A Sense of Privacy: How Privacy's Past May Shape Its Future

Authors: Alessandro Acquisti[1*], Laura Brandimarte[2], Jeff Hancock[3]

Affiliations:

[1]Heinz College, Carnegie Mellon University; 5000 Forbes Avenue, Pittsburgh, PA 15213-3890.

[2]Eller College of Management, The University of Arizona; 1130 E. Helen Street, Tucson, AZ 85721-0108.

[3]Department of Communications, Stanford University; 450 Jane Stanford Way, Stanford, CA 94305-2050.

[*]Corresponding author. Email: acquisti@andrew.cmu.edu.

Abstract: Expanding on the theory that a drive for privacy is a constant throughout human history and may have evolved from physical security concerns, this article argues that the evolutionary roots of privacy render common policy approaches to privacy protection inadequate. If humans evolved their sense of privacy from stimuli arising from their physical surroundings (such as detecting the presence of others), then the transition to a digital world disrupts individuals' regulation of privacy—because physical cues online are missing or easily manipulated. As such, the widely used policy framework of notice and consent, which places the burden and responsibility of privacy protection on the individual, is bound to fail. New frontiers of privacy management, using technologies to embed privacy-by-default into systems, offer a viable alternative, but the industry may have little economic incentive to implement them. The only option for meaningful privacy management therefore is policy intervention.

One-Sentence Summary: The evolutionary roots of privacy render current systems of privacy management that rely on user responsibilization inadequate, and highlight the need for determined regulatory interventions.

Main Text

To understand the current state of privacy and its likely future, we must study its history. Privacy has been construed as a modern anomaly—something our ancestors lacked and technological innovation and urban growth made possible (*1*). Research from history, anthropology, and ethnography suggests otherwise. The evidence for peoples seeking to manage the boundaries of private and public spans time and space, social class and degree of technological sophistication. From the Greek philosophers of the fourth century BC "discuss[ing] the concept of private life" (*2*, *p. 120*), to Chinese thinkers developing "a sharp distinction between the concepts of public and private" by the third century BC (*p. 221*); from lovers in ancient Rome who would need to "borrow the house of an indulgent friend" to escape the prying eyes of their servants (*3*, *p. 72*), to the Mehinacu tribe in Southern America, living in communal accommodations but building private houses miles away from the village to achieve some reclusion (*4*), up to 1950s Javanese culture, where little physical seclusion was available and people manifested their privacy needs through stiff social interactions (*5*). Evidence of privacy mores even arises from within the ancient holy books of monotheistic religions—from the Quran to the Talmud to the Bible (*6*). Norms about, opportunities for, and even definitions of privacy do fluctuate over time and, markedly, across cultures. And yet, strikingly, privacy-seeking behaviors emerge across peoples throughout history and across geography. As Irwin Altman concluded, privacy—understood not merely as hiding of data, but as the selective opening and closing of the self to others—appears to be both culturally specific and culturally universal (*7*).

But if so, why? What could explain the simultaneous universality and diversity of a human drive for privacy? And what could such an explanation teach us about privacy's digital future? An account of the evolutionary roots of privacy may offer an answer.

The Evolutionary Roots of Privacy

The notion that a concern for privacy may have evolutionary roots may surprise modern netizens, used to construing privacy and its invasions in terms of data rights and Internet surveillance. And yet, this notion—and examples supporting it—can be found in the writings of leading privacy scholars from half a century ago (*6*, *8*). Under this account, what we refer to today as privacy evolved from physical needs for security and self-interest. By using signals to assess threats and opportunities in their physical proximity, humans and other species enhance their chance of survival and evolutionary success. In our past, the ability to detect and react appropriately to the presence of others by selectively choosing exposure or protection, openness or withdrawal, would have offered an evolutionary advantage—for instance, distinguishing kin from stranger and adapting behavior from cooperative to protective accordingly. Under this account, the evolutionary advantage of being able to process and appropriately react to *sensorial cues* that suggest the presence of others is clear.

Over time, in humans, evolutionary adaptation developed more sophisticated forms of social cognition, balancing the tensions between individual freedom and collective welfare (*9)*. In the case of privacy, those developments may have followed various possible vectors. In one of them, our ancestral ability to detect and react to the physical presence of others adapted into reputation and impression management (*10*). Barkow (*11*), for example, suggests that the primary evolutionary function of the self could have been impression management. Selecting information about oneself to make favorable impressions on others (*12*) and using reticence and withdrawal to "reduce the risk of saying or doing something that others might regard negatively" (*13*, *p. 520*) are instruments, then and now, for such management (*14*).

There may, therefore, be a line connecting the seemingly diverse manifestations of privacy across human history: privacy as the selective, self-interested opening and closing of the self to others (*15*). This dialectic is a universal characteristic of the human species—our eternal tension between a "desire for seclusion" and a "need to depend on others" (*16, p. 5*). The boundaries between self and others take multiple forms—from physical to informational to regulatory. This explains the diversity in dimensions, definitions, and norms of privacy over time and across the literature (*17*).

The ability to leverage sensorial cues to detect others and modulate appropriate responses is advantageous not just in terms of protection from risk, but also extraction of benefits. Supporting the evolutionary account, both dynamics have been observed in animal species (*5, Chapter 1*), such as cats seeking seclusion under duress (protection against risk: *18*), and chimps lower in the social hierarchy concealing mating activities or a coveted source of food from higher-status males (extraction of benefits: *19, pp. 36-37* and *p. 62*; *20*). Both dynamics are also surprisingly consistent with modern economic accounts of *informational* privacy as a self-interested process of selective sharing: an individual may rationally want to share with a marketer their interests and preferences, so as to get beneficial targeted offers (extraction of benefits); but may not want to share with the marketer their willingness to pay for those interests, so as to avoid price discrimination (protection against risk) (*21*).

Sensorial cues that indicate that other humans are present still influence today how we define the boundaries between public and private. Consider the familiar feeling alerting us that someone is staring at us as we were lost in thought: a "sense" of privacy, so to say. In our own research, we found that experimental participants were less likely to divulge sensitive personal information in an online survey when they visually detected the presence of another person, either in the same room or—notably—in an adjacent room from which that person could not have seen what the participant was typing (*22*). Indeed, recent work has observed the power of visual cues in young children's conceptualizations of privacy risks with technology (*23*). Where visual cues are present, children have much more accurate understanding of how a technology works (for instance, why they see content recommendations), whereas they have no sense that data is flowing back to companies when visual cues are lacking.

A Privacy Gap

And herein lies the problem. If people do in fact rely, in part, on sensorial cues to navigate privacy choices, the more we transition from physical to digital interactions, the less equipped we may be for informed digital privacy decisions. At the extreme, offline sensorial cues we depend upon may, online, be absent, subdued, or even intentionally manipulated through so-called dark patterns (*24, 25*). Thus, privacy (and security) responses common in the physical world may be muted or entirely unavailable online. This may explain the unique hurdles we face in protecting privacy online, and the surprising observation of seemingly careless online behaviors by individuals who claim to care about their privacy (*26*). In our non-digital life, Altman (*7*) noted, behaviors that help us manage privacy boundaries are common and ubiquitous. We often engage in them with little conscious awareness. We lower our voice during intimate conversations and raise it when we want a large audience to hear us; we cover a document we are reading so it is protected from prying eyes, or raise it up and emphatically show it when we want to make a point. But online, the evolutionary mechanisms Westin (*5*) and Klopfer and Rubenstein (*8*) identified cannot help us: we do not *see* Google leaning over our shoulders to track our sensitive searches; we do not *hear* the NSA stepping closer to listen to our video-conferences. And yet, human beings may be wired to rely in part on those very sensorial cues to assess the privacy implications of their behaviors. This discrepancy may create a gap, or mismatch, in our ability to manage privacy sensitive scenarios in digital realms relative to the physical world - a hypothesis we have proposed (*22, 27*) and that Sharif*,* Green, and Jettinghoff have also independently developed (*28*).

This privacy mismatch has distinctive features vis a vis other evolutionary mismatches of modernity (*29*), such as the modern diet's mismatch with our ancestral history (*30*). Where the evolutionary mismatch between human evolved physiology and the modern diet arises entirely within physical systems, the very boundaries of privacy have evolved with the digital transition—from predominantly physical to predominantly informational. Through that transition, privacy's role and relevance in society have evolved as well—as did the value individuals and societies can harness from personal data. This mismatch does not imply that online disclosures are inherently damaging. Nor should we fall into a naturalistic fallacy of elevating senses to be sole arbiters of privacy decision-making, even online. Rather, a privacy mismatch implies that the ability to regulate privacy boundaries (choosing when to share and when to protect) is in an individual's best interest; and when that evolutionarily-rooted ability is impaired, individuals become more vulnerable, online, to novel threats: ubiquitous instances of micro-level influence on decision making which, in the aggregate, shape portentous macro-level dynamics. While offline and online surveillance risks differ along numerous dimensions (including the likelihood of material harm), the privacy mismatch is, in fact, happening at the same time when ramifications of informational privacy issues become consequential not merely at the individual but at the *societal* level. Privacy externalities (*31)*—the collective ramifications of individual disclosures—such as algorithmic personalization spawning filter bubbles (*32*), amplification of disinformation (*33*), and political polarization (*34*, *35*), with implications for democratic elections (*36*). And while the privacy mismatch arises from our transition to digital interactions, the consequences are not merely digital, but physical as well: consider the material threats associated with "doxxing" (*37*) and episodes of genocidal violence fostered by data-driven personalization algorithms (*38*).

An evolutionary account of privacy provides the underlying reason why a dominant approach to privacy management in the United States—notice and consent—has failed to address these problems. While popular with industry, an overreliance on notice and consent mechanisms, disjointed from baseline privacy safeguards, is ineffectual (*39*) and can backfire (*40*) —because we may be wired to react to privacy invasions viscerally, not just deliberatively. Those mechanisms burden citizens with "responsibilization" (*41*) for problems they did not create and cannot truly control. Responsibilization also creates unequal burdens, as it disadvantages certain groups more than others: to get the most out of notice and consent, for example, people need resources—time, education, the economic leverage to not consent to unfair policies, and so forth—and these resources are not equally distributed across the population. Even substantial interventions such as the General Data Protection Regulation in the EU resulted in a proliferation of consent mechanisms that burden users with manipulative interfaces and implied consent (*42, 43*). Notice and consent approaches fail to recognize that much of our privacy behavior is intuitive and visceral (*44*), rendering them archaic solutions for a modern problem.

Using Policy to Foster and Embed Privacy Technologies

So, what should be done? *If*, as a society, we determined that privacy was still valuable to us, then to maintain it we would have to embed it into the fabric of our digital systems as default. We cannot demand people to overcome an evolved sense of privacy strongly reliant on sensorial cues unavailable online. Any approach that places not just the *ability* to choose, but—ultimately—the *responsibility* to protect on individuals themselves, will—according to this account—fail. Privacy mismatches will keep rising in frequency and significance with the growth of the information economy, past the point where any approach relying upon individual responsibility alone could contain them (28). Individuals' personal information will keep being collected across too many instances and vectors (*45*) for our bounded cognitive resources to manage efficiently or effectively. And humans will keep tuning out even visible signs of electronic surveillance (*46*). This is why attempts to reproduce online the visceral cues of the physical world are unlikely, alone, to solve the problem: an evolutionary mechanism of explanation does not necessarily demand an evolutionary mechanism of change (*47*). For those same reasons, even well-meaning proposals such as data propertization from economics (*48, 49*), nudges from behavioral research (*25*), and simplified notices from usability research (*50*), while useful, may hardly make a difference in absence of a combination of technology and policy interventions that embed privacy safeguards into the core of our information systems.

Consider an analogy. Thanks to continuous technological improvements, the speed and acceleration of production cars kept growing over time. Once they reached velocities that rendered drivers' reaction times unreliable tools for avoiding collisions, the solution was not to teach individuals to develop faster reaction abilities, but rather to develop policy interventions (e.g., mandatory safety standards on cars for accident avoidance and damage reduction) and technological fixes (e.g., ABS, airbags) that countered the challenges arising from other technological progress. Better safety in cars was the result of deliberate policy intervention driving investment in technical and infrastructural changes, not merely of more driver education and market forces (*51*).

In the case of privacy, policy intervention can both instill baseline safeguards (such as those embedded in the OECD's Guidelines on the Protection of Privacy and Transborder Flows of Personal Data; *52*) and foster the deployment of technologies that make those safeguards possible without hurting modern society's reliance on data.

The last few decades have produced not only a burst of innovation around data as critical asset for economic and societal development; they have also generated innovations in statistics, cryptography, and computer science that may address the challenges of creeping surveillance while allowing value to keep being extracted from data. Protocols from differential privacy to homomorphic encryption to federated learning point at the possibility of protecting individuals' privacy while allowing beneficial analytics to advance (*53*). Tools from AI and machine learning deployed in privacy assistants suggest a future where computerized agents learn and then autonomously represent human preferences in interactions with services (*54, 55*). These developments portend a world where privacy by design (*56*) is possible without undermining the value that society can extract from flows of data. Indeed, economic research suggests that data protection is not inherently welfare-decreasing (see *26, 31*), and the use of differentially private algorithms to achieve the dual goal of producing accurate statistics while protecting individuals' privacy is being investigated (*57*). But these technologies are unlikely to fulfill their promise unless they are embedded at the core of our information infrastructures *by default*. And to achieve that, we need first to resolve through policy an inception problem: in absence of policy intervention supporting their deployment, isolated individual market agents may act, but the incentives for the vast array of industry players who control individuals' data may be insufficient to reach a critical mass of adoption (*26; 58*). Returning to the car safety analogy, the externalities of privacy call for regulatory intervention similarly to how externalities arising from unsafe driving led to policy and technological responses.

Privacy legal scholarship has also evolved with the times, expanding the notion of privacy protection from mere control over data flows to encompass issues of autonomy and protection from bias (*61*), and proposing novel approaches, such as construing data holders as data "fiduciaries" who hold legal obligations to act in the best interests of their customers (*62*). Those efforts are promising. In our view, any regulatory effort genuinely intent on addressing the novel challenges of privacy will have to shift away from an ineffectual focus on notice and consent mechanisms, which are easily gamed (*42, 43*), and at best provide necessary but insufficient conditions for privacy management. The problems with recent legislations like GDPR make the limitations of this approach abundantly clear, with consumers fatigued by the constant clicking of buttons to waive privacy rights. Instead, we argue that regulation concentrate on promoting impact analysis, efficient development, and ultimately mass-scale deployment of privacy technology. Mandating standards of compliance with user-centered privacy technologies in products and services (including intelligent agents representing user preferences) and mandating deployment of privacy-preserving, analytics-retaining algorithms among data holders should aim at minimizing user burden and the likelihood of coercion and manipulation.

The history of privacy tells us that the drive to maintain a private space may be as universal as the drive to commune (and that the two drives are in fact intimately related); and that humans invariably attempt to carve out those spaces even when odds are stacked against them by surveillance, be it digital (*63*) or physical (*64*). The reason why concerns over privacy endure, despite privacy being repeatedly declared dead, may be in part cultural and in part related to the visceral, evolutionary roots we discuss in the article. The current state of privacy also tells us, however, that those spaces have become unquestionably harder for individuals to manage. Solutions that predominantly rely on notice and consent mechanisms are inadequate—because they fail to take into account the visceral roots of our sense of privacy, and thus can be easily gamed by platforms and service providers. Understanding and then accounting for those ancestral roots of privacy may be critical to secure its future.

References and Notes

1. V. Cerf, Speech at the Federal Trade Commission Internet of Things Workshop, Washington, D.C., 19 November 2013 (https://www.ftc.gov/sites/default/files/documents/public_events/internet-things-privacy-security-connected-world/final_transcript.pdf).
2. B. Moore Jr, *Privacy: Studies in Social and Cultural History: Studies in Social and Cultural History* (Sharpe, 1984).
3. P. Aries, G. Duby, *A History of Private Life: From Pagan Rome to Byzantium*. (Belknap Press, 1987), Vol. 1.
4. T.A. Gregor, Exposure and seclusion: A study of institutionalized isolation among the Mehinacu Indians of Brazil. *Ethnology*. **9**, 234-250 (1970).
5. A. Westin, *Privacy and Freedom* (Simon & Schuster, 1967).
6. A. Acquisti, L. Brandimarte, G. Loewenstein, Privacy and human behavior in the age of information. *Science* **347**(6221), 509-514 (2015).
7. I. Altman, Privacy regulation: Culturally universal or culturally specific? *Journal of Social Issues.* **33**(3), 66-84 (1977).
8. P.H. Klopfer, D.I. Rubenstein, The concept privacy and its biological basis. *Journal of Social Issues*, **33**(3), 52–65 (1977).
9. S.L. Neuberg, M. Schaller, Evolutionary social cognition. M. Mikulincer, P. R. Shaver, E. Borgida, & J. A. Bargh Eds., (*APA Handbook of personality and social psychology, Vol. 1. Attitudes and social cognition),* pp. 3–45. (American Psychological Association, 2015).
10. D.M. Buss, The Evolutionary Psychology of Human Social Strategies. E. T. Higgins & A. W. Kruglanski Eds., (*Social Psychology: Handbook of Basic Principles*), pp. 3–38. (Guilford Press, 1996).
11. J.H. Barkow, *Darwin, Sex, and Status: Biological Approaches to Mind and Culture* (University of Toronto Press, 1989).
12. B.R. Schlenker, *Impression Management: The Self-Concept, Social Identity, and Interpersonal Relations* (Brooks-Cole, Monterey, CA, 1980).

13. R.F. Baumeister, M.R. Leary, The need to belong: desire for interpersonal attachments as a fundamental human motivation. *Psychological Bulletin*. ***117***(3), 497 (1995).
14. E. Goffman, *The Presentation of Self in Everyday Life* (London: Harmondsworth 1978).
15. I. Altman, *The Environment and Social Behavior: Privacy, Personal Space, Territory, and Crowding*. (Brooks/Cole Publishing Company, Monterey, CA, 1975).
16. R.F. Hixson, *Privacy in a Public Society* (Oxford University Press 1987).
17. D.J. Solove, A taxonomy of privacy. *University of Pennsylvania Law Review.* **154**(3), 477-560 (2005).
18. K. Kry, R. Casey, The effect of hiding enrichment on stress levels and behaviour of domestic cats (Felis sylvestris catus) in a shelter setting and the implications for adoption potential. *Animal Welfare*, **16**(3), pp. 375-383 (2007).
19. F. De Waal, *Chimpanzee Politics: Power and Sex Among Apes* (JHU Press, 2007).
20. Klopfer and Rubenstein (8, p. 64) conclude: "In sum, there are innumerable instances of nonhuman animals acting in a manner that characterizes humans seeking privacy."
21. H. R. Varian, Economic Aspects of Personal Privacy. *Privacy and Self-Regulation In The Information Age*, (National Telecommunications and Information Administration, US Department of Commerce, 1996).
22. A. Acquisti, L. Brandimarte, J.T. Hancock, Online Self-Disclosure and Offline Threat Detection. *Workshop on the Economics of Information Security*, Delft, The Netherlands, 2015.
23. K. Sun, C. Sugatan, T. Afnan, H. Simon, S. Gelman, J. Radesky, F. Schaub, "They See You're a Girl if You Pick a Pink Robot with a Skirt": A Qualitative Study of How Children Conceptualize Data Processing and Digital Privacy Risks. *Proceedings of the 2021 ACM Conference on Computer-Human Interaction*, Virtual, May 2021.
24. G. Conti, E. Sobiesk, Malicious interface design: exploiting the user. *Proceedings of the 19th International Conference on World Wide Web*, Raleigh, NC, April 2010.
25. A. Acquisti, I. Adjerid, R. Balebako, L. Brandimarte, L.F. Cranor, S. Komanduri, P.G. Leon, N. Sadeh, F. Schaub, M. Sleeper, Y. Wang, S. Wilson, Nudges for privacy and security: Understanding and assisting users' choices online. *ACM Computing Surveys*. **50**(3), 1-41 (2017).
26. A. Acquisti, L. Brandimarte, G. Loewenstein, Secrets and Likes: The Drive for Privacy and the Difficulty of Achieving It in the Digital Age. *Journal of Consumer Psychology* **30**(4), 736-758 (2020).
27. A. Acquisti, L. Brandimarte, J.T. Hancock, Are There Evolutionary Roots To Privacy Concerns? *Privacy Law Scholars Conference*, Berkeley, CA, 2013.
28. A. Shariff, J. Green, W. Jettinghoff, The Privacy Mismatch: Evolved Intuitions in a Digital World. *Current Directions in Psychological Science*. (*Forthcoming*).
29. J.E. Riggs, Stone-age genes and modern lifestyle: evolutionary mismatch or differential survival bias. *Journal of Clinical Epidemiology*. **46**(11), 1289-1291 (1993).
30. A.C. Logan, F.N. Jacka, Nutritional Psychiatry Research: An Emerging Discipline and Its Intersection With Global Urbanization, Environmental Challenges And The Evolutionary Mismatch. *Journal of Physiological Anthropology*. **33**(1), 1-16 (2014).
31. A. Acquisti, C.R. Taylor, L. Wagman, Liad, The Economics of Privacy. *Journal of Economic Literature*. 52(2), 442-492 (2016).
32. E. Pariser, *The Filter Bubble: How the New Personalized Web Is Changing What We Read and How We Think* (Penguin, 2011).

33. S. Vosoughi, D. Roy, S. Aral. The spread of true and false news online. *Science*. **359**(6380), 1146-1151 (2018).
34. C.A. Bail, L.P. Argyle, T.W. Brown, J.P. Bumpus, H. Chen, M.F. Hunzaker, J. Lee, M. Mann, F. Merhout, A. Volfovsky. Exposure to opposing views on social media can increase political polarization. *Proceedings of the National Academy of Sciences*. ***115***(37), 9216-9221 (2018).
35. R.E. Levy, Social media, news consumption, and polarization: Evidence from a field experiment. *American Economic Review.* **111**(3), 831-70 (2021).
36. S. Aral, D. Eckles, Protecting elections from social media manipulation. *Science.* **365**(6456), 858-861 (2019).
37. S. Eckert, J. Metzger-Riftkin, Doxxing. K. Ross, I. Bachmann, V. Cardo, S. Moorti & M. Scarcelli Eds., (*The International Encyclopedia of Gender, Media, and Communication*). (Wiley, 2020).
38. A. Warofka, “An Independent Assessment of the Human Rights Impact of Facebook in Myanmar”. *Facebook,* (2018; revised 2020; https://about.fb.com/news/2018/11/myanmar-hria/).
39. D.J. Solove, Introduction: Privacy self-management and the consent dilemma. *Harvard Law Review*. **126,** 1880-1903 (2012).
40. L. Brandimarte, A. Acquisti, G. Loewenstein, Misplaced confidences: Privacy and the control paradox. *Social Psychological and Personality Science*. **4**(3), 340-347 (2013).
41. M. Giesler, E. Veresiu, Creating the responsible consumer: Moralistic governance regimes and consumer subjectivity. *Journal of Consumer Research*. **41**(3), 840-857 (2014).
42. C. Utz, M. Degeling, S. Fahl, F. Schaub, T. Holz, (Un) informed consent: Studying GDPR consent notices in the field. *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, London, UK, November 2019.
43. M. Nouwens, I. Liccardi, M. Veale, D. Karger, L. Kagal, Dark patterns after the GDPR: Scraping consent pop-ups and demonstrating their influence. *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems,* Honolulu, HI, April 2020.
44. R. Calo, Against notice skepticism in privacy (and elsewhere). *Notre Dame Law Review*. **87**(3), 1027-1072 (2011).
45. H. Almuhimedi, F. Schaub, N. Sadeh, I. Adjerid, A. Acquisti, J. Gluck, L. Cranor, Y. Agarwal, Your location has been shared 5,398 times!: A field study on mobile app privacy nudging. *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, 2015.
46. A. Oulasvirta, A. Pihlajamaa, J. Perkiö, D. Ray, T. Vähäkangas, T. Hasu, N. Vainio, P. Myllymäki, Long-term effects of ubiquitous surveillance in the home. *Proceedings of the 2012 ACM Conference on Ubiquitous Computing,* Pittsburgh, PA, September 2012.
47. We thank Professor Neil Anthony Lewis for the formulation.
48. K.C. Laudon, Markets and privacy. *Communications of the ACM*. **39**(9), 92-104 (1996).
49. I. Arrieta-Ibarra, L. Goff, D. Jiménez-Hernández, J. Lanier, E.G. Weyl, Should we treat data as labor? Moving beyond" free". *AEA Papers and Proceedings*. **108** 38-42 (2018).
50. P.G. Kelley, J. Bresee, L.F. Cranor, R.W. Reeder, A "nutrition label" for privacy, *Proceedings of the 5th Symposium on Usable Privacy and Security*, Mountain View, CA, July 2009.
51. M.R. Lemov, *Car Safety Wars: One Hundred Years of Technology, Politics, and Death* (Rowman & Littlefield, 2015).

52. Organisation for Economic Co-operation and Development (OECD), "OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data", (1980; revised in 2013; http://www.oecd.org/sti/ieconomy/oecdguidelinesontheprotectionofprivacyandtransborderflowsofpersonaldata.htm).
53. R. Mendes, J.P. Vilela, Privacy-preserving data mining: methods, metrics, and applications. *IEEE Access*, **5**, 10562-10582 (2017).
54. A. Das, M. Degeling, D. Smullen, N. Sadeh, Personalized Privacy Assistants for The Internet of Things: Providing Users With Notice and Choice. *IEEE Pervasive Computing*. **17**(3), 35-46 (2018).
55. A. Grotto, R.E. Guadagno, J.T. Hancock, H. Mieczkowski, Applying the Principal-Agent Framework to AI-Mediated Communication: Psychological Implications, *AAAI-21 Workshop on AI For Behavior Change*, Palo Alto, CA, February 2021.
56. A. Cavoukian, Privacy by Design. *IEEE Technology and Society Magazine*. **31**(4), 18-19 (2012).
57. J.M. Abowd, I.M. Schmutte, An Economic Analysis of Privacy Protection and Statistical Accuracy as Social Choices. *American Economic Review*. 109(1), 171-202 (2019)
58. See, on the debate around AI and data-driven algorithmic personalization: "Facebook leadership has also repeatedly weakened or halted many initiatives meant to clean up misinformation on the platform because doing so would undermine that growth" (*59*). Also, consider Apple's roll out of an App Tracking Transparency (ATT) system, which allows individuals to opt out of being tracked by apps on their phone, and the backlash it has received from the advertising industry (*60*).
59. K. Hao, How Facebook got addicted to spreading misinformation, *MIT Technology Review*, (2021; https://www.technologyreview.com/2021/03/11/1020600/facebook-responsible-ai-misinformation).
60. K. Wagner. "Facebook Users Said No to Tracking. Now Advertisers are Panicking". *Bloomberg Technology*, *Bloomberg, (July 2021, https://www.bloomberg.com/news/articles/2021-07-14/facebook-fb-advertisers-impacted-by-apple-aapl-privacy-ios-14-changes).*
61. D.K. Mulligan, C. Koopman, N. Doty, Privacy is an Essentially Contested Concept: A Multi-Dimensional Analytic for Mapping Privacy. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*. 374(2083), 1-17 (2016).
62. J.M. Balkin, Information Fiduciaries and the First Amendment. *University California, Davis Law Review*. 49, 1185-1232 (2015).
63. D. Boyd, A.E. Marwick, Social privacy in networked publics: Teens' attitudes, practices, and strategies. *A Decade in Internet Time: Symposium on the Dynamics of the Internet and Society*. University of Oxford, September 2011.
64. P. Betts, *Within Walls: Private Life in The German Democratic Republic* (Oxford University Press, 2010).

Acknowledgements:

Funding:

The authors gratefully acknowledge support from the National Science Foundation through Awards 1228857 (Evolutionary Approaches to Privacy and Information Security, 2012) and #1514192 (Understanding and Exploiting Visceral Roots of Privacy and Security Concerns, 2015).

Author contributions:

Conceptualization and Writing: AA, LB, JH

Competing interests:

Authors declare that they have no competing interests.