

Viewpoint

Privacy Is Dead, Long Live Privacy

Protecting social norms as confidentiality wanes.

THE PAST FEW years have been especially turbulent for privacy advocates. On the one hand, the global dragnet of surveillance agencies has demonstrated the sweeping surveillance achievable by massively resourced government organizations. On the other, the European Union has issued a mandate that Google definitively “forget” information in order to protect users.

Privacy has deep historical roots, as illustrated by the pledge in the Hippocratic oath (5th century B.C.), “Whatever I see or hear in the lives of my patients ... which ought not to be spoken of outside, I will keep secret, as considering all such things to be private.”¹¹ Privacy also has a number of definitions. A now common one among scholars views it as the flow of information in accordance with social norms, as governed by context.¹⁰ An intricate set of such norms is enshrined in laws, policies, and ordinary conduct in almost every culture and social setting. Privacy in this sense includes two key notions: confidentiality and fair use. We argue that confidentiality, in the sense of individuals’ ability to preserve secrets from governments, corporations, and one another, could well continue to erode. We call instead for more attention and research devoted to fair use.

To preserve existing forms of privacy against an onslaught of online



working hard to develop *privacy-enhancing technologies* (PETs). PETs enable users to encrypt email, conceal their IP addresses, avoid tracking by Web servers, hide their geographic location when using mobile devices, use anonymous credentials, make untraceable database queries, and publish documents anonymously. Nearly all major PETs aim at protecting confidentiality: we call these *confidenti-*

can be good and helpful. But there is a significant chance that in many or most places, C-PETs will not save privacy. It is time to consider adding a new research objective to the community’s portfolio: preparedness for a post-confidentiality world in which many of today’s social norms regarding the flow of information are regularly and systematically violated.

Global warming offers a useful

ingly unstoppable human-induced disaster and a worldwide tragedy of commons. Scientists and technologists are developing a portfolio of mitigating innovations in renewable energy, energy efficiency, and carbon sequestration. But they are also studying ways of coping with likely effects, including rising sea levels and displacement of populations. There is a scientific consensus that the threat justifies not just mitigation, but preparation (for example, elevating Holland's dikes).

The same, we believe, could be true of privacy. Confidentiality may be melting away, perhaps inexorably: soon, a few companies and surveillance agencies could have access to most of the personal data of the world's population. Data provides information, and information is power. An information asymmetry of this degree and global scale is an absolute historical novelty.

There is no reason, therefore, to think of privacy as we conceive of it today as an enduring feature of life.

Example: RFID

Radio-Frequency IDentification (RFID) location privacy concretely illustrates how technological evolution can undermine C-PETs. RFID tags are wireless microchips that often emit static identifiers to nearby readers. Numbering in the billions, they in principle permit secret local tracking of ordinary people. Hundreds of papers proposed C-PETs that rotate identifiers to prevent RFID-based tracking.⁶

Today, this threat seems quaint. Mobile phones with multiple RF interfaces (including Bluetooth, Wi-Fi, NFC), improvements in face recognition, and a raft of new wireless devices (fitness trackers, smartwatches, and other devices), offer far more effective ways to track people than RFID ever did. They render RFID C-PETs obsolete.

This story of multiplying threat vectors undermining C-PETs' power—and privacy more generally—is becoming common.

The Assault on Privacy

We posit four major trends providing the means, motive, and opportunity

est sense. The adversaries include surveillance agencies and companies in markets such as targeted advertising, as well as smaller, nefarious players.

Pervasive data collection. As the number of online services and always-on devices grows, potential adversaries can access a universe of personal data quickly expanding beyond browsing history to location, financial transactions, video and audio feeds, genetic data⁴, real-time physiological data—and perhaps eventually even brainwaves.⁸ These adversaries are developing better and better ways to correlate and extract new value from these data sources, especially as advances in applied machine learning make it possible to fill in gaps in users' data via inference. Sensitive data might be collected by a benevolent party for a purpose that is acceptable to a user, but later fall into dangerous hands, due to political pressure, a breach, and other reasons. "Secondhand" data leakage is also growing in prevalence, meaning that one person's action impacts another's private data (for example, if a friend declares a co-location with us, or if a blood relative unveils her genome). The emerging Internet of Things will make things even trickier, soon surrounding us with objects that can report on what we touch, eat, and do.¹⁶

Monetization (greed). Political philosophers are observing a drift from what they term *having* a market economy to *being* a market society¹³ in which market values eclipse non-market social norms. On the Internet, the ability to monetize nearly every piece of information is clearly fueling this process, which is itself facilitated by the existence of quasi-monopolies. A market-

There is no reason to think of privacy as we conceive of it today as an enduring feature of life.

place could someday arise that would seem both impossible and abhorrent today. (For example, for \$10: "I know that Alice and Bob met several times. Give me the locations and transcripts of their conversations.") Paradoxically, tools such as anonymous routing and anonymous cash could facilitate such a service by allowing operation from loosely regulated territories or from no fixed jurisdiction at all.

Adaptation and apathy. Users' data curation habits are a complex research topic, but there is a clear generational shift toward more information sharing, particularly on social networks. (Facebook has more than one billion users regularly sharing information in ways that would have been infeasible or unthinkable a generation ago.). Rather than fighting information sharing, users and norms have rapidly changed, and convenience has trumped privacy to create large pockets of data-sharing apathy. Foursquare and various other microblogging services that encourage disclosure of physical location, for example, have led many users to cooperate in their own physical tracking. Information overload has in any event degraded the abilities of users to curate their data, due to the complex and growing challenges of "secondhand" data-protection weakening and inference, as noted previously.

Secret judgment. Traceability and accountability are essential to protecting privacy. Facebook privacy settings are a good example of visible privacy practice: stark deviation from expected norms often prompts consumer and/or regulatory pushback.

Increasingly often, though, sensitive-data exploitation can happen away from vigilant eyes, as the recent surveillance scandals have revealed. (National security legitimately demands surveillance, but its scope and oversight are critical issues.) Decisions made by corporations—hiring, setting insurance premiums, computing credit ratings, and so forth—are becoming increasingly algorithmic, as we discuss later. Predictive consumer scores are one example; privacy scholars have argued they constitute a regime of secret, arbitrary, and potentially discriminatory

A Post-Confidentiality Research Agenda

We should prepare for the possibility of a post-confidentiality world, one in which confidentiality has greatly eroded and in which data flows in such complicated ways that social norms are jeopardized. The main research challenge in such a world is to preserve social norms, as we now explain.

Privacy is important for many reasons. A key reason, however, often cited in discussions of medical privacy, is concern about abuse of leaked personal information. It is the potentially resulting unfairness of decision making, for example, hiring decisions made on the basis of medical history, that is particularly worrisome. A critical, defensible bastion of privacy we see in post-confidentiality world therefore is in the *fair use* of disclosed information.

Fair use is increasingly important as algorithms dictate the fates of workers and consumers. For example, for several years, some Silicon Valley companies have required job candidates to fill out questionnaires (“Have you ever set a regional-, state-, country-, or world-record?”). These companies apply classification algorithms to the answers to filter applications.⁵ This trend will surely continue, given the many domains in which statistical predictions demonstrably outperform human experts.⁷ Algorithms, though, enable deep, murky, and extensive use of information that can exacerbate the unfairness resulting from disclosure of private data.

On the other hand, there is hope that algorithmic decision making can lend itself nicely to protocols for enforcing accountability and fair use. If decision-making is algorithmic, it is possible to require decision-makers to prove that they are not making use of information in contravention of social norms expressed as laws, policies, or regulations. For example, an insurance company might prove it has set a premium without taking genetic data into account—even if this data is published online or otherwise widely available. If input data carries authenticated labels, then cryptographic techniques permit the construction of such proofs without revealing underlying algorithms

If we cannot win the privacy game definitively, we need to defend paths to an equitable society.

secrets (for example, see Ben-Sasson et al.¹). Use of information flow control¹² preferably enforced by software attested to by a hardware root of trust (for example, see McKeen et al.⁹) can accomplish much the same end. Statistical testing is an essential, complementary approach to verifying fair use, one that can help identify cases in which data labeling is inadequate, rendered ineffective by correlations among data, or disregarded in a system. (A variety of frameworks exist, for example, see Dwork et al.³)

Conclusion

A complementary research goal is related to privacy quantification. To substantiate claims about the decline of confidentiality, we must measure it. Direct, global measurements are difficult, but research might look to indirect monetary ones: The profits of the online advertising industry per pair of eyeballs and the “precision” of advertising, perhaps as measured by click-through rates. At the local scale, research is already quantifying privacy (loss) in such settings as location-based services.¹⁴

There remains a vital and enduring place for confidentiality. Particularly in certain niches—protecting political dissent, anti-censorship in repressive regimes—it can play a societally transformative role. It is the responsibility of policymakers and society as a whole to recognize and meet the threat of confidentiality’s loss, even as market forces propel it and political leaders give it little attention. But it is also incumbent upon the research community to contemplate alternatives to C-PETs, as confidentiality is

social evolution. If we cannot win the privacy game definitively, we need to defend paths to an equitable society. We believe the protection of social norms, especially through fair use of data, is the place to start. While C-PETs will keep being developed and will partially mitigate the erosion of confidentiality, we hope to see many “fair-use PETs” (F-PETs) proposed and deployed in the near future.¹⁵ ■

References

1. Ben-Sasson, E. et al. SNARKs for C: Verifying program executions succinctly and in zero knowledge. In *Advances in Cryptology—CRYPTO*, (Springer, 2013), 90–108.
2. Dixon, P. and Gellman, R. The scoring of America: How secret consumer scores threaten your privacy and your future. Technical report, World Privacy Forum (Apr. 2, 2014).
3. Dwork, C. et al. Fairness through awareness. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*. (ACM, 2012), 214–226.
4. Erlich, Y. and Narayanan, A. Routes for breaching and protecting genetic privacy. *Nature Reviews Genetics* 15, 6 (2014), 409–421.
5. Hansell, S. Google answer to filling jobs is an algorithm. *New York Times* (Jan. 3, 2007).
6. Juels, A. RFID security and privacy: A research survey. *IEEE Journal on Selected Areas in Communication* 24, 2 (Feb. 2006).
7. Kahneman, D. *Thinking, Fast and Slow*. Farrar, Straus, and Giroux, 2012, 223–224.
8. Martinovic, I. et al. On the feasibility of side channel attacks with brain-computer interfaces. In *Proceedings of the USENIX Security Symposium*, (2012), 143–158.
9. McKeen, F. et al. Innovative instructions and software model for isolated execution. In *Proceedings of the 2nd International Workshop on Hardware and Architectural Support for Security and Privacy*, Article no. 10 (2013).
10. Nissenbaum, H. *Privacy in Context: Technology, Policy, and the Integrity of Social Life*. Stanford University Press, 2009.
11. North, M.J. Hippocratic oath translation. U.S. National Library of Medicine, 2002.
12. Sabelfeld, A. and Myers, C. Language-based information-flow security. *IEEE Journal on Selected Areas in Communications* 21, 1 (2003), 5–19.
13. Sandel, M.J. *What Money Can't Buy: The Moral Limits of Markets*. Macmillan, 2012.
14. Shokri, R. et al. Quantifying location privacy. In *Proceedings of the IEEE Symposium on Security and Privacy* (2011), 247–262.
15. Tramèr, F. et al. Discovering Unwarranted Associations in Data-Driven Applications with the FairTest Testing Toolkit, 2016; arXiv:1510.02377.
16. Weber, R.H. Internet of things—New security and privacy challenges. *Computer Law and Security Review* 26, 1 (2010), 23–30.

Jean-Pierre Hubaux (jean-pierre.hubaux@epfl.ch) is a professor in the Computer Communications and Applications Laboratory at the Ecole Polytechnique Fédérale de Lausanne in Switzerland.

Ari Juels (juels@cornell.edu) is a professor at Cornell Tech (Jacobs Institute) in New York.

We would like to thank George Danezis, Virgil Gligor, Kévin Huguenin, Markus Jakobsson, Huang Lin, Tom Ristenpart, Paul Syverson, Gene Tsudik and the reviewers of this Viewpoint for their many generously provided, helpful comments, as well as the many colleagues with whom we have shared discussions on the topic of privacy. The views presented in this Viewpoint remain solely our own.