to FDA regulation, and the FDA's proposed regulatory approaches do not appear likely to preempt state product liability suits.[341] The state tort system thus could provide a more consistent, generally applicable framework to incentivize good practices: a framework that would apply whether or not the software receives FDA oversight as a medical device.

The policies discussed above aim to fill lingering gaps in the Privacy Rule's protections. They address Professor Balkin's call for information fiduciary requirements for controllers of AI/ML algorithms.[342] They implement the National Commission's recommendation that ethical use of people's data in research without consent

336. *See HHS, 1997 Recommendations*, *supra* note 38, § I.H; *see also supra* note 66 (quoting HHS's statements, in the preamble to the proposed HIPAA Privacy Rule, calling for legislative action to create a private right of action).

337. *See* U.S. GOV'T ACCOUNTABILITY OFF., *supra* note 121, at 30.

338. *See generally* Barbara J. Evans & Frank Pasquale, *Product Liability Suits for FDA-Regulated AI/ML Software*, *in* THE FUTURE OF MEDICAL DEVICE REGULATION: INNOVATION AND PROTECTION 22 (I. Glenn Cohen et al. eds., 2022) (proposing product liability causes of action to address bias in AI/ML medical software).

339. *See id.* at 32-34.

340. *See id.*

341. *See id.* at 26-27, 30.

342. *See* Balkin, *Three Laws*, *supra* note 49, at 1227.

should be subject to public benefit requirements.[343] They promote accountability by implementing HHS's 1997 recommendation to Congress that private rights of action are a necessary part of meaningful privacy protections.[344] Ideally, states should collaborate to develop a uniform model framework of more-stringent privacy requirements to strengthen the Privacy Rule's protections. Interstate coordination can reduce compliance burdens that come with an inconsistent patchwork of state privacy requirements.

CONCLUSION

The Privacy Rule was a prophecy that the control-over-information privacy theory popular after 1970 might cease to perform well on the altered landscape of twenty-first-century medicine, as a diverse patient population struggles for equity in a health system increasingly dependent on informational science as a source of medical truth. That truth is blurry and unreliable for those not included in the AI/ML training data from which truth is—increasingly—gleaned. Protecting privacy by empowering people to hide their data may, as an unintended consequence, expose them to deadly risks.

A flaw in the control-over-information theory is that it treats data privacy as something individuals can protect for themselves by limiting access to their personal data. In AI-enabled research and health care, patients and research participants are thrust into a dizzying mix of roles as data subjects, patients/participants, and ultimate users of unfamiliar technologies. The imbalances of power and control they face are so vast that consent, even when it feels ethically necessary, may not be sufficient to protect their interests. Software developers and other key players often are not subject to the medical ethics and state-law fiduciary duties that bind traditional health care providers to an ethic of responsible data handling. Going forward, law needs to place more of the burden of protecting people's privacy on those who design, implement, use, and control medical AI systems.

This Article explored ways the Privacy Rule's informational norms could promote greater inclusivity in AI/ML training data. The nagging question is whether data inclusivity is worth it, if it jeopardizes everybody's privacy. The Privacy Rule's informational norms offer pathways for reducing consent bias and promoting greater equity in AI/ML training data. To be ethically acceptable, however, these pathways should be pursued only in conjunction with policies that strengthen the alterative protections the Privacy Rule currently prescribes when data are used without consent. This Article recommended several policies to

343. *See* Protection of Human Subjects: Institutional Review Board; Report and Recommendations of the National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research, 43 Fed. Reg. 56174, 56181 (Nov. 30, 1978).

344. *See supra* note 318 and accompanying text.

address the Privacy Rule's current gaps. These recommendations were not an exhaustive list but mere examples to stimulate further discussion among scholars and policymakers.

To conclude, an important caveat must always be borne in mind: acquiring training data without consent can reduce consent bias, but consent bias is only one type of bias. The larger challenge—always—is to tackle the systemic/structural biases that pervade the U.S. health care system and infect real-world data drawn from it. Norms for acquiring data for AI/ML software—whether with or without consent—cannot eliminate systemic biases already stamped upon the data being acquired. This latter challenge demands work not just from privacy scholars but from all of us.