

Turning Privacy Inside Out

*Julie E. Cohen**

The problem of theorizing privacy moves on two levels, the first consisting of an inadequate conceptual vocabulary and the second consisting of an inadequate institutional grammar. Privacy rights are supposed to protect individual subjects, and so conventional ways of understanding privacy are subject-centered, but subject-centered approaches to theorizing privacy also wrestle with deeply embedded contradictions. And privacy's most enduring institutional failure modes flow from its insistence on placing the individual and individualized control at the center. Strategies for rescuing privacy from irrelevance involve inverting both established ways of talking about privacy rights and established conventions for designing institutions to protect them. In terms of theory, turning privacy inside out entails focusing on the conditions that are needed to produce sufficiently private and privacy-valuing subjects. Institutionally, turning privacy inside out entails focusing on the design, production, and operational practices most likely to instantiate and preserve those conditions.

INTRODUCTION

The problem of theorizing privacy moves on two levels, the first consisting of an inadequate conceptual vocabulary and the second consisting of an inadequate institutional grammar. Theories about privacy have a tendency to dissolve into contradictions. So, for example, one justification commonly asserted for privacy is that it promotes and protects individual autonomy, but making privacy serve autonomy effectively is impossible unless one confronts the constructedness of selfhood. Another common justification

* Mark Claster Mamolen Professor of Law and Technology, Georgetown Law. My thanks to participants in the Cegla Center Symposium on the Problem of Theorizing Privacy and the 2018 Privacy Law Scholars Conference for their very helpful comments. Cite as: Julie E. Cohen, *Turning Privacy Inside Out*, 20 THEORETICAL INQUIRIES L. 1 (2019).

for privacy is that it promotes and protects an essential degree of separation between self and society. That justification is implicitly predicated on the reality of social construction, but making privacy serve the construction of selfhood effectively is impossible unless one confronts privacy's social (*i.e.*, collective) value. A second layer of contradictions emerges in the prevailing formulations of legal entitlements and instruments intended to vindicate privacy rights. Formulations of privacy in the liberty-based language of human rights discourse are both difficult to dispute and operationally meaningless. Policy instruments intended to have operational effect are couched in the language of granular control and have returned over and over to subject-centered constructs like notice and consent, even though such constructs are widely recognized as both unilluminating and impracticable in the face of inscrutable, machine learning-driven algorithmic mediation.

From a different perspective, though, privacy's embedded contradictions represent a rich source of opportunity. In earlier work, I have characterized privacy as a paradigmatic information-era right, as it both exposes important shortcomings of now-conventional forms of rights discourse and points the way toward strategies for developing new forms better tailored to the political economy of informationalism.¹ Those strategies involve inverting both established ways of talking about privacy rights and established conventions for designing institutions to protect them — *i.e.*, turning privacy inside out.

A useful point of entry for the project of rescuing privacy theory is the metaphoric relation between the figure and the ground as used in cognitive theory. From a very young age, human beings learn to parse complex images to identify patterns. That process entails judgments about the interrelationship and importance of different elements — about which parts of the image constitute the relevant subject and which are inessential background. Pattern-recognition processes are subconscious and near-instantaneous and produce perceptions that are highly durable, but they are also malleable. Learning to spot different patterns can lead an observer to locate figure and ground differently.² By “turning privacy inside out,” I mean to suggest approaching the two important conceptual dyads in privacy theory — the self/society relation and the self/materiality relation — using precisely that sort of reversal. In

1 This Article extends the argument sketched in Julie E. Cohen, *Affording Fundamental Rights: A Provocation Inspired by Mireille Hildebrandt*, 4 CRITICAL ANALYSIS L. 78 (2017).

2 For an overview of the accumulated learning on figure-ground perception, see Johan Wagemans et al., *A Century of Gestalt Psychology in Visual Perception: I. Perceptual Grouping and Figure–Ground Organization*, 138 PSYCHOL. BULL. 1172 (2012).

terms of theory, turning privacy inside out entails focusing on the conditions that are needed to produce sufficiently private and privacy-valuing subjects. In terms of institutions, turning privacy inside out entails focusing on the design, production, and operational practices best suited to instantiate and preserve those conditions.

Part I reviews the conceptual and operational contradictions that have bedeviled privacy theory and policymaking. Part II develops the basis for a theoretical approach to privacy that decenters subjects and foregrounds social and material conditions. Part III connects the theoretical to the operational, outlining an approach to privacy that decenters consent and foregrounds requirements of low flow and operational accountability.

I. PRIVACY'S CONCEPTUAL AND INSTITUTIONAL CONTRADICTIONS

Both conceptually and institutionally, privacy is a construct built on contradictions. When the principal justifications asserted for privacy are probed with any rigor, each rapidly dissolves into its opposite. Institutional arrangements devised to protect privacy, meanwhile, tend to be organized around subject-centered constructs that are either overly vague or impossible to implement from an operational perspective.

A. Theoretical Constructs Purporting to Justify Privacy

Perhaps the dominant justification for privacy is that it promotes and protects individual autonomy.³ But making privacy serve autonomy effectively is impossible unless one confronts the constructedness of selfhood. As I have explained in my earlier work, the condition of autonomy is contingent and ultimately paradoxical.⁴ Human beings do experience ourselves as having identities and making choices. But that experience — and hence our ultimate selfhood — is socially shaped in many important respects. Human subjectivity is malleable and emergent and embodied; it evolves as individuals and

3 For an especially rich statement of this view, see BEATE RÖSSLER, *THE VALUE OF PRIVACY* (2d ed. 2018). See also Anita L. Allen, *Coercing Privacy*, 40 WM. & MARY L. REV. 723, 738-40 (1999).

4 For an early version of this argument, still working within the “autonomy” framing, see Julie E. Cohen, *Examined Lives: Informational Privacy and the Subject as Object*, 52 STAN. L. REV. 1373 (2000). For a later version departing more definitely from that framing, see Julie E. Cohen, *What Privacy is for*, 126 HARV. L. REV. 1904 (2013) [hereinafter Cohen, *What Privacy is for*].

communities engage in practices of mutually constituting self-definition that are both culturally embedded and open-ended. Important work developing the concept of “relational autonomy” engages with this paradox to an extent.⁵ But the subjectivity-shaping processes to which I mean to refer are broader and deeper, and involve cultures, materialities, and ideologies.

The related justification for privacy in terms of dignity confronts a similar paradox.⁶ Like experienced subjectivity, conceptions of dignity are themselves culturally constructed. Even if one posits a decontextualized, universal starting point, such as the Kantian categorical imperative, matters rapidly become more complex. Different societies articulate and perform commitments to dignity differently — for example, by adopting different norms about the extent to which various activities and functions may be discussed or observed.

Another justification commonly asserted for privacy is that it promotes and protects an essential degree of separation between self and society that permits dissent and critique.⁷ That justification is implicitly predicated on a partial engagement with the reality of social construction. So, for example, the rubric of the “chilling effect,” which traces its lineage to Jeremy Bentham’s design for the Panopticon, is so powerful precisely because it accepts that observation shapes behavior.⁸ Some prominent scholars working within the chilling effects framework have avoided engaging with social construction by attempting to distinguish between behavioral conditioning and “real” subjectivity.⁹ But how are we to tell which is which? Other theories predicated on the reality of the chilling effect answer that question by positing commitments to shared cultural values such as critical independence of mind.¹⁰ The decision to privilege such values, however, reveals social construction at work. Why is

5 See RÖSSLER, *supra* note 3; JENNIFER NEDELSKY, *LAW’S RELATIONS: A RELATIONAL THEORY OF SELF, AUTONOMY AND LAW* (2011).

6 See David Matheson, *Dignity and Selective Self-Presentation*, in *LESSONS FROM THE IDENTITY TRAIL* 319 (Ian Kerr et al. eds., 2009).

7 See, e.g., NEIL RICHARDS, *INTELLECTUAL PRIVACY: RETHINKING CIVIL LIBERTIES IN THE DIGITAL AGE* 103-08 (2015). See also Cohen, *What Privacy is for*, *supra* note 4, at 1917-18 (relating privacy to critical citizenship).

8 See JEREMY BENTHAM, *PANOPTICON OR THE INSPECTION HOUSE* (Dublin, Thomas Byrne 1787). See also MICHEL FOUCAULT, *DISCIPLINE AND PUNISH: THE BIRTH OF THE PRISON* (Alan Sheridan trans., Vintage Books 2d ed. 1995) (1977).

9 See, e.g., JEFFREY ROSEN, *THE UNWANTED GAZE: THE DESTRUCTION OF PRIVACY IN AMERICA* 166 (2000). See also RICHARDS, *supra* note 7, at 103-08 (2015) (positioning the patterns of intellectual activity that preexist surveillance as individual and authentic).

10 See, e.g., RICHARDS, *supra* note 7, at 96-103.

separation between self and society so essential, if not that it serves important values we as a society have chosen to privilege?

Justifications for privacy in terms of its social value, meanwhile, cannot explain why we should care about those gains without referring to conceptions of the self.¹¹ Why, for example, should we not subject job applicants to a requirement of full and complete disclosure about every foible and failing? One might attempt to answer that question by arguing that a full-disclosure requirement would foreclose the welfare gains to be realized from allowing all people to contribute to the full extent of their abilities. Why, though, do we value such contributions in the first place? Jobs and other resources are scarce and few potential candidates are truly unique; particularly as tools like personality testing and forecasts of criminal recidivism become more precise, why not use them as triage? One common answer is to duck the question by pointing out all of the ways that such tools are flawed and discriminatory in practice. (To be clear, such objections do not simply make the perfect the enemy of the good; they are significant and often deservedly fatal. But my point in this thought experiment is different.) On the assumption, though, that the defects in predictive frameworks could be cured or minimized, why should not such methods be used to optimize hiring? If one resists that conclusion, it can only be because of beliefs about the dignity and moral worth of human beings.

Legal scholarship about privacy has struggled to embrace these contradictions, returning over and over to assertions about autonomy, chilling effects, and welfare-enhancing tradeoffs as though such assertions were self-explanatory. But the contours of a right to privacy cannot be derived, like a sort of jurisprudential hypotenuse, from first-order philosophical commitments in a way that bypasses the need to make normative choices. Privacy must be chosen, and it is deeply intertwined with other societal commitments that also represent normative choices. I have discussed those reasons in other work.¹² In the balance of this Article, I treat the choice to value privacy — to prioritize the production of sufficiently private and privacy-valuing subjects — as having been made and ask whether, having so chosen, there are other kinds of conceptual gains to be made. To be sure, conceptual coherence isn't the sole or even the most important determinant of a functioning system of privacy protection.¹³ As one

11 For a representative selection of arguments about the social value of privacy, see generally *SOCIAL DIMENSIONS OF PRIVACY: INTERDISCIPLINARY PERSPECTIVES* (Beate Roessler & Dorota Mokrosinska eds., 2015).

12 See Cohen, *What Privacy is for*, *supra* note 4.

13 In legal discourses about privacy, the particular forms that the quest for theoretical consistency has assumed are artifacts of allegiance to particular philosophical

team of researchers has put it, privacy is an essentially contested concept and is profoundly generative for that reason.¹⁴ As we will see in Part II, however, more effective tools for theorizing privacy — tools that embrace and foreground privacy’s paradoxes and contradictions — do exist.

B. Institutional Arrangements Purporting to Safeguard Privacy

A second layer of contradictions emerges in the prevailing formulations of legal entitlements and instruments intended to vindicate privacy rights. Formulations of privacy in the liberty-based language of human rights discourse are grand, inspiring, and difficult to dispute but also operationally meaningless. For such formulations to bite meaningfully on the conduct of either governments or non-state entities, they must be translated into more specific mid-level rules. Policy instruments intended to have operational effect, however, have been largely ineffective in practice. This section summarizes the principal reasons for privacy’s operational failures.

The first and most important reason for failure is that notice-and-consent protections, which function as the principal regulatory tool in the U.S. system and as an increasingly important backstop in the European system, simply do not work.¹⁵ Meaningful consent requires meaningful notice, but the information provided about data collection, processing, and use tends to be vague and general. Equally important, such disclosures tend to conflate important distinctions between remembering users’ preferences, creating predictive profiles that may also include other, inferred data, using those preferences for targeted marketing, and tracking users across multiple websites, devices, and locations.¹⁶

and political traditions. See JULIE E. COHEN, *CONFIGURING THE NETWORKED SELF: LAW, CODE, AND THE PLAY OF EVERYDAY PRACTICE* 107-26 (2012) [hereinafter COHEN, *CONFIGURING THE NETWORKED SELF*].

14 Deirdre K. Mulligan, Colin Koopman & Nick Doty, *Privacy Is an Essentially Contested Concept: A Multi-Dimensional Analytic for Mapping Privacy*, 374 PHIL. TRANSACTIONS ROYAL SOC’Y. A 118 (2016), <http://rsta.royalsocietypublishing.org/content/374/2083/20160118>.

15 On the role and the impossibility of consent within the European system, see Bert-Jaap Koops, *The Trouble with European Data Protection Law*, 4 INT’L DATA PRIVACY L. 250 (2014); Alessandro Mantelero, *The Future of Consumer Data Protection in the E.U.: Rethinking the “Notice and Consent” Paradigm in the New Era of Predictive Analytics*, 30 COMPUTER L. & SECURITY REV. 643 (2014).

16 See Solon Barocas & Helen Nissenbaum, *Big Data’s End Run Around Anonymity and Consent*, in *PRIVACY, BIG DATA, AND THE PUBLIC GOOD: FRAMEWORKS FOR ENGAGEMENT* 44 (Helen Nissenbaum, Julia Lane & Victoria Stodden eds., 2014);

The range of potential future applications is most often left unspecified — often because the parties collecting and processing the information do not know and wish to leave their options open.

More basically, there is an intractable tension between disclosure requirements that aim to educate consumers and conventional wisdom about efficacy in marketing. Both marketing experts and consumer advocates have long recognized that even truthful disclosures about product quality and characteristics are easy to manipulate to induce consumers to buy; so too with disclosures about the collection, processing, and use of personal information that induce consumers to consent.¹⁷ The design of digital interactive environments introduces additional variables that can be adjusted to encourage both over-disclosure and broad forward-looking consent to processing and use.¹⁸ Recent revelations about politically motivated media manipulation have reminded us that the processes used to mediate access to news and information also play important roles in shaping what consumers think and believe and that opaque, advertiser-driven profit models do not produce informed publics.¹⁹

One way to address the notice failures that have become endemic in networked digital environments might be to require meaningful disclosures about the automated logics involved in processing personal information, including disclosures about both the kinds of information that such logics

Solon Barocas & Helen Nissenbaum, *On Notice: The Trouble with Notice and Consent*, in PROCEEDINGS OF THE ENGAGING DATA FORUM: THE FIRST INTERNATIONAL FORUM ON THE APPLICATION AND MANAGEMENT OF PERSONAL ELECTRONIC INFORMATION (2009), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2567409.

- 17 On the various failures of consent resulting from incompleteness and/or manipulation of mandated disclosures, see Alessandro Acquisti et al., *The Economics of Privacy*, 54 J. ECON. LITERATURE 442, 442-43 (2016); Kirsten Martin, *Privacy Notices as Tabula Rasa: An Empirical Investigation into How Complying with a Privacy Notice Is Related to Meeting Privacy Expectations Online*, 34 J. PUB. POL'Y & MARKETING 210 (2015); Lauren E. Willis, *Performance-Based Consumer Law*, 82 U. CHI. L. REV. 1309, 1322-25 (2015); Lauren E. Willis, *When Nudges Fail: Slippery Defaults*, 80 U. CHI. L. REV. 1155, 1170-1200 (2013).
- 18 See WOODROW HARTZOG, *PRIVACY'S BLUEPRINT: THE BATTLE TO CONTROL THE DESIGN OF NEW TECHNOLOGIES* 21-55 (2018).
- 19 See Carole Cadwalladr, *'I Made Steve Bannon's Psychological Warfare Tool': Meet the Data War Whistleblower*, THE GUARDIAN (Mar. 18, 2018), <https://www.theguardian.com/news/2018/mar/17/data-war-whistleblower-christopher-wylie-faceook-nix-bannon-trump>; Mike Isaac & Daisuke Wakabayashi, *Russian Influence Reached 126 Million Through Facebook Alone*, N.Y. TIMES (Oct. 30, 2017), <https://www.nytimes.com/2017/10/30/technology/facebook-google-russia.html>.

treat as significant and the ways that the results of processing will be used. Although the matter is not free from doubt, some argue that the new European General Data Protection Regulation (GDPR) should be read to impose such arrangements in all cases where automated processing involved.²⁰

Scholars have raised important questions about whether it is possible to explain certain types of machine learning-driven processes at all and about whether such explanations, if available, constitute meaningful remedies for complaints that are, at bottom, complaints about unfair treatment.²¹ At minimum, it is clear that operationalizing a requirement of meaningful disclosure will require new kinds of tools for audit, explanation, and translation, a matter to which I return in Part III below. If a requirement of meaningful disclosure is to have any teeth, moreover, it needs to be accompanied by skepticism toward broad trade secrecy claims that would shield key information about logics and consequences from disclosure. Additionally, a requirement of meaningful disclosure can function as a regulatory lever only if it is backstopped by robust consumer protection laws guaranteeing individual consumers meaningful access to goods and services on acceptable terms even after they decline to provide certain items of personal information. In the U.S., at least, this last point conflicts with deeply ingrained reflexes about freedom of contract, and so substantive protections for consumer rights are increasingly rare.

More generally, it is not clear how much consumers would benefit from the opportunity to navigate an additional layer of complexity in aid of making wide-ranging and imperfectly informed decisions about the future. Defaulting to broad forward-looking consent — now deemed fully “informed” — may seem to many to be the best option. Put differently, when disclosure is in aid of a regulatory regime predicated on consent, the cure may be worse than the disease.

In European Union member states, the purpose limitation principle and the prohibition on processing certain categories of sensitive data purport to offer more powerful tools for constraining processing and use of already-collected data.²² Proponents of those tools, however, tend to engage in considerable overclaiming. To begin with, information businesses have little incentive to respect

20 See Council Regulation 2016/679 2016 O.J. (L 119) 1 (EU), arts. 13(2)(f), 14(2)(g), 15(1)(h), 22 [hereinafter GDPR]. For a useful overview of the debate on this point, see Andrew Selbst & Julia Powles, *Meaningful Information and the Right to an Explanation*, 7 INT’L DATA PRIVACY L. 233 (2017).

21 Lilian Edwards & Michael Veale, *Slave to the Algorithm? Why a ‘Right to an Explanation’ Is Probably Not the Remedy You are Looking for*, 16 DUKE L. & TECH. REV. 18 (2017); Andrew D. Selbst & Solon Barocas, *The Intuitive Appeal of Explainable Machines*, 86 FORDHAM L. REV. 1685 (2018).

22 See GDPR, *supra* note 20, art. 5(1)(b), 9(1).

such restrictions and tend to use broad consent provisions systematically as a way of circumventing them, which returns us to the problems just discussed. More importantly, powerful new techniques for data-intensive, machine learning-based analysis and prediction create difficult implementation challenges that repeated exhortations to respect the restrictions — such as those offered by the Article 29 Working Party in its most recent draft guidelines on automated profiling and its intersection with the principle of purpose limitation — do not address.²³

Machine learning-based predictive tools use existing data to infer missing data and to extrapolate predictions based on both the data that are known and those that have been inferred. Put differently, such systems are designed both to detect nonobvious patterns within masses of data and to work around constraints created by the absence of other data.²⁴ Because such systems are constantly and creatively seeking ways around experienced constraints, both prohibitions and permissions must be designed differently if they are to be effective. Consider first prohibitions. A system forbidden to use race as a variable may use other data, such as media consumption or purchase of hair care products, to infer race; and it might use factors that themselves reflect preexisting patterns of discrimination, such as lower scores on standardized tests or longer commuting distances to the site of a new job, as decision-making proxies.²⁵ If instructed to avoid race-based disparate impact (which must be computationally defined using parameters that cannot be set without collecting information about race in the first place), it may adjust by burdening a particular subgroup — for example, black men aged 18-25 or Asian Muslims — more heavily. Data-driven machine learning systems, in other words, may respect specific prohibitions directly and clearly expressed in code, but that respect will not necessarily translate into a more general orientation toward equal treatment of all persons. Eliminating or minimizing racially disparate

-
- 23 Article 29 Data Protection Working Party, *Guidelines on Automated Individual Decision-Making and Profiling for the Purpose of Regulation 2016/679*, at 9-15 (Oct. 3, 2017), http://ec.europa.eu/newsroom/article29/item-detail.cfm?item_id=612053. See also Judith Rauhofer, *Of Men and Mice: Should the EU Data Protection Authorities' Reaction to Google's New Privacy Policy Raise Concern for the Future of the Purpose Limitation Principle?*, 1 EUR. DATA PROTECTION L. REV. 5 (2015); Tal Z. Zarsky, *Incompatible: The GDPR in the Age of Big Data*, 47 SETON HALL L. REV. 995, 1004-14 (2017).
- 24 See David Lehr & Paul Ohm, *Playing with the Data: What Legal Scholars Should Learn about Machine Learning*, 51 U.C. DAVIS L. REV. 653, 670-72 (2017).
- 25 See generally Solon Barocas & Andrew D. Selbst, *Big Data's Disparate Impact*, 104 CALIF. L. REV. 671, 677-93 (2016); Pauline T. Kim, *Data-Driven Discrimination at Work*, 58 WM. & MARY L. REV. 857, 874-90 (2017).

results in data-driven machine learning environments is possible only if the analytic tools are subject to continual audit and retraining. Because that process necessarily entails making tradeoffs among different training parameters, it also requires articulating the various conceptions of fairness that might be employed, defining those conceptions computationally, and choosing which one(s) to prioritize.²⁶ The design of machine-learning processes also includes a number of other steps that entail value-laden choices and that the idea of prohibitions on certain uses does not capture.²⁷

Now consider strategies for subjecting machine learning-based analytics to purpose-limited permissions. Imagine, for example, that a grocery chain collects information about its customers' purchases. The purpose limitation principle would restrict future uses of that information to "compatible" uses, including perhaps future food-related marketing in the form of coupon discounts. To be effective, that restriction would need to travel alongside the information wherever and however it is stored, limiting all of the various inferences and correlations that might be drawn from it. Practically speaking, such proposals for pervasive and enduring "privacy as control" entail continuing and intensified surveillance simply to ensure compliance.²⁸ And, once again, because machine learning-based systems are constantly and creatively seeking ways around experienced constraints, it still will be extraordinarily difficult to predict and foreclose all of the third- or fourth-order inferences that might be drawn from the availability or absence of particular items of information. The problem is not that implementing an appropriate permission structure would be a complex and costly endeavor — there would be costs, but the same is true of many other safety-related design features that we have come to understand as essential. It is that when the "privacy as control" paradigm intersects with the problem of continually reoptimizing machine learning systems, structures for control will grow exponentially more complex, will entail rapid proliferation of internal surveillance functionality, and still will not work.

Finally, privacy protections also fail when decision-makers believe that using collected personal information to enable fine-grained determinations about access to certain important resources is in fact the fairest method.

26 See generally Richard Berk et al., *Fairness in Criminal Justice Risk Assessments: The State of the Art* (May 28, 2017) (unpublished manuscript), <https://arxiv.org/abs/1703.09207>.

27 See Lehr & Ohm, *supra* note 24, at 669-702.

28 See Rula Sayaf, Dave Clarke & James B. Rule, *The Other Side of Privacy: Surveillance in Data Control*, in *PROCEEDINGS OF THE 2015 BRITISH HCI CONFERENCE* 184 (Shaun Lawson ed., 2015).

Over the course of the twentieth century and continuing into the twenty-first, frameworks for making decisions about a wide range of important issues — credit, life insurance, health insurance and health policy, government benefits — have become both more actuarial and more computationally sophisticated. Although the rhetoric of information-based innovation holds out the promise of more perfect personalization for some of these decisions, the underlying mechanism — pattern-driven machine learning — remains fundamentally actuarial because it is correlation-based rather than causation-based. One kind of objection to such decision-making is that it cannot escape the influence of discriminatory rubrics that are deeply embedded in the data because they are deeply embedded in our society.²⁹ A more fundamental objection is that, whether or not particular predictions are accurate and regardless of why they might or might not be accurate, there is an unbridgeable gap between actuarial decision-making and fairness. Actuarial decision-making treats human beings as collections of data points; even when the data point toward leniency with respect to credit or employment, they do so in a way that is objectifying. If that is the case, though, too much personalization is as problematic as too little. Sometimes, we will need to be satisfied with less than fully individualized treatment simply because individualized decision-making does not scale, but it does not follow that highly granular actuarial decision-making will be the best proxy, and faith in actuarial decision-making has become so dominant that the question about other possible proxies simply has not received sufficient attention.

It is important to underscore that institutional failures of privacy protection are overdetermined. Data harvesting and processing are one of the principal business models of informational capitalism, so there is little motivation either to devise more effective methods of privacy regulation or to implement existing methods more rigorously. Instead, the cultural and political discourses that have emerged around data-centered “innovation” work to position such activities as virtuous and productive, and therefore ideally exempted from state control.³⁰ Strategies for addressing those problems are outside the scope of this Article. My point here is more basic: Should the political will for more effective privacy regulation be mustered, it still would be necessary to develop a new regulatory toolkit.

29 See generally Barocas & Selbst, *supra* note 25, at 677-93.

30 See generally Julie E. Cohen, *The Surveillance-Innovation Complex: The Irony of the Participatory Turn*, in *THE PARTICIPATORY CONDITION IN THE DIGITAL AGE* 207 (Darin M. Barney et al. eds., 2016).

II. CONSTRUCTING AN INSIDE-OUT THEORY OF PRIVACY

Although some have concluded that privacy is dead or on life support, the better answer is that the failures of privacy theory and privacy institutions also present new opportunities to get it right. However, the projects of theorizing privacy properly — by which I mean in a way that acknowledges and incorporates its contradictions — and then operationalizing the results of that inquiry require a series of methodological inversions. In terms of theory, turning privacy inside out entails consciously abandoning theories organized around the presumptive autonomy of selves — the *figures* on whom privacy theory overwhelmingly has focused — and focusing instead on the conditions necessary to produce sufficiently private and privacy-valuing subjects.

Consider first the self-society relation, with all of its seemingly insuperable contradictions. As suggested in Part I.A above, processes of self-development do not conform to the idealized theoretical models preferred by liberal legal theorists, which revolve around the purposive exercise of expressive or market liberty.³¹ Selfhood is a product of both social shaping and embodied experience. People are born into networks of relationships, practices, and beliefs, and those networks profoundly shape the processes of self-articulation. Selfhood is also and importantly a product of serendipity. People find ways to push back against the particular institutional, cultural, and material constraints that they encounter in their everyday lives. In addition, they exploit the unexpected encounters and juxtapositions that everyday life supplies. The fact that processes of self-articulation defy neat theoretical simplification via the usual methods, however, does not mean there is nothing useful to say about them or about the role(s) that privacy plays. It simply counsels more careful attention to the social and environmental factors more conventionally understood as background.

To understand the relationship between privacy and self-articulation, it is useful to contrast surveillance and privacy as modes of social ordering. Surveillance — defined generically as attention that is purposeful, routine, systematic, and focused³² — is a mode of ordering predominantly concerned with producing predictable, rationalized behaviors and information flows. Privacy is a dynamic condition that is best described as breathing room for socially situated subjects to engage in processes of boundary management

31 For more detailed discussions of the misalignment, see COHEN, CONFIGURING THE NETWORKED SELF, *supra* note 13, at 107-26; Cohen, *What Privacy is for*, *supra* note 4, at 1906-11.

32 KIRSTIE BALL ET AL., SURVEILLANCE STUDIES NETWORK: A REPORT ON THE SURVEILLANCE SOCIETY 8 (David Murakami Wood ed., 2006).

through which they define and redefine themselves as subjects.³³ Systematic surveillance reduces breathing room; privacy preserves (degrees of) spatial, informational, and epistemological open-endedness.³⁴

Put differently, privacy is about both boundaries and boundary-crossings, and that perspective yields improved analytical purchase on the problem of automated data processing's predictable yet still surprising "creepiness." Advocates of data processing have argued that objections framed in terms of creepiness are both insufficiently rigorous to count analytically and insufficiently durable to matter politically.³⁵ According to the former argument, feelings of unease that defy reduction to more concrete terms have no place in policy debates; according to the latter, the feelings will recede as we become accustomed to the new functionalities and the transformative abilities they promise. But the problem of creepiness — or, in some literatures, the "uncanny valley" between the natural and the artificial³⁶ — has a rigor that sounds in cognitive science rather than philosophy or economics, and creepiness is not a result of uncertainty about threats but rather of uncertainty about categories. Phenomena that are creepy are those that transgress the boundary between animate and inanimate, violating principles of essentialism that are cognitively determined.³⁷ Like the replicants and Terminators of science fiction, machine-learning predictive analytics are non-natural but challenge that categorization

33 See COHEN, CONFIGURING THE NETWORKED SELF, *supra* note 13, at 148-52. See also Valerie Steeves, *Reclaiming the Social Value of Privacy*, in LESSONS FROM THE IDENTITY TRAIL 191 (Ian Kerr et al. eds., 2009).

34 On the contrasts between surveillance and privacy as modes of social ordering, see Julie E. Cohen, *Surveillance vs. Privacy: Effects and Implications*, in THE CAMBRIDGE HANDBOOK OF SURVEILLANCE LAW 455 (David Gray & Stephen E. Henderson eds., 2017).

35 See, e.g., Larry Downes, *A Rational Response to the Privacy "Crisis"* 26-31 (Cato Inst., Policy Analysis No. 716, 2013), <http://www.cato.org/sites/cato.org/files/pubs/pdf/pa716.pdf>; Adam Thierer, *The Pursuit of Privacy in a World Where Information Control Is Failing*, 36 HARV. J.L. PUB. POL'Y 409, 417-21 (2013). See also Omer Tene & Jules Polonetsky, *A Theory of Creepy: Technology, Privacy, and Shifting Social Norms*, 16 YALE J.L. & TECH. 59, 71 (2013) ("Naturally, identifying creep is more an art than a science. Hence, inductive reasoning based on anecdotal evidence may be the best way forward in theorizing this term.").

36 See Masahiro Mori, Karl F. MacDorman & Norri Kageki, *The Uncanny Valley*, 19 IEEE ROBOTICS & AUTOMATION MAG. 98, 98-100 (2012), <https://doi.org/10.1109/MRA.2012.2192811> (reprinting and translating Mori's original Japanese-language essay).

37 See David Livingstone Smith, *Paradoxes of Dehumanization*, 42 SOC. THEORY & PRAC. 416, 430-33 (2016).

by behaving in ways that seem to cross the animate/inanimate divide. As such systems begin to learn and respond autonomously, they join lizard-headed dogs and other inhabitants of nightmares in the uncanny valley not because they are unfamiliar but rather because they are simultaneously apparent and impossible.

So understood, the sensation of *creepiness* is trying to tell us something about the destabilizing nature of *certain types* of boundary violations, and that message is inescapably relevant to privacy. The problem is not that automated anticipation of our every want and need turns out not to be what we wanted after all. It is not even that automated anticipation of our inferred wants and needs in the service of extractive logics embeds us in a set of social and commercial relations that are not of our choosing — that is also true, but most people do not reason about their own experienced reality in such ways. Rather, it is that on the most fundamental level our own self-consciousness depends on interactions with human others to work properly.

Selfhood is a process, not a state, and that process is discursive and social; it is informed by a sense of the self as viewed from the perspective of others.³⁸ Interactions with automated logics disrupt processes of self-formation because the others whose perspective must be assimilated are so alien that their perspective cannot be imagined. Surrounded by rapidly proliferating simulacra of animate, human intelligences that are intermittently revealed to be lizard-headed dogs, phylogenetically speaking, we detect the pervasive and powerful operation of an alien rationality that does not appear to be in sympathy with or in aid of humanity at all. Two examples from science fiction can usefully underscore the point. The poster child for technological creepiness is HAL, the renegade operating system for the spacecraft in *2001: A Space Odyssey*, which operates according to an internally coherent rationality of its own; its opposite number is Star Trek's "computer," which is always named and addressed in a way that reaffirms its inalterable otherness and its status as subordinate to human goals and needs. By contrast, today's artificial intelligence-based, virtual assistants have names like Siri and Alexa and, like HAL, are named and addressed as singular selves. They exemplify design and marketing strategies that seek to mask and domesticate (and gender) their

38 Many roads lead to this conclusion. See, e.g., COHEN, CONFIGURING THE NETWORKED SELF, *supra* note 13, at 129-30 (discussing self-performance for imagined publics); Mireille Hildebrandt, *Privacy as Protection of the Incomputable Self: From Agnostic to Agonistic Machine Learning*, 20 THEORETICAL INQUIRIES L. 83 (2019) (framing self-perception as shaped by the hermeneutics of language and text, which enable us to posit both "I" and "you"/"them").

alien rationalities, even though the idea of a singular Siri or Alexa who exists for us alone is an insupportable fiction.³⁹

The second needed reversal of the figure and the ground in privacy theory involves the relationship between selfhood and materiality. Until relatively recently, privacy theory, like rights discourse more generally, has operated with a set of unstated and often unexamined assumptions about the material environment's properties — assumptions both about constraint (e.g., the physical impossibility of universal surveillance) and about lack of constraint (e.g., the correspondingly open-ended possibilities for constructing arrangements characterized by privacy and/or confidentiality). Those assumptions have enabled materiality to fade into the background. Advances in networked digital communication and information processing have drawn that approach into question, making clear that it is a mistake to take materiality for granted. Yet, just as with the self-society relation, the project of foregrounding materiality within privacy theory remains incomplete.

What I mean by this is that *acknowledging* the relevance of materiality to theoretical understandings of privacy is not the same thing as *incorporating* materiality-based considerations within theoretical frameworks. The emergence of networked digital information and communications technologies and the varying affordances of such technologies — for both expression and control of expression and for both enhanced privacy and enhanced surveillance — have elicited an outpouring of scholarship and have prompted the United Nations to commission a series of special investigations and reports.⁴⁰ But

39 On the gendering of virtual digital assistants, see Ian Bogost, *Sorry, Alexa is Not a Feminist*, THE ATLANTIC (Jan. 24, 2018), <https://www.theatlantic.com/technology/archive/2018/01/sorry-alexa-is-not-a-feminist/551291/>; Adrienne LaFrance, *Why Do so Many Digital Assistants Have Feminine Names?*, THE ATLANTIC (Mar. 30, 2016), <https://www.theatlantic.com/technology/archive/2016/03/why-do-so-many-digital-assistants-have-feminine-names/475884/>. On the Star Trek episode in which the computer gained HAL-like self-awareness, with predictably calamitous consequences, see *The Ultimate Computer*, WIKIPEDIA, https://en.wikipedia.org/wiki/The_Ultimate_Computer (last visited May 19, 2018).

40 See, e.g., Ben Emmerson (Special Rapporteur on the Promotion and Protection of Human Rights and Fundamental Freedoms while Countering Terrorism), *Fifth Rep. on the Promotion and Protection of Human Rights and Fundamental Freedoms While Countering Terrorism*, General Assembly, U.N. Doc. A/70/371 (Sept. 18, 2015); David Kaye (Special Rapporteur on the Promotion and Protection of the Right to Freedom of Opinion and Expression) *Rep. on the Promotion and Protection of the Right to Freedom of Opinion and Expression*, Human Rights Council, U.N. Doc. A/HRC/29/32 (May 22, 2015); Ben Emmerson (Special

there does not yet seem to be any serious discussion about how to construct a vernacular for rights discourse that would incorporate notions of constraint and affordance as core conceptual building blocks. That omission should prompt reconsideration of prevailing approaches to theorizing fundamental rights generally, but nowhere is that need more urgent than it is for the project of theorizing privacy.

Consider the example of another important recent reorientation within fundamental rights theory. Fundamental rights to political and economic self-determination are made available, partly by the content and institutional structure of the applicable legal regime but also partly by access to the resources and capabilities needed to equip people to exercise their rights fully and effectively. Growing recognition of the importance of *capabilities for human flourishing* — understood to encompass the resources required not only for physical wellbeing but also for intellectual, cultural and political participation and self-determination — ultimately engendered a reorientation of rights discourse on the level of theory.⁴¹ Now we have come to recognize that fundamental rights also are made available, partly by the constraints and affordances of the physical environment, and that recognition too frames an important choice on the level of theory. When our background assumptions about the material environment fail to hold, we can choose to tolerate a basic level of hypocrisy about the conditions of possibility for, *e.g.*, privacy or self-expression (as is the case with liberty-based rights discourse that ignores the

Rapporteur on the Promotion and Protection of Human Rights and Fundamental Freedoms while Countering Terrorism) *Fourth Rep. on the Promotion and Protection of Human Rights and Fundamental Freedoms While Countering Terrorism*, General Assembly, U.N. Doc. A/69/397 (Sept. 23, 2014); Frank La Rue (Special Rapporteur on the Promotion and Protection of the Right to Freedom of Opinion and Expression), *Rep. on the Promotion and Protection of the Right to Freedom of Opinion and Expression*, Human Rights Council, U.N. Doc. A/HRC/23/40 (Apr. 17, 2013); Frank La Rue (Special Rapporteur on the Promotion and Protection of the Right to Freedom of Opinion and Expression), *Rep. on the Promotion and Protection of the Right to Freedom of Opinion and Expression*, Human Rights Council, U.N. Doc. A/HRC/17/27 (May 16, 2011); Martin Scheinin (Special Rapporteur on Human Rights and Counter-Terrorism), *Rep. on the Promotion and Protection of Human Rights and Fundamental Freedoms While Countering Terrorism*, Human Rights Council, U.N. Doc. A/HRC/13/37 (Dec. 28, 2009).

41 See AMARTYA SEN, DEVELOPMENT AS FREEDOM (2d. ed. 2001); Amartya Sen, *Elements of a Theory of Human Rights*, 32 PHIL. & PUB. AFF. 315 (2004); MARTHA C. NUSSBAUM, CREATING CAPABILITIES: THE HUMAN DEVELOPMENT APPROACH 31-36 (2011).

problem of capabilities in an era of vast and growing economic inequality), or we can extend rights discourse — and, along with it, our notions of what counts as a theory about rights — into the realm of the material.

Because the very idea of a materiality-based rights discourse is both unfamiliar and deeply unorthodox, it is worth spending a moment on terminology. By *affordance*, I mean to refer to the concept developed by environmental psychologist James Gibson as it has been translated into technology studies through Donald Norman's work on design but then to play a bit with that concept's implicitly individualized framing.⁴² Gibson coined the term "affordance" to refer to the enabling properties of physical environments and more specifically to particular kinds and ways of enablement whether or not such enablement is consciously apprehended or remarked. So, for example, a meadow is walk-able or run-able and perhaps eat-able but a lake is not; it is swim-able and perhaps drink-able. An organism may be optimized to take advantage of the affordances of a meadow or a lake, or it may not, which is another way of saying that affordance inheres in the relationship between environment and organism rather than being a separate property of either one.

By analogy, an artifact's affordances are the kinds and ways of uses that it enables whether or not such enablement is consciously apprehended or remarked; for example, a hammer is bang-able while a knife is slice-able or stab-able; a park bench is sit-able or sleep-able; and a mobile phone is pocket-able and connect-able and confers locate-ability (in multiple senses) on the one who carries it. Unlike natural environments, however, built environments and artifacts may also be designed to disafford certain uses, either by prohibiting them outright or employing other strategies to minimize disfavored uses. A park bench is sleep-able only if fixed dividers have not been added and it is less sleep-able if it has been contoured so that sleepers will roll off onto the ground. A mobile phone may disafford locational privacy in order to afford tracking, and that result may be hard-coded or it may be achieved via an interface design that makes privacy-enhancing configurations difficult to discover and implement. (One might therefore say that park benches and mobile phones "regulate" their users, but a comprehensive understanding of human-artifact relations must account for enablement and co-construction as well as prohibition.)

Critically, the idea of an affordance does not reduce either to liberty (because affordances can also constrain) or to capability (because affordances need not translate into skill or improved flourishing); it is concerned simply with the range of uses that are possible. It also does not reduce either to the

42 JAMES J. GIBSON, *THE ECOLOGICAL APPROACH TO VISUAL PERCEPTION* 127-37 (1979); DONALD A. NORMAN, *THE DESIGN OF EVERYDAY THINGS* 9-11, 81-92 (1988).

kinds and ways of use that are perceivable or to the uses that an artifact's design suggests as most appropriate, although both of those questions are also very important.⁴³ Affordances also, and necessarily, have collective (population-based) dimensions and implications. Meadows and prairies and desert scrublands are for eating in different ways and with different implications for the populations that want or need to eat them. Mobile phones are for connecting and being located in different ways and with different implications for the populations wanting connectivity.⁴⁴

Questions about affordances for fundamental rights also cannot simply be subsumed into capabilities discourse. To define a right in terms of capabilities is to specify a minimum threshold (of material wellbeing, literacy, and/or some other good) below which people cannot as a practical matter enjoy the civil and political rights they are presumed to possess. By contrast, to define a right in terms of materiality is to focus on the ways that its enjoyment is informed by the built environment's systemic tolerances and prohibitions. Access to information and communications capabilities may, of course, be distributed differentially — and so some kinds of claims about access to networked digital resources ought to figure prominently in capabilities-based formulations of fundamental rights⁴⁵ — but other types of questions about networked communication and information processing protocols are centrally concerned with how particular functionalities are achieved. So, for example, if access to credit or employment increasingly is mediated in ways that produce racial or socioeconomic segmentation, a liberty-based approach would highlight the discrimination and the resulting relative disadvantage to disfavored groups; a capabilities approach would highlight the disadvantaged groups' diminished access to essential resources and the resulting functional handicap; and an affordance-based approach would focus on the infrastructural configurations that enable market segmentation to proceed and to evade oversight.⁴⁶

43 See Donald A. Norman, *Affordance, Conventions, and Design*, 6 INTERACTIONS 38, 39 (1999).

44 Cf. Karen E.C. Levy, *The User as Network*, 20 FIRST MONDAY (2015) (“An individual use model . . . fails to illuminate the full range of social and political entanglements that underlie and mediate technological engagement.”), doi: <http://dx.doi.org/10.5210/fm.v20i11.6281>.

45 For an illustrative list of information-related capabilities, see Lea Bishop Shaver, *Defining and Measuring A2K: A Blueprint for an Index of Access to Knowledge*, 4 I/S: J.L. POL'Y INFO. SOC'Y 235 (2008). On a fundamental right to Internet access, see Christoph B. Graber, *Bottom-Up Constitutionalism: The Case of Net Neutrality*, 7 TRANSNAT'L LEGAL THEORY 524 (2016).

46 For examples of the sorts of analysis I have in mind, see Mary Madden et al., *Privacy, Poverty, and Big Data: A Matrix of Vulnerabilities for Poor Americans*,

One caveat: It is important to avoid trying to make a concept intended to explicate a particular dimension of the relationship between organisms and their environment do too much.⁴⁷ There are other constructs for getting at the ways that the mutually reinforcing interactions between design, business models, organizational imperatives, and communities of practice produce results that are both enabling and choice-foreclosing for users later on. They include, for example, the idea of a sociotechnical system — an arrangement reflecting the interaction of technical and social factors — and the related idea of an assemblage — “a mode of ordering heterogeneous entities so that they work together for a certain time.”⁴⁸ Such arrangements reflect and reproduce

95 WASH. U. L. REV. 53 (2017); Pauline T. Kim, *Data-Driven Discrimination at Work*, 58 WM. & MARY L. REV. 857 (2017).

- 47 For example, Ryan Calo uses the term affordance broadly to describe not only the enabling effects of artifacts but also those of cultural practices and norms. Ryan Calo, *Can Americans Resist Surveillance?*, 83 U. CHI. L. REV. 23 (2016); Ryan Calo, *Technology, Law, and Affordance: A Review of Smart Technologies and the End(s) of Law*, 4 CRITICAL ANALYSIS L. 72 (2017). See also Mireille Hildebrandt, *Law As an Affordance: The Devil Is in the Vanishing Point(s)*, 4 CRITICAL ANALYSIS L. 116, 121-22 (2017). So used, the idea of an affordance is nearly all-encompassing; one might say, for example, that a doctor’s office affords doctor-patient confidences or that a bazaar affords bargaining while a supermarket does not. That seems generally right but also not very precise. In the doctor’s office or the bazaar, culture is doing almost all of the work (in the doctor’s office, law is doing important work as well). A supermarket is more complicated, because it is in part a sociotechnical system consisting of scannable bar codes, cash wraps, and credit authorization uplinks that have been designed with a particular set of desired behaviors in mind. The behaviors in question do not include bargaining so the artifacts do not afford it (though they tend to have other affordances, such as the ability to receive coupon codes, that may enable similar discount-driven behaviors to occur). But the reasons for designing the artifacts that way are economic and cultural, and the layout of a supermarket also reflects other economic and cultural influences which in turn shape (constrain or direct) shopper behavior.
- 48 Martin Muller, *Assemblages and Actor-Networks: Rethinking Socio-Material Power, Politics, and Space*, 9 GEOGRAPHY COMPASS 27, 28 (2015); See generally Benjamin K. Sovacool & David J. Hess, *Ordering Theories: Typologies and Conceptual Frameworks for Sociotechnical Change*, 47 SOC. STUD. SCI. 703 (2017). For an illustration of the way the idea of an assemblage can illuminate issues at the intersection of law, business, and technology, see Tony Porter, *Tracing Associations in Global Finance*, 3 INT’L POL. SOC. 334 (2013). See also Madeleine Akrich, *The De-Description of Technical Objects*, in SHAPING TECHNOLOGY, BUILDING SOCIETY: STUDIES IN SOCIOTECHNICAL CHANGE 205 (Wiebe E. Bijker &

political and economic power, and they may crystallize into institutions that become more durable still. Intervening to restructure the affordances and disaffordances of particular artifacts may entail contending with the relative stickiness of sociotechnical systems, assemblages, and institutions, a point to which I will return in Part III below. The idea of an affordance simply surfaces a register — materiality — in which rights discourse must learn to operate.

As we saw in Part I, privacy rights have never fit particularly well within the implicit parameters of more conventional forms of discourse about fundamental rights precisely because privacy-related expectations and practices are relational, contextual, and spatial in character. An affordance-based approach to privacy promises greater taxonomic clarity. As one illustration of the difference that such a shift might make, consider the debate among European scholars over whether data protection is best understood as a separate fundamental right or as a way of implementing certain aspects of the fundamental right to privacy.⁴⁹ The answer is both — and neither. The “right to privacy” is a liberty-based formulation. The “right to data protection,” which is concerned with the conditions under which personal data may be collected, processed, used, and retained, is an entitlement better suited to articulation within an affordance-based discourse. This point also helps to explain why the seemingly inexorable drift toward notice and consent as a universal legitimating condition for satisfaction of data protection obligations is a strategy that cannot hope to succeed; consent is a liberty-based construct, but effective data protection is first and foremost a matter of design.

An affordance-based approach also promises to lend new rigor to the articulation and justification of entitlements to privacy. From the standpoint of affordances, privacy is most usefully described not as an abstract right or a static good to be traded off against other possible goods, but rather as an environmental condition and a related entitlement (or set of entitlements) relating to that condition. In my own earlier work, I have argued that the condition of privacy entails dynamic maintenance of breathing room for socially situated subjects to engage in the processes of boundary management through which they define and redefine themselves as subjects. Put differently,

John Law eds., 1992) (describing technologies as embedding behavioral scripts for heterogeneous networks of actors).

49 See, e.g., GLORIA GONZALEZ FUSTER, *THE EMERGENCE OF PERSONAL DATA PROTECTION AS A FUNDAMENTAL RIGHT OF THE EU* (2014); Lee A. Bygrave, *Data Protection Pursuant to the Right to Privacy in Human Rights Treaties*, 6 INT’L J.L. & INFO. TECH. 247 (1998); Orla Lynskey, *Deconstructing Data Protection: The “Added-Value” of a Right to Data Protection in the EU Legal Order*, 63 INT’L & COMP. L.Q. 569 (2014).

because subjectivity emerges within the interstices of social shaping, the nature and quality of the interstices are of paramount concern. Social shaping is a constant, but people also exercise agency with regard to self-articulation and may do so more fully to the extent that the interstices are larger and the architectural and informational linkages between them less complete. I have referred to the condition of incomplete linkage as *semantic discontinuity* and have argued that a sufficient quantum of semantic discontinuity is indispensable to human wellbeing.⁵⁰

One kind of strategy for developing an affordance-based approach to privacy involves taking the condition of semantic discontinuity as a lodestar and systematically elaborating the entitlements needed to preserve it. A signal virtue of this approach is its relatively stronger orientation toward practice and, in particular, toward design. In Part III we will see that detaching privacy from figures and reorienting it toward environmental conditions opens up spaces of operational possibility that a narrower focus on subjects forecloses (or has enabled us to ignore), even as the benefits of such an approach redound to those subjects. It directs our attention to the essential roles of gaps, barriers, breakdowns, and failures of translation in producing the conditions that render selves incomputable.⁵¹ Taking semantic discontinuity seriously also promises to offer some traction on the problem of actuarial decision-making that Part II described, because it opens the way for reasoning about how and why certain kinds of decision-making about individual claims and interests may be more dignifying *even when* — and indeed *because* — they treat those claims and interests in aggregate without attempting to subdivide them.⁵²

A complementary strategy for developing the elements of an affordance-based approach to fundamental rights involves retheorizing the necessary agency relationships to materiality in ways that depart from the traditional narrow emphasis on individual choice and consent. Mireille Hildebrandt's compelling new formulation — “the right to co-determine how we will be read” — is a promising starting point for that project because it is framed in terms of accountability for the processes through which we are rendered legible by and to our intelligent artifacts.⁵³ Although we cannot entirely escape the constitutive force-fields generated by our technologies — and hence it would

50 COHEN, CONFIGURING THE NETWORKED SELF, *supra* note 13, at 239-41.

51 *See generally* Hildebrandt, *supra* note 38.

52 For the beginning of an argument about both the theoretical possibility and the value of dignifying aggregation, see COHEN, CONFIGURING THE NETWORKED SELF, *supra* note 13, at 250-52.

53 MIREILLE HILDEBRANDT, SMART TECHNOLOGIES AND THE END(S) OF LAW: NOVEL ENTANGLEMENTS OF LAW AND TECHNOLOGY 102-03 (2015).

be intellectually dishonest to speak of a right to “determine” our own legibility to other human and nonhuman actors — we can and should expect to have a say. At the same time, the idea of co-determination rights leaves open the precise nature of the balance to be struck between agential co-determination and system-wide protections that sound in semantic discontinuity (or in some other general value).

The terminology of co-determination rights, however, threatens to direct our attention back toward individualized consent and divert it from broader questions about collective self-determination of conditions, and so I prefer to think of the idea of co-determination rights in terms of *operational accountability*. As Lisa Austin has observed, certain common relationships in contemporary commercial and civic life simply cannot be theorized in terms of individual choice and consent in any meaningful way; those relationships are about power, and privacy theory should acknowledge that fact and move on.⁵⁴ Accountability to the subjects of privacy also is about much more than just the opportunity to make choices or to indicate “consent.” Accountability, like affordance, is a relational and collective construct; it entails taking responsibility for outcomes in regard to communities of stakeholders, and relations of accountability are (or should be) relations of respect. A regime of operational accountability grounded in appropriate respect for users and civil society more generally must afford a say in the conditions of our own legibility, and it sometimes will be more meaningful and more effective to afford that say collectively. A regime of operational accountability also must provide users and communities of users adequate and meaningful levels of operational transparency about the sociotechnical systems within which they are enmeshed.⁵⁵

Notably, neither a semantic discontinuity principle nor an operational accountability principle is likely to produce results that correspond well to the boundaries of privacy as traditionally understood within liberty-based rights discourse. Both semantic discontinuity and operational accountability encompass privacy-related considerations and other considerations, such as freedoms of speech, association, and thought. But that objection is not really an objection at all. As capabilities discourse illustrates, there is no reason to imagine that an affordance-based approach to privacy and other fundamental rights will follow a pattern of strict one-to-one correspondence, nor is there any need to impose such a requirement. Just as, for example, the capabilities-based rights

54 Lisa M. Austin, *Enough About Me: Why Privacy Is About Power, Not Consent (or Harm)*, in *A WORLD WITHOUT PRIVACY? WHAT CAN/SHOULD LAW DO?* 131 (Austin Sarat ed., 2014).

55 On the importance of operational transparency see COHEN, *CONFIGURING THE NETWORKED SELF*, *supra* note 13, at 234-39.

to health or literacy have implications for the enjoyment of multiple liberty-based rights, so affordance-based rights conceptualized with privacy in mind will have implications for other liberty-based rights as well. The relationship between privacy and data protection is a case in point; the two rights overlap in coverage, but the overlap is incomplete. Effective data protection also serves a number of non-privacy-related interests, such as interests in freedom of thought, belief, and association, and effective protection of breathing room for self-development requires more than just data protection. As capabilities discourse also illustrates, there also is no reason to suppose that affordance-based rights need be framed as attaching exclusively to individuals rather than to communities. To the contrary, conceptualizing privacy in a way that foregrounds conditions rather than subjects underscores the extent to which privacy rights are inherently communal, as several burgeoning strands of the privacy literature have suggested.⁵⁶ Securing breathing room for self-articulation requires universally applicable material and operational guarantees.

Within networked digital environments, new affordance-based correlates to liberty-based rights of privacy, autonomy, and self-determination — formulated in terms such as the right to a sufficient baseline level of semantic discontinuity and operational accountability — seem likely to emerge as core protections for fundamental rights in the digital era. Both formulations offer more than just new kinds of abstract rhetoric about the importance of human freedom. They are ways of directing attention to required sets of sociotechnical conditions and demanding their effective realization. They envision reconfiguring rights discourse all the way down, so that it speaks with effective force to new kinds of material and operational considerations. They therefore direct our attention inexorably to the other half of the theory-practice relation — to privacy’s institutional grammar.

III. OPERATIONALIZING AN INSIDE-OUT APPROACH TO PRIVACY

An affordance-based approach to conceptualizing fundamental rights also demands attention to the kinds of infrastructural and operational details with which theories and institutions designed around liberty-based and capabilities-based approaches generally have not engaged. The theoretical moves described in the previous section suggest complementary sets of principles directed toward privacy institutions and privacy-related design practices.

⁵⁶ See generally *GROUP PRIVACY: NEW CHALLENGES OF DATA TECHNOLOGIES* (Linnet Taylor, Luciano Floridi & Bart van der Sloot eds., 2017).

We saw in Part I.B that conventional data protection approaches framed in terms of consent by individual subjects and command-and-control oversight by designated data controllers have failed to provide effective privacy protection. The discussion in Part II helps to explain why: conventional data protection instruments emphasize subject-centered control of discrete information flows rather than protection of privacy-preserving boundaries and conditions. We also saw in Part I.B that the European conception of purpose binding, which comes closest to decentering subjects and foregrounding conditions, confronts significant operational challenges that limit its ability to accomplish either perspectival shift effectively.

Protecting privacy effectively requires a willingness to depart more definitively from subject-centered frameworks in favor of condition-centered frameworks — and to refrain from labeling the latter as offensive because they are “paternalistic.” Design has consequences that limit the horizons of possibility for choice. Not all such consequences can or should be retheorized within neoliberal frameworks that cast them as infringements on liberty. Design for seamless data harvesting subjects individuals to the rhythms of organizations seeking to use their data for extractive purposes. Design for semantic discontinuity and operational accountability would disrupt those organizational rhythms. One design ethos is not self-evidently more paternalistic than the other.

The ideas of semantic discontinuity and operational accountability supply foundational condition-centered design principles for the production of privacy-friendly environments. A semantic discontinuity principle would require the deliberate pursuit of durable strategies for foreclosing and disrupting information flow in ways that frustrate seamless legibility and manipulation. One example of what that project might look like, in microcosm, is the work by Paul Ohm and Jonathan Frankle on “desirable inefficiency” in the design of digital systems and artifacts.⁵⁷ Ohm and Frankle identify design practices that conventional, efficiency-driven thinking would disfavor and link those practices to specific regulatory functions and values. Another set of examples appears in the growing body of research on “seamful design,” which focuses on deliberately interrupting information flows in ways that bring mediation and artificiality to users’ attention and furnish meaningful points of intervention.⁵⁸ An operational accountability principle would direct regulatory attention to

57 Paul Ohm & Jonathan Frankle, *Desirable Inefficiency*, 70 FLA. L. REV. (forthcoming 2018).

58 See Matthew Chalmers & Ian MacColl, *Seamful and Seamless Design in Ubiquitous Computing*, 8 WORKSHOP AT THE CROSSROADS: THE INTERACTION OF HCI AND SYSTEMS ISSUES IN UBIComp (2003), <http://www.techwondo.com/external/>

matters such as the nature and quality of explanations about data collection and processing and the possibilities for giving both individual users and society more generally a say in the sorts of legibility we are willing to underwrite.

To be effective, however, constraints originating in principles of semantic discontinuity and operational accountability must operate on the same scale at which design decisions affecting privacy are made. As Seda Gürses and Joris Van Hoboken explain, the scholarly literatures on privacy and privacy governance have focused on the demand side of the equation and therefore have not grappled with the ways that larger changes in practices of systems design appear to stack the deck against data protection. Contemporary design practices emphasize modularity, continual rewriting and run-time upgrades, and seamless flow across platforms and applications, and all of those characteristics make effective data protection much more difficult.⁵⁹ This observation returns us by a different route to the point about sociotechnical systems and assemblages raised in Part II above. Although Gürses and Van Hoboken do not fully engage the point, what they characterize as the “agile turn” in software development also draws momentum from the political economy of data harvesting. Networked digital environments have important affordances for continual rewriting and seamless flow, but they also have been configured over time to generate as much appropriable data as possible in the interest of new models of profitability.⁶⁰ This means that efforts to instantiate semantic discontinuity throughout the networked digital environment generally must confront not only the norms of continuous rewriting but also more intractable obstacles that are both organizational and conceptual.

With regard to organizational structure, important work by Kenneth Bamberger and Deirdre Mulligan demonstrates that privacy oversight has become well integrated into the corporate risk management landscape.⁶¹

pdf/reports/2003-chalmers.pdf; Janet Vertesi, *Seamful Spaces: Heterogenous Infrastructures in Interaction*, 39 *SCI. TECH. & HUM. VALUES* 64 (2014).

59 Seda Gürses & Joris van Hoboken, *Privacy after the Agile Turn*, in *CAMBRIDGE HANDBOOK OF CONSUMER PRIVACY* 579 (Evan Selinger, Jules Polonetsky & Omer Tene eds., 2018).

60 On the sociotechnical, legal, and economic dimensions of that process of (re)-configuration, see JOSEPH TUROW, *THE AISLES HAVE EYES: HOW RETAILERS TRACK YOUR SHOPPING, STRIP YOUR PRIVACY, AND DEFINE YOUR POWER* (2017); Julie E. Cohen, *The Biopolitical Public Domain: The Legal Construction of the Surveillance Economy*, 31 *PHIL. & TECH.* 213 (2018); Shoshana Zuboff, *Big Other: Surveillance Capitalism and the Prospects of an Information Civilization*, 30 *J. INFO. TECH.* 75 (2015).

61 KENNETH A. BAMBERGER & DEIRDRE K. MULLIGAN, *PRIVACY ON THE GROUND: DRIVING CORPORATE BEHAVIOR IN THE UNITED STATES AND EUROPE* (2015).

Academic researchers in computer science and information studies have begun to use Helen Nissenbaum's theory of privacy as contextual integrity as a lodestar for more privacy-aware development practices that would guard against or minimize unforeseen cross-contextual information flows.⁶² Such efforts also draw on the academic literature on values in design, which creates blueprints for integrating normative considerations into the design process.⁶³ But translating that embryonic sensibility into private-sector workplaces is a project of a very different order. Corporate internalization of the need for privacy oversight does not necessarily result in effective internalization of privacy-related values and imperatives by the technologists who design networked digital products and services. Preliminary research into the attitudes of employees directly responsible for the design and ongoing maintenance of networked digital products and services suggests different and more adversarial understandings of privacy compliance requirements.⁶⁴

One way to pursue greater internalization of privacy-friendly design values involves rethinking the ways that future technologists are educated and trained. Some have argued that this problem can be approached by adopting codes of ethics for technologists and technology companies and incorporating those codes into professional training and development at every level. Improving ethics-related education for future technologists and on-the-job training for

62 HELEN NISSENBAUM, *PRIVACY IN CONTEXT: TECHNOLOGY, POLICY, AND THE INTEGRITY OF SOCIAL LIFE* (2010); *see, e.g.*, Gordon Hull, Heather Richter Lipford & Celine Latulipe, *Contextual Gaps: Privacy Issues on Facebook*, 13 *ETHICS & INFO. TECH.* 289 (2011); Louise Barkhuus, *The Mismeasurement of Privacy: Using Contextual Integrity to Reconsider Privacy in HCI*, in *PROCEEDINGS OF THE SIGCHI CONFERENCE ON HUMAN FACTORS IN COMPUTING SYSTEMS* 367 (2012); Pan Shi, Heng Xu & Yunan Chen, *Using Contextual Integrity to Examine Interpersonal Information Boundary on Social Network Sites*, in *PROCEEDINGS OF THE SIGCHI CONFERENCE ON HUMAN FACTORS IN COMPUTING SYSTEMS* 35 (2013); Frances S. Grodzinsky & Herman T. Tavani, *Privacy in "the Cloud": Applying Nissenbaum's Theory of Contextual Integrity*, 41 *SIGCAS COMPUTERS & SOC'Y* 38 (2011); Primal Wijesekera et al., *Android Permissions Remystified: A Field Study on Contextual Integrity*, in *PROCEEDINGS OF THE 24TH USENIX SECURITY SYMPOSIUM* 499 (2015).

63 *See* HUMAN VALUES AND THE DESIGN OF COMPUTER TECHNOLOGY (Batya Friedman ed., 1997).

64 Ari Ezra Waldman, *Designing Without Privacy*, 55 *Hous. L. Rev.* 659 (2018); *cf.* JESSICA SILBEY, *THE EUREKA MYTH: CREATORS, INNOVATORS, AND EVERYDAY INTELLECTUAL PROPERTY* 184-220 (2015) (exploring the potential for breakdown of relationships between in-house IP lawyers and their clients).

practicing technologists surely are urgent projects.⁶⁵ But ethical guidelines often replicate the principal failings of liberty-based formulations of privacy rights — they tend to be couched in vague, general terms that are difficult to translate into more concrete design principles. Giving ethics review processes teeth — for example, by requiring approval of privacy-related design aspects by in-house IRBs — risks returning by a different route to the problem of intra-organizational adversarialism.⁶⁶ There is also considerable risk that such projects may be understood as useful primarily for managing public perception rather than for implementing meaningful change. So, for example, the American data-analytics company Palantir has constituted a blue-ribbon advisory board composed of prominent privacy scholars, but it's unclear whether that move has resulted in significant alterations to its core products and services, which are designed to give federal, state, and local law enforcement the ability to conduct pervasive, cooperative, long-term dataveillance of populations.⁶⁷

Changing the culture of an occupation and an industry, particularly in the face of what has become the industry's *de facto* profit model, is a longer-term project that requires a more fundamental rethinking of what constitutes "good" design. Those who doubt that design is always-already infused with values need look no further than the cybersecurity debates in which many privacy professionals also participate. The pathbreaking "end to end" design

65 For a thoughtful discussion, see Susan Landau, *Educating Engineers: Teaching Privacy in a World of Open Doors*, 12 IEEE SECURITY & PRIVACY 66 (2014). Arguably, a much broader educational program is needed. See, e.g., John Naughton, *How a Half-Educated Tech Elite Delivered Us Into Chaos*, THE GUARDIAN (Nov. 19, 2017), <https://www.theguardian.com/commentisfree/2017/nov/19/how-tech-leaders-delivered-us-into-evil-john-naughton> ("[T]he new masters of our universe are people who ... have had no exposure to the humanities or the social sciences, the academic disciplines that aim to provide some understanding of how society works, of history and of the roles that beliefs, philosophies, laws, norms, religion and customs play in the evolution of human culture.").

66 See Jules Polonetsky, Omer Tene & Joseph Jerome, *Beyond the Common Rule: Ethical Structures for Data Research in Non-Academic Settings*, 13 J. TELECOMM. & HIGH TECH. L. 333 (2015).

67 Palantir, *Announcing the Palantir Council on Privacy and Civil Liberties*, PALANTIR, <http://www.palantir.com/2012/11/announcing-the-palantir-council-on-privacy-and-civil-liberties/> (Nov. 2012); Mark Harris, *How Peter Thiel's Secretive Data Company Pushed Into Policing*, WIRED (Aug. 9, 2017), <https://www.wired.com/story/how-peter-thiels-secretive-data-company-pushed-into-policing/>; Ali Winston, *Palantir Has Secretly Been Using New Orleans to Test Its Predictive Policing Technology*, THE VERGE (Feb. 27, 2018), <https://www.theverge.com/2018/2/27/17054740/palantir-predictive-policing-tool-new-orleans-nopd>.

of technical protocols for the Internet reflected solid technical judgment about robustness to certain kinds of disruptions and also encoded the generally libertarian commitments of the original Internet pioneers. As a result, although the Internet overall is extraordinarily resistant to disruptions of service, it has proved extraordinarily hospitable to other kinds of threats that exploit networked interconnection. The vulnerabilities of the Internet's present are encoded in the value choices of its past.⁶⁸

Operationalizing semantic discontinuity at scale requires multiple and overlapping sets of broadly distributed strategies for durably interrupting (rather than temporarily disrupting) networked information flows. The growing body of research on differential privacy aims to produce such interruptions computationally, introducing noise into data sets in ways that foreclose reidentification of individuals while still enabling the data sets to be used as research tools.⁶⁹ It is important to recognize, however, that differential privacy is not a magic bullet. To begin with, practitioners of differential privacy sometimes seem loath to acknowledge that different kinds of data aggregation have different values and different politics attached to them. Techniques for differential privacy are especially valuable where the underlying research has social value rather than merely private value. There are differences between using population data to study epidemiology and using it to refine capabilities for psychometric targeting, and there are differences between using epidemiological research to underwrite techniques for pharmaceutical price discrimination and using it to identify risk factors for particular diseases or to counteract systemically embedded environmental racism. Additionally, the idea of differential privacy as a mode of epistemological target-hardening breaks down at its endpoints — the places where third-party apps plug into platforms and where users connect with devices via interfaces.

For both reasons, even universal mandatory adoption of differential privacy techniques would be an insufficient strategy for operationalizing semantic discontinuity. In particular, there remains a need for strategies aimed at undermining seamless functionality by limiting flows across devices and applications in hard-coded ways, setting nonnegotiable (and perhaps randomly

68 For a prescient early treatment of this problem, see JONATHAN ZITTRAIN, *THE FUTURE OF THE INTERNET—AND HOW TO STOP IT* 36-57 (2008).

69 See Kobbi Nissim et al., *Differential Privacy: A Primer for a Non-technical Audience (Preliminary Version)* (May 7, 2017) (unpublished manuscript), https://www.ftc.gov/system/files/documents/public_comments/2017/11/00023-141742.pdf; Cynthia Dwork & Aaron Roth, *The Algorithmic Foundations of Differential Privacy*, 9 *FOUND. & TRENDS THEORETICAL COMPUTER SCI.* 211 (2014); Cynthia Dwork, *Differential Privacy: A Survey of Results*, in *THEORY AND APPLICATIONS OF MODELS OF COMPUTATION 1* (Manindra Agrawal et al. eds., 2008).

varying) limits on the extent of the networked convenience it is possible to have. Interrupting practices organized around seamless convenience also might advance the goal of boundary-setting described in Part II. Functional interruptions linked to well-designed cognitive “breaks” would foreground the automated processing running continuously behind our interfaces, reminding us of the irreducible nonhumanness of our automated tools.

For some, the idea of saying no to data flow will be powerfully counterintuitive. Here practice should return to theory, and to the lessons about technologies, social values, and path dependence illustrated in the cybersecurity example above. Different levels and types of flow create different benefits and produce different vulnerabilities. Borrowing from the environmental framing that has become familiar in debates about privacy theory and policy, such restrictions might perhaps be understood as analogous to rules requiring “low-flow” showerheads and toilets for environmental reasons.⁷⁰ No-flow mandates in turn evoke measures ranging from firewalls for critical systems to border restrictions on agricultural imports designed to protect local biomes against colonization by invasive species. The parameters of *low-flow and no-flow directives for personal information harvesting and use*, however, would need to be specified in terms designed to secure compliance by automated machine-learning systems. As Part I.B explained, that project requires active oversight; it entails implementation strategies designed to stand up to machine workarounds and adequate attention to all of the stages through which machine-learning processes are designed and refined. Additionally, it requires complementary strategies for incentivizing compliance with low-flow and no-flow directives. Over the decades, it has become easier to make water more expensive. Data harvesting is cheap and difficult to police, but perhaps it need not remain that way.

Other kinds of disruptions might challenge totalizing knowledge frameworks. We saw in Part I.B that conventional approaches to data protection do little to dislodge certain practices of decision-making about access to important resources that privilege informational granularity as a species of epistemological due process. As just discussed, research on differential privacy has made important progress toward operationalizing ideals of obfuscation in ways that limit the discoverability of individual information while preserving the possibility of knowledge discovery within databases. In theory, techniques for obfuscation also could be used to make certain kinds of allocative decision-making fuzzier.

70 Cf. A. Michael Froomkin, *Regulating Mass Surveillance as Privacy Pollution: Learning from Environmental Impact Statements*, 2015 U. ILL. L. REV. 1713 (2015); Dennis D. Hirsch, *The Glass House Effect: Big Data, the New Oil, and the Power of Analogy*, 66 ME. L. REV. 373 (2014).

But differential privacy does not address normative questions, and there are many situations in which policymakers and for-profit actors would contend that granular, fully-identified information should be preserved and used for reasons of both efficiency and fairness.

Here again, practice can and should return to theory; for initiatives directed toward knowledge obfuscation within decision-making frameworks to be taken seriously, it is necessary to explain why and when incomplete, fuzzy, or even inaccurate personalization can be dignifying. When a highly granular decision about an individualized program of medical treatment must be made, the case for personalization is strong. When the question concerns social or commercial judgments about desert — as is the case, for example, with determinations about eligibility for credit or disability benefits — the calculus might be different. Foundational commitments to the rule of law typically demand generally applicable rules applied in reasoned ways and require that like cases be treated alike. The notion of like cases presupposes both the ability and the need to make certain kinds of distinctions, but at the same time data-driven predictive profiling now threatens to strain both the very idea of likeness and the commitments to general applicability and reasoned application to the breaking point. Efforts to articulate the normative value of lumpiness in practices surrounding resource allocation might draw upon basic norms of legality, blending old theories with new techniques to construct a framework for allocative fairness based on good-enough decision-making.

Many of these strategies for operationalizing semantic discontinuity also feed into the pursuit of operational accountability for organizational choices that affect the legibility of users both individually and collectively. Achieving operational accountability in a way that moves beyond subject-centered conceptions of choice and consent requires that new techniques for ensuring the accountability of algorithmic and machine-learning processes be paired with new regulatory competencies capable of harnessing those techniques on behalf of the publics affected by those processes. Growing literatures address both sets of questions, defining a range of possible pathways for institutional change.⁷¹ In connection with that project, it will also be important to consider

71 On algorithmic accountability, see Joshua A. Kroll et. al., *Accountable Algorithms*, 165 U. PA. L. REV. 633 (2017). On entry points for accountability in machine learning, see Lehr & Ohm, *supra* note 24. On new regulatory competencies for the information era, see Kenneth A. Bamberger, *Technologies of Compliance: Risk and Regulation in a Digital Age*, 88 TEX. L. REV. 669 (2010); Julie E. Cohen, *The Regulatory State in the Information Age*, 17 THEORETICAL INQUIRIES L. 369 (2016); Cary Coglianese & David Lehr, *Regulation by Robot*, 165 U. PA. L. REV. 1147 (2017); Rory Van Loo, *Rise of the Digital Regulator*, 66 DUKE L.J. 1267 (2017).

how to harness new algorithmic and machine-learning technologies rather than merely seeking to constrain them. An example of the former sort of approach is Hildebrandt's proposal to deploy machine learning agonistically, producing competing sets of patterns and interpretations that might open a broader and more open-ended space for policymaking.⁷² Another is the work by the technology policy organization Robinson + Yu (now Upturn) exploring alternative data-driven techniques for endorsing the reliability of lower-income consumers who have lacked access to conventional sources of credit.⁷³ Both proposals express respect for and responsibility to the subjects of privacy, and those principles should be touchstones of operational accountability more generally.

CONCLUSION

In the networked information era, preserving effective privacy protection for the subjects of privacy entails decentering them both within theoretical frameworks and in the design of privacy institutions. Developing institutions and practices for operationalizing an affordance-based approach to privacy requires a design ethos informed by careful attention to the relationship(s) between materiality and practices of self-articulation, and especially to the importance of boundaries, gaps, and discontinuities for those practices. It also requires a robust conception of operational accountability that moves beyond individualized choice and consent to emphasize responsibility, respect, and new modalities for effective regulatory oversight of algorithmic and data-driven processes.

72 Hildebrandt, *supra* note 38, at 105.

73 ROBINSON + YU, KNOWING THE SCORE: NEW DATA, UNDERWRITING, AND MARKETING IN THE CONSUMER CREDIT MARKETPLACE (2014), https://www.teamupturn.org/static/files/Knowing_the_Score_Oct_2014_v1_1.pdf.

