Protecting health privacy even when privacy is lost

T J Kasperbauer

Center for Bioethics, Indiana University School of Medicine, Indianapolis, Indiana, USA

Correspondence to

Dr T J Kasperbauer, Center for Bioethics, Indiana University School of Medicine, Indianapolis, IN 46202, USA; tkasperb@iu.edu

Received 2 October 2019
Revised 17 November 2019
Accepted 20 November 2019
Published Online First
5 December 2019

ABSTRACT

The standard approach to protecting privacy in healthcare aims to control access to personal information. We cannot regain control of information after it has been shared, so we must restrict access from the start. This 'control' conception of privacy conflicts with data-intensive initiatives like precision medicine and learning health systems, as they require patients to give up significant control of their information. Without adequate alternatives to the control-based approach, such data-intensive programmes appear to require a loss of privacy. This paper argues that the control view of privacy is shortsighted and overlooks important ways to protect health information even when widely shared. To prepare for a world where we no longer control our data, we must pursue three alternative strategies: obfuscate health data, penalise the misuse of health data and improve transparency around who shares our data and for what purposes. Prioritising these strategies is necessary when health data are widely shared both within and outside of the health system.

INTRODUCTION

The predominant conception of privacy in healthcare focuses on controlling and limiting access to one's personal information. On this view, privacy is something that cannot be reobtained after it is lost. Once personal information has been shared, there is no getting it back, and no way to regain control of the inferences people might make from that information.

For example, advocates of 'privacy by design' in healthcare focus on privacy as 'an individual's ability to exercise control over the collection, use, disclosure and retention of his or her personal information, including personal health information'.¹ Similarly, the Nuffield Council of Bioethics' report on data ethics claims that 'respect for persons' requires 'recognition of a person's profound moral interest in controlling others' access to and disclosure of information relating to them held in circumstances they regard as confidential'.² This conception of privacy is also featured in Beauchamp and Childress's classic introduction to bioethics.³ On all these accounts, controlling the flow of one's data is the primary tool for protecting privacy in healthcare.

When taken seriously, this 'control' conception of privacy threatens to undermine many promising initiatives in healthcare. For instance, one could take this to mean that patients should be asked to consent each time their data are shared within the health system.⁴ This would impose a significant barrier on initiatives like learning health systems and precision medicine, which require significant amounts of patient data from both within and outside the health system.^{5–7} There are also implications for programmes like the National Institutes of Health's All of Us initiative in the USA, which

aims to build a national research cohort consisting of a wide range of more than one million people's data.⁸ This includes their biospecimens, electronic health record, mobile health data and potentially much else, all of which will be shared for research purposes for years to come. The control conception would seem to conflict with these practices.

This control view of privacy is shortsighted and overlooks important ways to protect information even when it is widely shared. This paper outlines three ways of protecting health data,¹ such that pervasive data sharing would not automatically entail a loss of privacy. They include data obfuscation, penalising data misuse and data transparency. These strategies are not new; many privacy advocates encourage their usage. However, they are rarely described in a healthcare context as providing a potential solution to problems raised by pervasive data sharing.

In broad terms, here are the characteristic features of each of these strategiesⁱⁱ:

Control: I cannot get your data without your permission.

Obfuscation: I can get your data, but I cannot make meaningful inferences from it.

Penalization: I can get your data, but using it against you is likely to get me into trouble.

Transparency: I have your data, but you and everyone else will easily see if I misuse the data.

Obfuscation, penalisation, and transparency are typically seen as secondary to controlling and limiting access to personal data. In data-intensive healthcare, however, they will need to take priority. Each of these three strategies also entail certain drawbacks, as well as ethical and policy trade-offs that have not been fully recognised. While we must prioritise these strategies to protect health data going forward, policy makers will also need to grapple with the implications of their implementation.

WHY CONTROL IS INADEQUATE

The control view of privacy is not just a threat to data-intensive initiatives. It is also ultimately inadequate for protecting privacy. I will briefly outline two reasons to think that the control conception is inadequate before entering into a fuller discussion of better alternatives.

ⁱBy 'health data', I mean information that is regularly used to make inferences about individuals' health. This includes information that is stored and collected by healthcare providers, as well as information collected by third parties outside the health system.

ⁱⁱConsider 'I' to be a third party with an interest in health data and 'you' to be the patient or the person whom the data are about.

© Author(s) (or their employer(s)) 2020. No commercial re-use. See rights and permissions. Published by BMJ.

To cite: Kasperbauer TJ. *J Med Ethics* 2020;**46**:768–772.

The first reason is that so many incentives exist for people to share their data that they will likely share regardless of the risks. Patients seem to understand and accept that sharing their data can lead to health benefits. A recent meta-analysis found that people are largely willing to give up control of their health-related data in exchange for such benefits.⁹ They are even more likely to do so when there is a clear consequence for their personal health. A survey of over 2000 people with a medical condition found that nearly all would be willing to share social media information with their doctors (94%), researchers (92%) and drug companies (84%) if it would improve their care.¹⁰

Patients also tend not to be willing to pay more for increased privacy. Trachtenberg *et al* found that 95% of patients surveyed (n=834) would not restrict the information they share with providers, even for sensitive conditions like HIV, and that if given the choice, 68% of patients would put money toward reducing medical costs rather than improving privacy in health-care.¹¹ Even if we insist on a strong control regime for health data, it seems unlikely that patients will actually restrict access in a meaningful way.

Second, there is already so much personal health data outside of the health system that people have already lost significant control. Health information is widely shared outside of the protections provided by health-related privacy policies.¹² Healthcare providers as well as insurance companies routinely sell anonymised data from medical records to third parties, including professional 'data brokers'.¹³

This information is also sometimes taken directly from people with medical conditions. Huesch found that 13 popular health websites (eg, WebMD) used tracking software, and seven of those sites allowed people's search terms to be shared with third parties.¹⁴ Over 70% of mobile health apps have also been found to regularly share data with third parties,¹⁵ and 19 of the 24 most popular mobile health apps were recently reported to share data with companies like Amazon, Facebook and Google.¹⁶ This is setting aside all the inferences that could be made about health based on people's genetic data, which are also widely available. In short, it is too late to rely on control as the primary tool for protecting health data.

ROUTES FOR PROTECTING PRIVACY

Regardless of whether you agree with the limitations of the control conception of privacy just outlined, there are good reasons to pursue alternatives. Let us look at the alternatives as well as the ethical and policy trade-offs involved with their implementation.

Data obfuscation

The first route to protecting health privacy is to make data obscure and thereby useless for making inferences about any particular individual. As Brunton and Nissenbaum define this strategy, 'Obfuscation...is the production of noise...in order to make a collection of data more ambiguous, confusing, harder to exploit, and therefore less valuable'.¹⁷ The basic idea behind obscurity and obfuscation is to make personal information hard to interpret. Obscurity is a helpful strategy for protecting privacy because it imposes transaction costs on people who want to make meaningful inferences from personal information.¹⁸ Someone who wants to use my information against me has to overcome obfuscation methods, which will be a deterrent to many adversaries.

Anonymisation is probably the best-known example of obfuscation in healthcare. My personal disease history could be shared all across the health system, but so long as all identifying

information is removed, it is difficult to trace it back to me. I have lost control of the information, but it is hard for people to know that is *my* information.

However, standard approaches to deidentification that simply remove identifying information are often inadequate for the reasons mentioned earlier. Sweeney, for example, was able to reidentify patients by comparing newspaper stories about hospitalisations (eg, due to accidents) to publicly available deidentified health records. Forty-three per cent of the news stories she analysed could be connected to an individual patient record.¹⁹

More sophisticated approaches modify patient data such that individuals' information is obscure, but the overall value of the dataset is not lost.²⁰ One especially promising option is to create synthetic databases, where each patient's data are thoroughly modified without sacrificing the statistical properties of the overall dataset.²¹ These databases are complex and can be expensive, but they retain both patient privacy and the research value of the original database. For example, over a million synthetically generated patient health records are available online through Synthea, enabling public health research without sacrificing privacy.^{22 23}

Synthetic databases create noise within the data itself. However, another option is to create noise when accessing or using a dataset.²⁴ For example, the programme Aircloak allows people to ask questions about data in a dataset but introduces noise to the answers to protect private information.²⁵ This leaves the data itself intact, while preventing others from readily making inferences about individuals. For example, if a researcher asks, 'How many people in Polk County have both diabetes and Alzheimer's disease?' the programme provides an answer that takes into account the uniqueness of those in the dataset who have both conditions. The programme produces a random value if it determines that an accurate number would be too revealing (ie, uniquely identifying).

Trade-offs and drawbacks to data obfuscation

There are a number of challenges with implementing these obfuscation methods. The main challenge is that they make the data less useful and less beneficial for patients. It is hard to share data to improve patient care if you no longer know whose data these are. Researchers will also sometimes need individually identifiable data to conduct the right analyses. Programmes like Aircloak also limit the types of questions that providers and researchers can ask with any hope of receiving accurate information.

Patients might also object to how their data are used even when thoroughly obscure. For example, suppose that a synthetic database I have contributed to is used to analyse opioid addiction in my neighbourhood. Further suppose that policy makers looking at the data conclude that a cap should be placed on certain prescription medicines for everyone in the neighbourhood, including me. The public health intervention is more attuned to the group that is negatively affected by opioids, of which I am a member, than whether I personally am addicted to opioids. While intuitions likely differ here, it is not unreasonable to feel that my privacy has been violated even though nobody could possibly identify me in the dataset.²⁶ ⁱⁱⁱ If I am in fact addicted to opioids and wanted that

ⁱⁱⁱ As Barocas and Nissenbaum describe this problem, 'Even when individuals are not 'identifiable', they may still be 'reachable', may still be comprehensibly represented in records that detail their attributes and activities, and may be subject to consequential inferences and predictions taken on that basis'.

to remain secret, the dataset was used to learn things about me that were meant to be private.

Despite these challenges, obfuscation is still necessary for protecting privacy when health data are widely distributed. While obfuscation reduces the value of certain datasets, that is still better than no data, which is a likely consequence of control-based methods. Recognising the challenges and associated trade-offs can also be helpful when deciding how to implement obfuscation methods. For example, precision health programmes could assure patients that these methods will be used whenever data are shared outside of the health system. While identifiable information would be required for targeted treatments, any information sent to researchers or other third parties would be obscure, thereby protecting patient privacy.

Penalising data misuse

The second route is to criminalise and create stronger civil penalties for exploiting and misusing personal data. If data sharing cannot be controlled, it must be made less harmful. As we lose control of our data, we will need better tools to defend against and penalise those who use our data against us.

Currently, most laws aim to prevent access to health data, and there is very little that can be done if access restrictions are breached. Typically, privacy violations are not considered to cause harm under the law, especially in the USA. For example, increasing the risk of financial injury or increasing anxiety from the release of personal information is not sufficient to count as harm.^{27 28} This reduces the penalties that violators could receive as well as the remuneration individuals could obtain when their information is stolen or used against them. Solove and Citron review numerous cases where leaked data were used to steal people's identity, including financial information, but courts decided no harm was caused because there was no 'imminent threat of financial injury'.²⁹ Despite illegal acts leading to the loss of privacy, there was very little victims could do to prevent future misuses of their information.

An essential first step in penalising privacy violations is thus to establish that they can cause harm. For example, failing to take adequate measures to prevent a data breach could perhaps be viewed as a harm. Even if no damage occurs at the time of the breach, the breach increases the chance of harm in the future. Calo argues that courts have applied an especially high standard when determining whether privacy violations caused harm (compared with violations like assault).³⁰ Laws that more clearly define privacy harms to include emotional distress or increased risk of financial injury would lower the threshold, thereby making it easier to hold people accountable when data are misused. To take another example, we could consider it harmful to sell health data to entities that are known to engage in health insurance fraud. The act of selling does not directly injure patients, but it does significantly increase their exposure to various risks.

Another option is to make it easier to take civil action in response to privacy violations and provide financial compensation to victims of those violations. To do this, Contreras argues that we need a liability framework for privacy harms instead of the current control-based and property-based framework.^{31 32} Current laws focus on protecting data as property, when they should focus on holding people (and organisations) accountable when they use health data in harmful ways. Stronger rules regarding proper use of health information would allow individuals to take civil action (eg, class action lawsuits) against those who break the rules. Under a liability framework, those harmed

by data misuse would have greater recourse to receive monetary damages.

Trade-offs and drawbacks to penalisation

The main drawback to penalisation is that it is unclear whether the above-mentioned proposals could fit within existing legal frameworks. For example, the Health Insurance Portability and Accountability Act (HIPAA) and its extensions under the Health Information Technology for Economic and Clinical Health (HITECH) Act are primarily concerned with regulating access to health data. A challenge these frameworks face is that they do not protect us from the enormous amount of health data processing that already occurs outside of healthcare contexts.^{33 34} As many have pointed out, 'shadow' health records already exist outside of HIPAA protections, which make HIPAA and HITECH largely irrelevant to privacy protection.^{12 35} Including new rules about liability and privacy harms to the HIPAA framework would likely still fail to address these other types of health data.

Similar problems apply to other frameworks many have hoped would improve privacy protections, including the California Consumer Privacy Act (CCPA) and, when sharing data outside of the USA, the European Union's General Data Protection Regulation (GDPR). The GDPR requires companies to obtain individual consent before processing that individual's data, while the CCPA requires companies to tell people when their data are collected, sold or shared, and allow them to opt out. However, data brokers can easily avoid these requirements with access to enough of the right deidentified information. They can still make inferences about an individual's health while plausibly maintaining that they do not know whose information it is (which also prevents withdrawal for any particular individual's data).

It is also unclear whether the above-mentioned proposals could fit within the Genetic Information Nondiscrimination Act (GINA), which might otherwise seem like a model for how to better penalise data misuse. Under GINA, employers and health insurers are not allowed to access genetic information or to use even proxies of genetic information in decision making, like someone's family medical history. If they do, they can be sued or fined by the Equal Employment Opportunity Commission, which provides a reason for at least some third parties to think twice before violating people's genetic privacy. However, as Prince and Schwarcz argue, advances in artificial intelligence and machine learning enable insurance companies to make inferences in ways that overcome that feature of the law.³⁶ Health insurance algorithms are designed to find proxies for genetic risks, even if they are not specifically trained to do so. Such proxies enable insurance companies to discriminate in a way that does not violate GINA. Data brokers would likely take advantage of the same loophole to overcome the above-mentioned proposals, if applied to GINA.

In short, existing legal frameworks seem unhelpful for implementing the above-mentioned proposals. The limitations of current frameworks suggest that a broader, more sweeping privacy law would be required to adequately penalise the misuse of health data. This new legislation would have to clarify how privacy violations can cause harm while also covering the many ways health data are used outside of the health system. Without stronger laws, adversaries have little to deter them from misusing personal health data.

Transparency of health data

The third route focuses on improving transparency. We may not be able to prevent others from accessing our personal information, but we can still track who has that information and how

it is used. Obtaining this knowledge helps to alter the incentives for exploiting information for personal gain. The reputational damage of being known as an organisation that exploits personal data is, in many cases, not worth any other potential benefit. Numerous studies have shown that consumer behaviour changes in response to privacy concerns and that corporations feel the effects of such changes.^{37 38} For example, an analysis of 150 companies' annual disclosure statements found that 95% of those with privacy issues were primarily concerned about the reputational damage that could occur.³⁹ A more transparent system would also facilitate law enforcement action.

The typical approach to improving transparency of health data is to increase individuals' visibility into their own health data. The HITECH Act, for example, requires that electronic health records be released to patients in a timely manner,⁴⁰ in addition to requiring 'audit trails', noting every instance a patient's records have been accessed.⁴¹ California's recently passed privacy legislation takes this a step further by requiring third parties outside of the health system to tell people what information is collected about them and why. It also grants individuals the right to request a copy of their personal information held by companies in a 'readily useable format'.⁴²

However, we must also address the 'shadow health record' problem mentioned previously, or these forms of increased visibility will be inadequate. As a possible remedy, some states are currently trying to increase visibility into the records held by data brokers. Vermont, for instance, requires data brokers to register with the state so people at least know which companies are likely processing their data.⁴³

Another option to potentially avoid data brokers is to give patients control over their data.⁴⁴ In principle, this could allow patients, rather than hospitals or health insurance companies, to sell data on their own. For example, companies like Humantix help people to sell their own data directly to other companies.⁴⁵ Data brokers would still possess shadow records but individuals' own records may be more valuable, which would disrupt data brokers' current business model. Individuals could even negotiate their own deals with data brokers in order to ensure that they understand exactly which of their personal information is publicly available.

A problem with these individual-based options is that they do not adequately publicise how health data are used. They rely too much on individuals to notify others of misuse, if they notice it at all. Under a model of pervasive data sharing, it would likely be preferable for governmental entities responsible for oversight to conduct this monitoring instead. Significant time and resources are required to properly track and publicise the various ways that third parties share and process health data.

Trade-offs and drawbacks to transparency

There are two fundamental challenges to improving transparency with health data. The first is that people may not understand the implications of who has their data or how the data are used. Health data are shared with dozens of entities, often through automated systems, and with varying degrees of administrative and billing information that people probably will not understand or care about. Achieving complete visibility into where health information is going is likely to be a mix of boring and confusing. This suggests that health data must be made intelligible as well as transparent.

Giving patients complete control over their data only enhances the challenge. Even if pervasive data sharing is the norm, individuals might still have the power to limit data sharing with providers, insurers, researchers and public health entities.

Assuming individuals do not fully understand how health data are used, they likely will not realise the negative impact such control would have on things like healthcare costs and public health.⁴⁶

The second fundamental challenge is that increased transparency may not lead to changes in behaviour. So many companies could be involved in pervasive data sharing for so many different questionable purposes that people could become desensitised to misuse. Complete transparency might be perfectly compatible with systematic exploitation.

Nonetheless, transparency is the last line of defence if we have lost all control of our health data. The obfuscation and penalisation methods outlined previously are ineffective if we do not even know that the data are being shared or that the databases already exist. Moreover, the combination of these methods can help overcome the potential disconnect between transparency and action. It matters less that individuals are unconcerned about data sharing practices if the data cannot clearly be connected to them and there are strong legal protections against misuse.

CONCLUSION

To protect our privacy in a world where we no longer control our data, we must obfuscate health data, penalise the misuse of health data, and improve transparency around who shares our data and for what purposes. Changing laws around privacy is perhaps the most important of the three strategies but also the most difficult. GINA, GDPR and related laws were decades in the making. Recent measures, such as California's new privacy law and other laws concerning data brokers, may enable quicker action to ensure protection of sensitive health data. However, as suggested earlier, a more sweeping privacy law is likely needed to address data processing outside of healthcare contexts.

Transparency and obfuscation methods may be more immediately implementable, as the information technology infrastructure and data management obstacles they face are continuously less burdensome. Synthetic databases, for instance, are becoming more cost-effective and are an increasingly attractive alternative to cumbersome deidentification methods. It is also becoming easier to track patient data in a way that can be easily communicated back to patients. Tests of these methods in their early implementation stage will provide crucial information on whether they can in fact compensate for the inadequacies of control-based methods.

Acknowledgements Many thanks to David Wright for his thoughtful feedback.

Contributors TJK has contributed to all aspects of this research and is the sole contributor.

Funding The authors have not declared a specific grant for this research from any funding agency in the public, commercial or not-for-profit sectors.

Competing interests None declared.

Patient consent for publication Not required.

Provenance and peer review Not commissioned; externally peer reviewed.

ORCID iD

T J Kasperbauer <http://orcid.org/0000-0003-0216-7632>

REFERENCES

- 1 Cavoukian A, Fisher A, Killen S, *et al*. Remote home health care technologies: how to ensure privacy? build it in: privacy by design. *Identity Inf Soc* 2010;3(2):363–78.
- 2 Nuffield Council on Bioethics. Biological and health data, 2015. Available: http://nuffieldbioethics.org/wp-content/uploads/DataEthics_Chapter5.pdf
- 3 InBeauchamp T, Childress J. *Principles of biomedical ethics*. 7th edn. New York: Oxford University Press, 2013: 312–4.

- 4 Belli L, Schwartz M, Louzada L. Selling your soul while negotiating the conditions: from notice and consent to data control by design. *Health Technol* 2017;7(4):453–67.
- 5 Darcy AM, Louie AK, Roberts LW. Machine learning and the profession of medicine. *JAMA* 2016;315(6):551–2.
- 6 Dhindsa K, Bhandari M, Sonnadara RR. What's holding up the big data revolution in healthcare? *BMJ* 2018;363:k5357.
- 7 Naylor CD. On the prospects for a (deep) learning health care system. *JAMA* 2018;320(11):1099–100.
- 8 National Institutes of Health. All of US research program. Available: <https://allofus.nih.gov/>
- 9 Clayton EW, Halverson CM, Sathe NA, *et al.* A systematic literature review of individuals' perspectives on privacy and genetic information in the United States. *PLoS One* 2018;13(10):e0204417.
- 10 Grajales F, Clifford D, Loupos P, *et al.* Social networking sites and the continuously learning health system: a survey, 2014. Available: <https://nam.edu/wp-content/uploads/2015/06/VSRT-PatientDataSharing.pdf>
- 11 Trachtenbarg DE, Asche C, Ramsahai S, *et al.* The benefits, risks and costs of privacy: patient preferences and willingness to pay. *Curr Med Res Opin* 2017;33(5):845–51.
- 12 Price WN, Kaminski ME, Minssen T, *et al.* Shadow health records meet new data privacy laws. *Science* 2019;363(6426):448–50.
- 13 Allen M. Health insurers are vacuuming up details about you—and it could raise your rates, 2018. Available: <https://www.propublica.org/article/health-insurers-are-vacuuming-up-details-about-you-and-it-could-raise-your-rates>
- 14 Huesch MD. Privacy threats when seeking online health information. *JAMA Intern Med* 2013;173(19):1838–9.
- 15 Andrews L. A new privacy paradigm in the age of apps. *Wake Forest L Rev* 2018;53:421–78.
- 16 Grundy Q, Chiu K, Held F, *et al.* Data sharing practices of medicines related apps and the mobile ecosystem: traffic, content, and network analysis. *BMJ* 2019;22.
- 17 Brunton F, Nissenbaum H. *Obfuscation: A user's guide for privacy and protest*. Cambridge: MIT Press, 2015: 38.
- 18 Hartzog W. *Privacy's blueprint*. Cambridge: Harvard University Press, 2018: p. 96.
- 19 Sweeney L. Only you, your doctor, and many others may know. *Tech Sci* 2015.
- 20 El Amam K. *Guide to the de-identification of personal health information*. Boca Raton, FL: CRC Press, 2013.
- 21 Foraker R, Mann DL, Payne PRO. Are synthetic data derivatives the future of translational medicine? *JACC Basic Transl Sci* 2018;3(5):716–8.
- 22 Walonoski J, Kramer M, Nichols J, *et al.* Synthea: an approach, method, and software mechanism for generating synthetic patients and the synthetic electronic health care record. *JAMIA* 2017;25(3):230–8.
- 23 Chen J, Chun D, Patel M, *et al.* The validity of synthetic clinical data: a validation study of a leading synthetic data generator (Synthea) using clinical quality measures. *BMC Med Inform Decis Mak* 2019;19(1):44.
- 24 Dankar FK, El Amam K. Practicing differential privacy in health care: a review. *Trans Data Privacy* 2013;5:35–67.
- 25 Aircloak. Anonymization with Aircloak: how it works. Available: <https://aircloak.com/solutions/how-it-works/>
- 26 Barocas S, Nissenbaum H. Big data's end run around anonymity and consent. In: Lane J, Stodden V, Nissenbaum H, eds. *Privacy, big data, and the public good*. Cambridge: Cambridge University Press, 2014: 45.
- 27 Terry N, Wiley LF. Liability for mobile health and wearable technologies. *Ann Health Law* 2016;25:62–97.
- 28 Solove D, Citron D. Risk and anxiety: a theory of data breach harms. *Tex Law Rev* 2018;96:737–86.
- 29 Solove D, Citron D. Risk and anxiety: a theory of data breach harms. *Tex Law Rev* 2018;96.
- 30 Calo R. Privacy harm exceptionalism. *Colo Tech Law J* 2014;12:361–4.
- 31 Contreras J. Genetic property. *Georgetown Law J* 2016;105:1–54.
- 32 Contreras J, Nordfalk F. Liability (and) rules for health information. *Health Matrix* 2019;29(1):179–223.
- 33 Clayton EW, Evans BJ, Hazel JW, *et al.* The law of genetic privacy: applications, implications, and limitations. *J Law Biosci* 2019;6(1):1–36.
- 34 Cohen IG, Mello MM. HIPAA and protecting health information in the 21st century. *JAMA* 2018;320(3):231–2.
- 35 Pasquale F, Ragone TA. Protecting health privacy in an era of big data processing and cloud computing. *Stanford Tech Law Rev* 2014;17:595–654.
- 36 Prince A, Schwarcz DB. Proxy discrimination in the age of artificial intelligence and big data. *Iowa Law Rev* 2019.
- 37 Goldberg R. Lack of trust in Internet privacy and security may deter economic and other online activities, 2016. Available: <https://www.ntia.doc.gov/blog/2016/lack-trust-internet-privacy-and-security-may-deter-economic-and-other-online-activities>
- 38 PWC. How consumers see cybersecurity and privacy risks and what to do about it, 2017. Available: <https://www.pwc.com/us/en/services/consulting/library/consumer-intelligence-series/cybersecurity-protect-me.html>
- 39 International Association for Privacy Protection. Privacy risk study 2017. Available: <https://iapp.org/resources/article/privacy-risk-study-2017-pii-remains-top-information-risk/>
- 40 45 CFR § 164.524.
- 41 45 CFR § 170.210.
- 42 Cal. Civ. Code §1798.100.
- 43 Melendez S, Pasternack A. Here are the data brokers quietly buying and selling your personal information, 2019. Available: <https://www.fastcompany.com/90310803/here-are-the-data-brokers-quietly-buying-and-selling-your-personal-information>
- 44 Topol E. *The patient will see you now*. New York: Basic Books, 2015.
- 45 Neville S. How patients can turn their medical data into money, 2019. Available: <https://www.ft.com/content/34add2b8-2eba-11e9-8744-e7016697f225>
- 46 Contreras J. The false promise of health data ownership. *NYU Law Rev*;94:624–61.