The Value—and Limits—of Distributive Justice in Information Privacy

Jeffrey Alan Johnson
Utah Valley University

Paper Presented at the
Digital Sociology Mini-Conference,
Eastern Sociological Society 2016 Annual Meeting

and the

Western Political Science Association 2016 Annual Meeting

Abstract: Information privacy is rife with unanswered questions of distributive justice. Privacy at least contributes to achieving just distributions of social goods. Information flow models of privacy assume answers to questions about justice in acquisition and transfer that may be indefensible on their own or incompatible with each other. Rights approaches often assume rather than articulate a justification of privacy itself. And there are considerable implicit differences between the two in what is to be distributed. It should not be seen as practical to address questions of information privacy without considering these questions. But Young's critique of the distributive paradigm reveals deeper problems with understanding the question of justice in information privacy. Information privacy is as much a matter of social structure as it is of distributing material or moral goods, and the focus on distribution obscures the ways in which information privacy violations challenge the ability to participate in determining one's actions. These critiques suggest that a more productive line of inquiry would be to pursue information justice as a matter of primarily structural rather than distributive justice.

I appreciate the helpful suggestions of Jade Davis, Melissa Frost, Michael Minch, Annette Moulder, Christopher Robinson, Jathan Sadowski, Steve Vanderheiden, and Sharon Yamen in the development of this paper.

Jeffrey Alan Johnson, Interim Director
Institutional Effectiveness and Planning
Utah Valley University

jeffrey.johnson@uvu.edu
@the_other_jeff

800 West University Parkway
Mail Stop 272
Orem, Utah 84107

The Value—and Limits—of Distributive Justice in Information Privacy

1 INTRODUCTION

One of the few strains of coherence in the study of information privacy[1] is acceptance of its incoherence. Woodrow Hartzog and Evan Selinger argue that in both academic theory and U.S. law, "despite the widespread concern and extensive academic treatment of surveillance issues, the language and framing used in surveillance debate is diverse, inconsistent, and over-generalized (2015, 1344);" H. Jeff Smith, Tamara Dinev, and Heng Xu suggest that in information systems research "the findings and the theories that emerged have often relied on overlapping constructs nestled within loosely bounded nomological networks. This has resulted in a suboptimal cumulative contribution to knowledge (2011, 990)." Daniel Solove's pronouncement has an air of definitiveness[2] about it: "Privacy, however, is a concept in disarray. Nobody can articulate what it means (2008, 1)."

This problem echoes other fields of applied social and political philosophy, notably that of environmental justice. David Schlosberg notes that the latter field is marred by a weak understanding of justice itself, making "most theories of environmental justice . . . incomplete theoretically (2004, 517)." Relying on the work of Iris Marion Young and of Nancy Fraser, he

[1] For the purpose of this paper, I will adhere, with some modification as this paper develops, to the distinctions among general privacy, physical privacy, and information privacy offered by Smith, Dinev, and Xu (2011, 990–991). *Physical privacy* refers to the privacy of "an individual and/or the individual's surroundings and private space" where *information privacy* refers to the privacy of information about one or more individuals, whether as individuals or as a group. General privacy refers to both. Unless otherwise specified, I will use *privacy* as a shorthand for information privacy, and specify when I refer to physical or general privacy. However, I depart from the authors above in confining privacy in any form to issues of access, for reasons that will be developed below.

[2] A false air, or course, coming at the beginning of a book Solove begins by saying "I am ready to set forth my theory of privacy (2008, ix)."

argues for a more expansive understanding of justice composed not only of the common distributive framework but also of a need to secure recognition of and participation by all groups in society. Using this framework, he is able to develop a framework of "critical pluralism" for environmental justice that makes sense not only philosophically but also in light of claims by social movements dedicated to securing environmental justice.

Schlosberg's success in framing a more complex vision of justice in environmental issues suggests the value of a similar framing for privacy, both generally and in information privacy specifically. Drawing on Schlosberg's conclusion that "justice itself is a concept with multiple, integrated meanings (2004, 536)," it may be possible to engage the challenges of information privacy effectively from a justice-centered perspective. That is to say, we can build a more effective theory of privacy by addressing privacy in relation to, as Serge-Christopher Kolm defines justice, "the central ethical judgment regarding the effects of society on the situation of social entities (1993, 438)."[3] Justice is the primary standard by which social and political structures, actions, and practices are evaluated. Echoing Aristotle, John Rawls calls justice "the first virtue of social institutions, as truth is of systems of thought (2005, 3);" Young considers justice "the primary subject of political philosophy (1990, 3)." Information privacy can be understood as a specific kind of political situation or condition of a social entity, that regarding information about the entity, which is affected by the situations and actions of other social entities. A framework for information privacy can thus be developed that evaluates situations

[3] Kolm adds the caveat, "with respect to each entity's valuation of its own situation for its own purposes (1993, 3)." This is, to be sure, a common enough feature of most modern theories of justice, especially those rooted in classical liberal political thought. But it cannot be a definition of the problem of justice in itself, as many theories of justice explicitly reject the idea that the entity's own valuation is central. Plato's theory of justice is the paradigmatic case: in arguing that justice is "the minding of one's own business and not being a busybody (Plato 1991, 111 [433a])," i.e., of fulfilling one's naturally ordained role in society, Plato explicitly rejects as unjust the pursuit of one's own purposes and holds instead that justice is to be judged with respect to nature's valuation of an entity's situation rather than the entity's valuation.

according to judgments about the rightness of that situation, and that (hopefully) promotes information practices that tend toward right situations.

To do so, of course, requires one or more underlying theories of justice itself. Schlosberg's success in environmental justice comes by recognizing that justice exists in multiple forms, and that our knowledge of justice flows from both abstract theory and social practice. This is true of privacy as well. The terrain of justice in information privacy, however, differs slightly from that which Schlosberg identifies—a condition that is entirely consistent with his claim in favor of localized and contextualized theories of justice. I suggest three main ideas of justice in information privacy: an instrumental view in which information privacy is valuable for its contribution to justice in other spheres, the common distributive paradigm that forms the basis for much contemporary work on information privacy, and an enhanced structural form composed of fully integrated recognition and participatory dimensions. In this paper, I primarily examine the first approaches. These approaches prove to have significant value in understanding privacy, but also have significant limitations, especially in light of criticisms offered by the third approach, that suggest that information privacy is most likely to be engaged effectively by a structural approach to justice.

2 INSTRUMENTAL JUSTICE IN INFORMATION PRIVACY

Information privacy may be related to justice only tangentially. One might ask what harm there is simply in having information; from a philosophical perspective, this is tantamount to asking if information privacy is a requirement of justice in itself or simply useful in the pursuit of justice in other forms. The latter is certainly a common attitude within the information technology industry, where many leaders have happily discarded information privacy as a concern. Facebook founder and chief executive Mark Zuckerberg has observed that information privacy is no longer a "social norm" as expansion of the Internet has made people more comfortable with sharing

information and more open about themselves: “That social norm is just something that has evolved over time,” Zuckerberg believes (B. Johnson 2010). At the 2015 World Economic Forum in Davos, Switzerland, computer science professor Margo Seltzer stated, “How we conventionally think of privacy is dead” because “Privacy as we knew it in the past is no longer feasible”; tech entrepreneur Anthony Goldbloom observed “I trade my privacy for the convenience. Privacy is not something that worries me (Carter 2015).” All of these suggest that privacy is not morally valuable in itself: it can be traded for convenience, abandoned because it is impractical, evolve out of existence.

That is not to say, however, that privacy is irrelevant to justice. One could argue that, even if information privacy can be abandoned without inherent injustice, it should be maintained because it is a valuable tool for protecting justice. The harm caused by violations of one’s privacy, from such a perspective, is not in the information transfer itself but in the (actual or potential) effects of the information transfer on the achievement of justice; holding private information about another is not unjust, but what one does with the information may be exceptionally so. Solove, for instance, finds arguments about the intrinsic value of privacy problematic, but concludes, “The value of ameliorating privacy problems lies in the activities that privacy protections enable” (Solove 2008, 85).[4] Certain kinds of injustice, one might with good reason argue, require information about the objects of action. Privacy is to be maintained as a way of preventing those who would use information unjustly from having one of the necessary conditions for the implementation of injustice. Privacy is thus, from this perspective, of solely—but by no means inconsiderable—instrumental value.

[4] To be clear, Solove does not support a strictly instrumental theory of privacy, arguing rather than privacy is a *sine qua non* for certain kinds of activities that are of the essence of justice. He does nonetheless endorse the idea that privacy is of value even if it is not a question of justice itself.

This is the approach to information privacy at work in employment discrimination law in the United States. It is generally considered a best practice for potential employers not to solicit information from job candidates regarding membership in "protected classes," those groups protected by the body of employment discrimination law. Those include "race, color, religion, sex, or national origin" under Title VII of the Civil Rights Act of 1964, persons over 40 years old under the Age Discrimination in Employment Act of 1967, those with disabilities under the Americans with Disabilities Act of 1990 (ADA), and genetic information under the Genetic Information Nondiscrimination Act of 2008 (GINA). The practice of restricting the flow of information about membership in protected classes can be seen as a fairly typical information privacy protection.

Federal employment discrimination law in the United States varies widely in its restrictions on collecting information. The Civil Rights Act does not, in fact, prohibit transfer of information about protected classes at all. The essential restrictions on employers under the Civil Rights Act are:

> (a)(1) to fail or refuse to hire or to discharge any individual, or otherwise to discriminate against any individual with respect to his compensation, terms, conditions, or privileges of employment, because of such individual's race, color, religion, sex, or national origin; or
>
> (2) to limit, segregate, or classify his employees or applicants for employment in any way which would deprive or tend to deprive any individual of employment opportunities or otherwise adversely affect his status as an employee, because of such individual's race, color, religion, sex, or national origin. (42 U.S.C. § 2000e-2)

This restriction is repeated nearly verbatim with exception of the description of the protected classes in question for the Age Discrimination in Employment Act (29 U.S.C. § 623), the ADA (42 U.S.C. § 12112), and GINA (42 U.S.C. § 2000ff-1). Indeed, many employers routinely

collect information about membership in protected classes from job applicants but isolate that information from hiring decisions. The advice against soliciting such information with other hiring information, and thus to further the protection of job applicants' privacy, should be seen as a protection of justice for both the applicant (from unlawful discrimination) and the employer (from unfounded claims of unlawful discrimination): employers cannot discriminate on the basis of a condition about which they have no knowledge.

The ADA and GINA do include provisions barring the collection of information itself. The ADA prohibits employers from:

> (6) using qualification standards, employment tests or other selection criteria that screen out or tend to screen out an individual with a disability or a class of individuals with disabilities . . . ; and
>
> (7) failing to select and administer tests concerning employment in the most effective manner to ensure that, when such test is administered to a job applicant or employee who has a disability that impairs sensory, manual, or speaking skills, such test results accurately reflect the skills, aptitude, or whatever other factor of such applicant or employee that such test purports to measure, rather than reflecting the impaired sensory, manual, or speaking skills of such employee or applicant . . . . (42 U.S.C. 12112(a)(6)-(7))

Moreover, subsection (d) extends the anti-discrimination provisions of subsection (a), which beyond the paragraphs quoted above generally mirror those of the Civil Rights Act, to include collecting information through medical examinations. Employers are specifically barred from "conduct[ing] a medical examination or mak[ing] inquiries of a job applicant as to whether such applicant is an individual with a disability or as to the nature or severity of such disability" (subparagraph (d)(2)(A)); a similar provision (subparagraph (d)(4)(A)) protects employees.

All of these provisions, however, include exemptions allowing the collection of information about the ability of a person to perform job related functions. In the case of

employment discrimination, the use of tests is permitted where "the standard, test or other selection criteria, as used by the covered entity, is shown to be job-related for the position in question and is consistent with business necessity" (paragraph (a)(6)) and where impaired skills "are the factors that the test purports to measure" (paragraph (a)(7)). Similarly, within the restrictions of section (d) an employer "may make preemployment inquiries into the ability of an applicant to perform job-related functions"; may require exams before an employee begins work and make offers condition on the results of those exams; may require medical examinations for current employees where "such examination or inquiry is shown to be job-related and consistent with business necessity"; and may conduct voluntary exams as part of benefits programs or that evaluate employees' abilities to perform their jobs. All of these exceptions are subject to the requirement that such exams are administered regardless of disability, maintained as a separate and confidential medical record, and used only as permitted elsewhere under the ADA.

GINA similarly makes it unlawful to "request, require, or purchase genetic information with respect to an employee or a family member of the employee" (42 U.S.C. § 2000ff-1(b)). But it, too, contains a lengthy list of exceptions that allow employers to collect such information for a number of non-discriminatory circumstances: "inadvertent" requests for family history, the operation of wellness programs (with protections for consent, confidentiality, and anonymity), certifying family leave requests, publicly available information that includes family history, monitoring workplace safety (again with protections for purposes, consent, confidentiality, and anonymity), and the unique needs of forensic DNA analysis labs. The information collected under these exceptions is considered a confidential medial record, protected from disclosure and to be kept separate from employment records (42 U.S. Code § 2000ff–5).

The most stringent protections of information privacy in the employment discrimination regime are thus shown to be the most clearly instrumental: information privacy is valuable not in itself but as a protection from unjust discrimination. Information privacy is a question related to justice in employment discrimination; the latter is a matter of justice and so practices that further

or inhibit discrimination raise questions of justice. But information privacy is instrumental to justice in employment and not a requirement of it. A government of Madisonian angels could be trusted with such information without undermining justice, a question that is moot if mere possession of the information is unjust. Among mere mortals, medical and disability information can be used where it is relevant to the ability to perform job duties given reasonable accommodations, and genetic information can be used for the administration of benefits or protecting workers from environmental hazards. Employment discrimination law does not see the possession or transfer of information as the harm, but as the means to harm. The remedy is to limit the possession and use of information such that it can only be used in the interests of justice.

This would explain the comfort of many information technology evangelists with the loss of privacy. It is not that they necessarily see privacy as inherently without value. But it is not an inherent right or essential feature of justice. It is valuable just to the extent that it is useful to protect those social features that are seen as inherent rights or essential features of justice. If information privacy is no longer feasible or has evolved out of existence, then one must look elsewhere for such protections—a belief that their own virtue upholds a system of perfect meritocracy, perhaps. If a calculus of utility overwhelms the protective value of information privacy with conveniences gained by giving it up, then it can be cast aside without moral loss—at least for those with the power to influence decisions about a society's information privacy practices. It is easy to be comfortable with the idea that privacy is dead even as one bemoans "the dawn of the age of genetic McCarthyism," as historian of science Sophia Roosth told the Davos panel, when one takes an instrumentalist view in which the central moral consideration behind privacy is that "By and large, tech has done more good than harm," as Seltzer argued (Carter 2015).

To be sure, it is hard to argue that information privacy does not have instrumental value aside from any claims about its intrinsic value. To take but one recent example, Terrell et al.

(2016) found that women contributing to GitHub, an online software repository that provides code management and version control, are less likely to have their code accepted in software development projects than men when the contributors are outside of the project and their gender is known, but more likely to have code accepted than men when their gender is not identifiable. The privacy of the code contributors in this case clearly supports more equitable outcomes in the open source software development community. One might further argue that the instrumental dimension of justice in information privacy is underappreciated by the theories of privacy that are analyzed below; as Schlosberg argues in his criticism of theories of justice generally. An understanding of the necessary, sufficient, and contributory social conditions for justice should be as much a part of its study as the abstract principles of it; information privacy is not trivialized by suggesting that it has value for more than its own sake. But this is not at all to say, as those like Zuckerberg seem to, that privacy is valuable only instrumentally, that it can be disregarded if justice can be achieved by other means or if does not further justice. A justice-centered view of information privacy—especially one that takes the practice of privacy as seriously as the philosophy of it—certainly must consider the instrumental value of privacy but, as the following two sections demonstrate, this view is inadequate in itself.

3 DISTRIBUTING INFORMATION PRIVACY

In philosophy, law, technology, and business practice contemporary approaches to privacy are predominantly built around controlling access to information. Smith et al. characterize privacy concerns as:

> grounded in the growing “art of the possible” in the technological realm. The spread of ubiquitous computing and the seemingly unbounded options for *collecting, processing, distributing, and using* personal information trigger consumer worries,

and define information privacy strictly in terms of "*access* to individually identifiable personal information (2011, 990, emphasis added)." Similarly, Lita van Wel and Lambèr Royakkers argue, "Informational privacy mainly concerns the control of information about oneself. It refers to the ability of the individual to protect information about himself (2004, 130);" implicitly, one controls others' access to information about one's self and protects information about one's self from being disclosed to others. To frame information privacy as a matter of access is thus to make it a matter of the distribution of information, and thus justice in information privacy is a matter of distributive justice.

Distributive justice, the dominant philosophical framework for understanding justice, is focused on the just distribution of material and social goods. It is telling that the two major collections reviewing the state of social and political philosophy over the past 25 years (Goodin and Pettit 1993; Gaus and D'Agostino 2013) both include chapters or sections on specifically distributive justice but not on other forms, as if this is the only form that justice takes. Kolm (1993) holds that questions of distributive justice arise when the issue is how to arbitrate among competing claims by opposing groups; while this most often concerns scarce material goods, theorists such as Rawls (2005) have applied distributive frameworks to primary social goods such as liberty, political rights, and social positions as well. One can approach distributive justice from two perspectives; the most common in the study of justice holds that a distribution is just if the pattern of distribution that results from the distributive processes is just (regardless, for the most part, of what the distributive process looks like). This approach is less consistent with most theories of privacy however. A minority of distributive theorists focus on the distribution arising historically through a just process and hold that any distribution actually resulting from a just process is thereby just regardless of what that distribution looks like, an approach more consistent with the dominant view of privacy.

3.1 Process Distributions

In process theories of justice, the process by which a distribution actually arises historically is an end in itself; the actual pattern of distribution of goods resulting from it is not relevant to evaluation. Robert Nozick's theory of justice is exemplary of process theories of distributive justice.[5] For Nozick, "a distribution is just if it arises from another just distribution by legitimate means. . . . Whatever arises from a just situation by means of justice is itself just (2013, 151)." Contra Rawls, the end state is irrelevant; a distribution that could be reached by just means but was not is not just (the thief, who could have received goods from the victim by gift rather than theft but did not, is the paradigmatic case). Nozick posits two dimensions of justice in holdings. A principle of justice in original acquisition holds that "A person who acquires a holding in accordance with the principle of justice in acquisition is entitled to that holding"; a principle of justice in transfer of holdings holds that "A person who acquires a holding in accordance with the principle of justice in transfer, from someone else entitled to the holding, is entitled to the holding." By iterative application of these two principles, goods can be continuously redistributed in a just society without need for authoritative reallocation to remedy deviations from the just end state: "The complete principle of distributive justice would say simply that a distribution is just if everyone is entitled to the holdings they possess under the distribution (Nozick 2013, 150–151)."

[5] "Process theories of distributive justice" are distinct from distributive theories of procedural justice. The former determine the distribution of material and social goods by repeated application of a just process over the history of a good or society, while the latter determine entitlements to political, legal, or social process by according to the principles of just distribution of such entitlements. In general, it is initially sufficient for the purposes of this section to regard common uses of the term *procedural justice* as cases of the latter rather than as a distinct form of justice. But while this is a common interpretation of the idea, the implication of section 4 is clearly that procedural justice is better understood as a set of social structures, relations, and processes that guarantee one's ability to participate in determining one's actions and their circumstances, and thus as a species of structural justice as described in that section.

Nozick's principles of justice in original acquisition and transfer are primarily rooted in property rights and market mechanisms. Liberty is a common basis for principles of just transfer, with such theories returning often to Locke's claim that the state of nature is "a state of perfect freedom to order their actions, and dispose of their possessions and persons, as they think fit, within the bounds of the law of nature, without asking leave, or depending upon the will of any other man (1980, sec. 4)." This leads to the purely economic theories of Hayek (1994) or Friedman (2002), and is most easily applied in the distribution of material goods than of social goods. But the claim that justice in transfer is rooted in liberty of action is also important to many areas of social life in which consent is central; one could interpret the crime of rape, for example, as a failure to justly transfer a right to sexual interaction under a consent-based theory of justice in transfer.

Many theories and practices of privacy can be understood from a process-distributive perspective. In this view, privacy is protected to the extent that transfers of information are limited to those permitted under some principle of justice in transfer. The simplest form of this is seen in corporate privacy policies, in which consent is the implied principle of justice in transfer. The principle behind a publically available privacy policy is that consumers can understand how the receiver of the information intends to collect and use information about the consumer, allowing the consumer to choose whether to provide the information either directly, or by undertaking actions such as making a purchase or visiting a web site that will allow the receiver to collect the information in the course of the action. The information collected in accordance with, for example, Google's privacy policy is not a privacy violation because it was collected with the informed consent of the consumer.

This approach to privacy is present in another set of federal laws regarding information privacy in the United States, one that developed somewhat later than employment discrimination law but that nonetheless had enough overlap to suggest the complexities of information privacy in U.S. legislation. The Family Educational Rights and Privacy Act of 1974

(FERPA) and the Health Insurance Portability and Accountability Act of 1996 (HIPAA), together with their associated implementing regulations, are philosophically different than the employment protections described in section 2. Here, we see a series of information privacy acts in which privacy is an end in itself rather than an instrumental protection of some other justice concern.

FERPA bars, in principle, funding to "any educational agency or institution which has a policy or practice of permitting the release of education records (or personally identifiable information contained therein…) of students without the written consent of their parents (20 U.S. Code § 1232g(b)(1))," with the consent rights shifting to the students themselves if they are at least 18 years old or attend a post-secondary institution (subsection (d)). Paragraph (1), of course, includes the usual lengthy list of exceptions for school and government program administration. Its implementing regulations' purpose is explicitly the protection of privacy (34 C.F.R. 99.2), and the largest category of exception, the disclosure of "directory information," is defined as "information contained in an education record of a student that would not generally be considered harmful or an invasion of privacy if disclosed (34 C.F.R. 99.3)." HIPAA does not specifically define standards but rather required the Secretary of Health and Human Services (HHS) to develop such standards, which led to the HHS Privacy Rule. The rule permits insurers and health care providers "to use and disclose protected health information" only for specific purposes related to medical practice and payment, prohibits the sale of such information, and requires that such discloses be limited to the minimum necessary information for the purpose of the disclosure. It exempts de-identified information from the rule entirely, but protects the information for 50 years after a person's death (45 CFR 164.502).

These are significantly different approaches to what is seen in employment discrimination law. FERPA does not at all restrict the collection of information about students by educational agencies or institutions, either formally as does the ADA or GINA or by presenting litigation risks as under the Civil Rights Act. Nor does it contain any restrictions on who can hold

such information or under what circumstances they can hold it, as the ADA and GINA do. FERPA simply requires that any transfer of information outside of the exceptions be done with the student's or parent's consent. The same is true of HIPPA's relationship with insurers and health care providers, who can collect information virtually at will but operate under even more significant restrictions on the transfer of individually identifiable information. HIPAA demonstrates even more clearly its purpose in protecting privacy as an end in itself with its wide scope for the use of de-identified information and protections of privacy even after death. FERPA and HIPPA are not instrumental protections of some other goal. They act to create an operational principle of justice in transfer for educational and medical information.

A more complex version of the consent principle holds information privacy to be a commodity that consumers exchange for services. Individuals determine the value of services received—for example, free email from Google—and choose whether the value of the services exceeds the value of the information that they will provide to the service provider. If the individual determines that Gmail is a good enough product to justify allowing Google to read one's email and use that to target advertisements, they "purchase" the service with information rather than money. (Campbell and Carlson 2002) The service provider thus gains a right to the consumer's information through a just transfer (i.e., a consensual and mutually beneficial exchange of considerations) of information from the consumer, a right constrained by the specific terms of the exchange. In either the pure consent or the commodity exchange framework, information privacy constitutes an injustice to the extent that information is collected or used without or in ways contrary to the consent of the individual from whom it is collected.

Process theories of distributive justice add considerable depth to both consent and commodity theories of privacy. Common dissatisfactions with privacy policies often focus on violations of an implicit theory of just transfer. Various failures of informed consent such as the length, complexity, accessibility, and incomprehensibility of the policies; the inability of consumers to negotiate the terms of the policy with receivers; or the necessity of the

transactions governed by the policy entail that the distribution of information has arisen from a (presumably) just initial acquisition of information (i.e., the individual has all of the information about them and others have none) but undermined by unjust transfers. The Apple iTunes Terms of Service (2015) agreement, for example, runs nearly 21,000 words, suggesting a reading time of between 98 and 136 minutes without accounting for the difficulty of the text (Trauzettel-Klosinski and Dietz 2012). Analyzing it in Microsoft Word 2013 shows it has a Flesch-Kincaid Grade Level score of 16—a college graduate—and a Flesch Reading Ease score of 31.4, putting it above the reading comprehension level of the more than 70% of Americans without a bachelor's degree. Considering consent explicitly as a theory of just transfer calls attention to the need to articulate specifically the conditions under which consent to information transfer is meaningful at a level of specificity similar to that developed to govern informed consent in medical practice. A legalistic notion of consent might dismiss these concerns with a simple "caveat emptor," but a justice-centered view of privacy must take seriously whether the capacity to consent is absent from such a case and whether that violates a principle of just transfer of information.

The process approach to justice also exposes a serious weakness in consent-based and especially commodity models of privacy. Few such models offer a serious attempt to explain the principle of justice in original acquisition at work. FERPA and HIPAA assume that the information held by educational institutions and health care providers is held legitimately according to some unstated principle of justice in original acquisition.[6] It is only in the transfer (and in some cases under HIPAA, use) of information that privacy protections come into play

[6] To be sure, this does not imply that the laws assume it is held accurately. FERPA grants extensive access rights to educational information to the student or their parents, and requires institutions to have formal processes in place that allow students to challenge and correct inaccurate or misleading information. One can consider this at best a minor element of an otherwise absent theory of justice in original acquisition, to the effect that inaccurate information is not held justly. It is surely inadequate in itself as a principle of just acquisition.

under these privacy laws. Campbell and Carlson make a compelling and critical argument that the transfer of information beyond that immediately necessary for a transaction is rooted in a kind of panoptic self-surveillance in which the threat of exclusion from market benefits entices cooperation, but cite those who hold that "users should willingly allow for information collection in return for economic benefits (2002, 593)." Both sides, however, start from the premise that information beyond minimum transactional data is originally the property of the individual. It is not at all clear that this is so, or that it is adequate as a principle of just original acquisition. Some kinds of inferences about individuals' personal information are central to most social interactions; traditionally at least, in formal English it is necessary to know the gender of someone one wishes to refer to in the third person even where their[7] gender is unknown or they may not see the gender binary as meaningful to them. It seems strained to say that information about one's gender is the "property" of that individual and cannot be transferred without that individual's consent.

At the same time, a theory of privacy based on procedural distributive justice seems to give the individual no rights with regard to such information as is minimally necessary for a transaction and blurs the line between transfer and acquisition. No reasonable theory of just

[7] The use of *they* as a gender neutral, indeterminate gender, or gender diverse singular pronoun is increasingly common as a response to exactly this problem. The singular they has a long history in English usage; it is nearly ubiquitous in spoken English, quite common in the writings of prominent authors, and came to be stigmatized by prescriptive grammarians only in the 18th Century. As Jürgen Gerner notes, "As there are no third-person personal, possessive, or reflexive pronouns in English which are both singular and gender-neutral, complete grammatical agreement is not possible [when the antecedent is both singular and not marked for gender, e.g., *anyone*]. The choice is between a violation of gender concord or a violation of number concord (Gerner 2000)." That number agreement is the obvious answer is now increasingly questioned

That this is (and, with the exception of the two-century interlude of prescriptive grammarians insisting that the only objective solution was number agreement, always has been) acceptable, however, does not solve the problem of ownership of one's gender information; it merely expands the range of options that a process theory of justice allows one to own. That one prefers *they*, or even *ze*, remains within the scope of owned information that then requires a theory of just transfer before one can speak of the person. This seems an exceptionally complicated way of understanding the matter.

original acquisition could conceivably argue that an online store has not justly acquired records of the items purchased through their site, the dates of purchase, and the means by which the purchase was made. But using only this data, without any transfer of data to another party, retailer Target was able, infamously, to predict which customers were pregnant, a finding that has stirred great controversy among privacy advocates (Hill 2012). I have argued elsewhere (J. A. Johnson 2014a) that so-called "big data" methods can result in privacy violations by creating new data about an individual that the individual would not have revealed themselves, whether or not the actually knew the information. A student might not be willing to disclose their probability of failing a course to their professor at the beginning of the semester; inferring information about membership in protected social categories (e.g., by identifying proxies for race) gets around laws barring the direct collection of such data. But the common sense perspective on acquisition and transfer would see this as unquestionably a matter of acquisition: the data collector has created the data themselves from other data that they held legitimately. Hence what would be a violation of privacy if it arose through transfer is achieved in a technically legitimate way by the creation of new data from data previously transferred, which is to say, by original acquisition. Transfer becomes acquisition, further separating the subject of the data from their rights regarding it.

Theories of consent, driven implicitly by a process approach to justice, only cover justice in transfer; justice in original acquisition is simply assumed. But a privacy framework that derives more explicitly from a process theory of justice will require not principles justifying transfer but justifying acquisition as well. Helen Nissenbaum's (2010) idea of contextual integrity can be interpreted as a partial solution here. Nissenbaum argues that the context of information flows is as important to privacy as the content of them; privacy violations occur when information flows beyond the "values, goals, and ends of the context" in which it exists. One could read the context as an initial limit on justice in original acquisition: the information was acquired not universally but for use in a specific context. To use it beyond that context, whether

by transfer to another party or by the initial possessor, is to hold the information contrary to the principle of justice in original acquisition, a violation of the principles behind process theories of justice, since a just distribution can only arise by repeated operation of the principles of justice in original acquisition and of justice in transfer. Target would thus be justified in using transactional data to execute the transaction and to manage its ability to repeat the transaction (on the assumption that the context of a transaction at Target includes some notion of iteration) but not to generate personal information about the consumer. Certainly, a principle of contextual integrity would be part of a principle of just transfer as well.

3.2 Pattern Distributions

While it is undoubtedly the most common way of engaging information privacy from the perspective of distribution, not all distributive approaches to privacy are oriented toward process. Hartzog and Selinger argue:

> Ideas about preventing the surveillance society from going too far usually focus on three desirable outcomes: (1) prevent certain groups from ever having access to certain types of information; (2) prevent certain groups from being able to use certain types of information in select contexts or in certain ways; and (3) make it harder for certain groups to be able to access or interpret information (2015, 1345).

In each of these outcomes, information flow is restricted not by stating terms under which it can be (un)justly acquired or transferred but by describing a pattern of distribution following the distributive process that is (un)just: one in which certain groups have inappropriate access to information or use it inappropriately. The result is the ability to call a particular pattern of distribution that is the end-state of a distributive process just regardless of the conditions or processes under which it arose historically, and to argue for information privacy practices as pragmatic policy solutions to bring about that distribution.

A theory of distributive pattern justice must generally consider two dimensions of distribution. Consider, for example, the basic utilitarian justice principle of maximizing utility. Kolm suggests that the basic framework for theories of distributive justice in which material and social goods are distributed according to one or more "directly relevant ethical variables" (1993, 438) that may range from a basic political equality to individuals' varying needs for material goods. In the utilitarian framework, the directly relevant ethical variable is utility. However, in spite of his framing it is clear from Kolm's presentation (as well as the wide range of theories of distributive justice reviewed in Gaus and D'Agostino's (2013) eight chapters on distributive justice) that one must consider not simply the variable on which the distribution is based but also the distributive model itself. *Distribute goods according to utility* is itself inadequate, as there are many possible distributions in which utility is the only variable. Hence the utilitarian principle includes a specific distribution as well, that of maximizing utility for the greatest number of individuals. One could equally offer the argument that goods ought to be distributed such that utility was equal across individuals or such that total utility is maximized without violating the principle that utility is the only directly relevant ethical variable. Rawls offers a more complex theory that nonetheless addresses both dimensions, distributing material goods on the basis of a maximin solution (one that maximizes the minimum outcome) but primary social goods on a greatest equality solution (one that maximizes outcomes subject to the constraint that all receive the same outcome), all within an ethical framework with three directly relevant ethical variables (individualism, utility, and consent).

While process is important to some pattern theories of distributive justice, in those theories the process is usually instrumental to either reach or justify a particular pattern of distribution. Rawls' original position and veil of ignorance are well-known procedural solutions to the problem of how to ensure that individuals choosing principles of justice will do so without regard for their own interests. But Rawls is clear that the original position is intended to justify a set of distributive principles:

> The original position is not, of course, thought of as an actual historical state of affairs, much less a primitive condition of culture. It is understood as a purely hypothetical situation characterized so as to lead to a certain conception of justice. (Rawls 2005, 12)

It is a distribution of goods that conforms to the final distributive principles that Rawls identifies through the process of reasoning from the original position, not the process of choosing them, that define the just society. Indeed, the process determines only the principles of justice to which the end distribution must conform and not the actual distribution of goods.

The most straightforward pattern distributive framework for information privacy attempts to define a right to information privacy analogous to existing rights to personal privacy framed, initially, as a "right to be left alone" (Smith, Dinev, and Xu 2011, 994). Such frameworks essentially posit a right to privacy as a primary social good, delineate the basis of the right as the directly ethically relevant variable, and use, presumably, equality as the principle of distribution. Rebecca Greene, for example, argues for a right to obscurity in political information about an individual:

> Political obscurity therefore describes a broader right than anonymity: it is the fundamental right to exist without one's political preferences being continuously recorded and, consistent with the right articulated in [*United States Department of Justice v. Reporters Committee for Freedom of the Press* (489 U.S. 749 (1989))], a right against state-facilitated cataloguing of one's political preferences (Greene 2013, 373–374).

This is a limited right that is equally distributed across all persons by virtue of being inherent in individual autonomy; in its absence, deliberate self-governance becomes exceptionally problematic. By distributing political obscurity as a primary social good individual autonomy and citizen participation is reinforced.

Often apparently rights-based approaches create only legal rights that put into effect an access control regime rather than creating a right in itself. The Canadian Personal Information Protection and Electronic Documents Act (S.C. 2000, c. 5) establishes 10 principles for the protection of personal information to which organizations must comply. Those principles give individuals a range of legal rights against organizations holding information about them but do not create a general right to which one can appeal beyond the specific situations governed by the act. Similarly, the U.S. Department of Health and Human Services states that HIPAA "gives you rights over your health information, including the right to get a copy of your information, make sure it is correct, and know who has seen it" (U.S. Department of Health & Human Services Office for Civil Rights 2013). But these rights largely protect flows of information: they require the disclosure of purposes, the consent of the individual, limited collection and use, and individual access to the information among others. The acts do not create a general right to privacy that can be distributed; they create procedural protections against specific information flows, and so are best understood as either a means of implementing a process approach to justice in information privacy or as an attempt to delineate and operationalize an implicit right to information privacy.

The latter is probably the best interpretation of European Union law on information privacy. The European Data Protection Directive (Council Directive 95/46/EC) is among the earliest and most extensive legal regimes for protecting information privacy specifically as a right. The directive was specifically created to "protect the fundamental rights and freedoms of natural persons, and in particular their right to privacy with respect to the processing of personal data" (art. I, para. 1), in particular aiming to harmonize information privacy protections across member states in the face of increasing flows of personal information, in both government and commerce, across member states (recitals 5, 7, 10, and 11). The right to privacy is not, in fact, articulated substantively, but recital 10 references everyone's "right to respect for his private and family life, his home and his correspondence" under Article 8 of the Convention for the

Protection of Human Rights and Fundamental Freedoms, tying the substance of the Directive to the existing body of privacy rights recognized within the EU. Recital 10 states that the Directive gives substance to and amplifies previous protections of privacy: Data collection and processing under the Directive must ensure both data quality and legitimacy, the latter secure by either consent or limited notions of transactional necessity; in general, processing of sensitive personal data is barred by Article 8 of the Directive; extensive information and access rights to an individual's own data are specified in Articles 10 through 12; and the controllers of data are subject to substantial regulation under Articles 16 through 20. This structure, of regulations that operationalize a principled right to general privacy as it applies to issues of personal information, is preserved in the proposed General Data Protection Regulation expect to take effect in 2016.

Claims of a fundamental right to specifically information privacy in the United States are less well developed. Many organizations concerned with information privacy have asserted that information privacy protected as a legal or moral right. The American Library Association's interpretation of its "Library Bill of Rights" argues that privacy is implicit in the bill's Article IV on resisting "abridgment of free expression and free access to ideas" and cites a chilling effect on those principles from breaches of privacy. But the basis claimed for this right relies primarily on legal precedents related to either receiving information in a library or to general privacy cases (American Library Association 2002) without a clear argument to that effect, which is complicated by the diverse legal bases for privacy claims under U.S. law. It may well be that one's choices to access information is included within the zone of personal behavior free from unreasonable state intrusion, implicit in ordered liberty, basic to a free society (*Mapp v. Ohio* [1961], 367 U.S. 643); is within "penumbras, formed by emanations" from specific guarantees within the U.S. Bill of Rights (*Griswold v. Connecticut* [1965] 381 U.S. 479, 484), or are made with a reasonable expectation of privacy (*Katz v. United States*, 389 U.S. 347 (1967) Harlan, J., concurring). But which doctrine forms the basis of information privacy and how the connection is made will certainly matter for the practical contours of a right to information privacy.

U.S. courts appear increasingly willing to accept such claims, but a definitive doctrine has yet to emerge. The U.S. Supreme Court's recent decision in *Riley V. California* (573 U.S. ____ (2014); Docket No. 13-132) held that the information contained on a mobile phone enjoyed the protection of the Fourth Amendment's requirements for reasonableness in searches and thus could not be searched incident to the arrest of a person in possession of the phone. Observing that mobile phones "are now such a pervasive and insistent part of daily life that the proverbial visitor from Mars might conclude they were an important feature of human anatomy," the Court relied on the extensive information contained in smartphones to conclude that the digital content of such devices is categorically different from treating them as the kinds of physical objects assumed by the search incident to arrest doctrine articulated in *Chimel v. California* (395 U.S. 752 (1969)), *United States v. Robinson* (414 U.S. 218 (1973)), and *Arizona v. Gant* (556 U.S. 332 (2009)). With mobile phones, the Court held, "The sum of an individual's private life can be reconstructed through a thousand photographs labeled with dates, locations, and descriptions" as "the more than 90% of American adults who own a cell phone keep on their person a digital record of nearly every aspect of their lives—from the mundane to the intimate," a situation aggravated by the fact that a mobile phone is connected to data stored elsewhere through cloud applications. The court here clearly presents a defense not of personal or general privacy but of information privacy specifically, distinguishing between the permissible physical search of a mobile phone to determine if it poses a physical threat (e.g., if the arrestee concealed a weapon such as a razor blade in it) and a search of the information contained on it.

This remains, however, a weak theory of information privacy. It springs from general privacy; the question before the court is how one specific legal doctrine, the search incident to arrest rule, applies to the information carried on a device. The court links the issue to the traditional defense of Fourth Amendment protections, that of an overreaching exercise of police powers by the state. The court's ruling does not lead to any definitive implications beyond criminal process. For example, the court remains silent on whether there is a reasonable

expectation of privacy in the contents of a mobile phone or simply whether the privacy implications are sufficient to overcome a weak case on behalf of the government's interest in the traditional justifications for searches incident to arrest, concluding that a search supported by a warrant is the proper course because that is the general principle of reasonableness in searches.[8] That would be important in determining, for example, whether the state can demand the right to search one's mobile phone in the course of a voluntary administrative process where significant discretion is allowed, such as approving membership in any one of the U.S. government's Trusted Traveler programs. And it is of course says nothing about intrusions on privacy from commercial actors. Google, among many others, already maintains "the sum of an individual's private life" for every individual who uses a Gmail account or an Android phone. Amazon's examination of the "digital record of nearly every aspect of their lives" may be less problematic than the state's, but it is by no means obviously unproblematic.

I will argue below that viewing rights from a distributive perspective is especially problematic. But it seems puzzling that other distributive approaches to privacy are nearly nonexistent. This is certainly not because such claims cannot be made. Rawls' first principle of justice is that "each person is to have an equal right to the most extensive basic liberty compatible with a similar liberty for others" (2005, 60). Rawls here treats rights and liberties as essentially synonymous, so we could conceive of privacy as part of a system of basic rights and

[8] There is good reason to believe a general right to privacy with regard to information might flow from this decision. The court's recognition of the scope of information available through a mobile phone, including information held by third parties, was a major part of its reasoning that a substantial privacy interest was implicated. This would challenge the major obstacle to a broad doctrine of information privacy in existing case law, the third-party doctrine (*Smith v. Maryland*, 442 U.S. 735 [1979]), by recognizing that individuals retain a privacy interest in data shared with a third party service provider. The U.S. Court of Appeals for the Second Circuit, in considering the legality of bulk telephone metadata collection in *ACLU v. Clapper* (No. 14-42-cv, slip op. at 83-90), suggested (in a judicial dictum, as it had already disposed of the case on non-constitutional grounds) that the third party doctrine is brought under new scrutiny by the technology of bulk collection, relying in part on the concurring opinions of Justices Alito and Sotomayor in *U.S. v. Jones* (132 S. Ct. 945 [2012]) arguing that Global Positioning System tracking on a vehicle exceeds the reasonable expectation of privacy.

liberties that ought to be maximized subject to the constraint that all have the same liberty. This may seem like a trivial or even substanceless requirement, but it actually illuminates a quite serious problem in information technology, that of online harassment and domestic violence, especially toward women.

In 2013, after feminist critics of video game culture raised concerns about the depiction of women (and lack thereof) in games, gamers responded with an astounding level of online harassment of their critics. The harassment, almost entirely either anonymous or pseudonymous, included quite specific rape and death threats (for example, including critics' addresses and the times they would be assaulted), driving at least three women from their homes and prompting FBI investigations. An especially serious tactic used to silence critics was that of "doxing," or publishing personal information about the critics that would make them vulnerable to violence or harassment in the physical world (Dewey 2014). One critic, Anita Sarkeesian, had to cancel a speech at Utah State University after the university, constrained by state law prohibiting it from barring legally carried concealed weapons on campus, could not respond to an anonymous threat to carry out "the deadliest school shooting in American history" targeting both the lecture and the university women's center if she spoke (McDonald 2014). The threat read, in part:

> Anita Sarkeesian is everything wrong with the feminist woman, and she is going to die screaming like the craven little whore that she is if you let her come to USU. I will write my manifesto in her spilled blood, and you will all bear witness to what feminist lies and poison have done to the men of America. (Neugebauer 2014)

Supporters of Gamergate counter that the movement supports ethics in video game journalism.

To be sure, there are many deep problems with Gamergate—some of which I address in section 4—that contributed to this harassment. Among them is an asymmetry of privacy. The critics of games could not make their critiques from a position of anonymity or pseudonymity;

their critiques are part of their professional identities and were unavoidably done from a public stance. In order for Zoe Quinn, Brianna Wu, and Sarkeesian to offer their criticisms publicly as professionals in the game development industry they must have adopted a public identity with all of the constraints and risks that entails. The lack of such an identity is equally essential to their harassers. Protected by anonymity, both personal and technical, they could level threats that were of an unquestionably criminal nature (indeed, as the example above demonstrates, an unquestionably inhuman nature) without substantial risk. The Gamergate harassers thus exercise a degree of privacy not available to their critics. A Rawlsian principle of "greatest equal privacy" could provide guidance to those developing systems that could mitigate the asymmetry of privacy, taking away the protection that anonymity provides online harassment and domestic violence. Facebook, for example, maintains a real name policy that requires users to register and use their real names rather than pseudonyms in part to cut down on online harassment. One might make similar arguments about privacy from the perspective of need or capabilities (Robeyns 2013; Brock 2013).

Of course, such a principle would be open to challenge as well, with anonymity being a core principle of privacy, especially in technical solutions. If the greatest possible amount of privacy compatible with all having a similar privacy prohibits anonymity entirely, then the critics who argue that privacy is a dated concept have a strong argument. The Electronic Frontier Foundation considers the TOR web browser, which routes web browsing through multiple relays making users anonymous to the sites they visit, an essential privacy tool. TOR might allow a user to gather information about protest movements without being subject to state surveillance, but it would also allow anonymous access to an email account from which a threat like those made to Sarkeesian could be sent without fear of legal consequences. These kinds of tradeoffs are part of why EFF considers threat modeling an essential practice for surveillance self-defense. That, however, further undermines a greatest *equal* privacy principle, since different threats will entail different levels of privacy, not only for end users but for system designers.

Perhaps a needs or capabilities approach would be an improvement here, though it is not entirely clear that this is so. Ultimately this may suggest that there are serious limits to what pattern approaches to distributive justice can offer a theory of privacy.

Finally, one notices a significant shift in the object of distribution in moving from process to pattern concepts of justice. In the former, the concern is with the distribution of information, while in the latter the concern is the distribution of privacy rights. It is not at all clear that they are the same thing. Distributions of information present an intersection of concerns with privacy and with open data (J. A. Johnson 2014b) in that aims of controlling data flows in the name of privacy directly conflict with commitments to expanding them in the name of transparency, a problem that suggests a solution of balancing the competing concerns. Distributing privacy rights, however, delineates a clear obligation that trumps the practicality arguments that prevail among advocates of open data. Sunshine may be the best disinfectant, but it certainly harms the patient. The two approaches to justice give very different answers to how tolerant of such harms we should be.

The process framework addresses information rather than privacy in part because it must. This is more than a definitional issue in which privacy is defined as a question of distributing information. The process approach simply does not work effectively when the issue at question is a right held to be universal and inalienable. In such cases, the principles of justice in acquisition and in transfer become trivial: the right is originally acquired when one achieves personhood (in whatever capacity one wishes to use, but in any case well before one can make decisions about one's privacy related to one's original acquisition) and the right cannot be transferred under any circumstances. In order for a process approach to distribute privacy rights as such rather than information, it must first show that privacy rights are somehow alienable. To say that they are wholly so does great violence to the concept of rights generally, so the question then becomes what aspects of one's privacy rights can be transferred, a question that

is, in practice, no different than stating the kinds of information that can be justly transferred and the circumstances under which it can be transferred.

It is less clear that one could not develop a pattern theory that distributes information rather than privacy rights, but it is no easy task. Though not focused on justice in any strong sense and taking a wide range of positions on privacy, the open data movement is very much one positing an ideal pattern of information distribution. This might be an especially good focus for distributive theories rooted in need of capabilities (See, e.g., Britz et al. 2012). But information privacy poses an unusual challenge to pattern theories of justice. In most cases, pattern theories of distributive justice assume that the goods being distributed are in principle either universal (in the case of primary social goods) or scarce (in the case of material goods). In these cases, the problem is to ensure that everyone receives their fair share of the goods in question. To be sure, this is a problem in information justice more broadly, as those questioning the justice of open data argue (Slee 2012; Raman 2012; Donovan 2012; Gurstein 2011). But information privacy poses the opposite problem: restricting the distribution of a good that can be, in the age of electronic reproduction, initially produced and then reproduced infinitely at near-zero cost. A pattern distribution of information that protects privacy would thus need to both create a solution to a novel type of issue that runs counter to existing theories while still supporting such theories in other areas of information practice. It is not at all clear how this might be done other than by restricting the flow of information.

4 BEYOND DISTRIBUTIVE PRIVACY

Concepts of distributive justice have revealed some useful insights into specific aspects of information privacy. In summary:

1. Information privacy is of at least instrumental value in pursuing distributive justice generally.

2. The information flow paradigm must pay attention to the principles for justice in information transfer, as they typically do; for justice in original acquisition, which they typically do not; and for distinguishing between the two.
3. Robust conditions for consent are critical to making the information flow paradigm produce meaningful principles of justice in transfer.
4. Rights-based distributive approaches to information privacy are easily confused with legally established information flow approaches, and are most coherent as systems of information justice when they are used to delineate and operationalize an existing framework of privacy rights.
5. Theories of distributive justice distinct from distributions of rights may be able to provide some useful, practical guidance for information privacy.
6. The justice implications of distributing information and of distributing privacy rights are significantly different.

Nonetheless, distributive justice has not proven itself capable of providing a strong framework for a general theory of information privacy. None of these considerations fundamentally remake our understanding of privacy, nor do they do much by themselves to bring coherence to the field. And they have raised as many problems as they have solved.

The distributive paradigm is not, however, the only way of understanding justice. The same alternatives that Schlosberg found so useful in reconstructing environmental justice can serve as the basis for a useful justice-driven conception of privacy. Young's critique of distributive frameworks of justice provides some insight into why such frameworks seem to offer relatively little to theories of privacy as well as a starting point for a more effective justice-driven approach to privacy. Young's critique focuses on how distributive justice "regards persons as primarily possessors and consumers of goods" rather than considering "action, decisions about action, and provision of the means to develop and exercise capacities (1990, 16)." This leads to two related failures on the part of the distributive paradigm connected to social structure.

The first is that it obscures the structural conditions that underlie the distribution of material goods. The distribution of employment, for example, is a common object for the study of distributive justice, as seen in the analysis of employment discrimination law above. Young argues, however, that asking about the just distribution of jobs tends to assume rather than examine structures like the hierarchical division of labor, social stratification, and commodification that tend to determine the distribution of jobs. In many cases, Young argues, controversies over the distribution of goods are, in fact, controversies over these structures themselves: citizens oppose a hazardous waste treatment plant or a plant closing not because they see the distribution of environmental economic burdens *per se* but because they lack a voice in decisions that affect their lives. Distributive justice fails to capture these aspects of justice in the distribution of material goods, thus restricting the scope of claims to justice (1990, 18–24).

This is a central issue in the question of information privacy, and one that receives scant attention. The distribution of information is a problem, but the problem is not simply a maldistribution of information but also, perhaps even more so, the structures that make it so. Ownership and commodification of information cannot be a solution to information flows because it is at the heart of the problem: Target can identify sensitive medical information that it has no business knowing—there would be no question of it being a privacy violation under HIPAA for Target to have determined which customers were pregnant by buying their medical records—because of the structure of economic activity in a free market system. The unequal positions of enterprise and consumer drive the latter's willingness to be surveilled by the former:

> Clearly, the exchange of information between consumers and suppliers is not equitable, as large corporations do not in the same transaction generally reveal to customers detailed information regarding their internal structure or operations. . . . [I]t is this very inequality in the relationship between consumers and suppliers of goods and services in the marketplace that

> compels individuals to provide personal information. The ability of the producer or supplier to set the terms of the contract that the consumer can only accept or decline defines the transaction as inherently inequitable. . . . The consumer is ultimately a "contract taker, rather than a contract maker," and thus provides the information in the belief that it represents a reasonable transaction cost. . . . [I]ndividuals are not necessarily aware of the degree of inequalities in their relationship with suppliers because marketers and advertisers have effectively concealed the consumerist Panopticon. (Campbell and Carlson 2002, 591–592)

Similarly, state surveillance is unlikely to be seriously addressed with the kinds of legalistic concepts of criminal process rights and limited government seen in *Riley* and in *ALCU v. Clapper* when terrorism, the surveillance state's *raison d'etre*, exists as a state of exception to law (Agamben 2005): the U.S. prison at Guantánamo Bay, Cuba exists solely because it is considered outside of U.S. legal jurisdiction (see *Rasul v. Bush*, 542 U.S. 466 (2004)).

These structural conditions run much deeper than the social contexts in which information is distributed. I have previously argued that information is substantively influenced by a translation regime that transforms underdetermined observations into a single data state (J. A. Johnson 2015). Information exists as a form of communication in which the translation regime encodes the information and a nexus of problems, models, and interventions decodes it. Both the translation regime and the problem-model-intervention nexus are constructed by social actors with social interests in mind, whether deliberately or as unconscious assumptions. This is not just to say that there are errors in information that can be remedied by the kinds of protections seen in FERPA or the European Data Protection Directive; the data is directly constructed by standards that are inevitably biased but, because the standards at least appear to be mechanical and objective, cannot be challenged as erroneous. The result is that

information privacy cannot be about the distribution of information without also being about the sources and construction of that information.

This failure to highlight structural conditions is compounded, Young argues, when distributive justice is extended beyond material goods. The kinds of moral goods that distributive theorists—and privacy rights theorists—address distributively are "better understood as functions of rules and relations than as things," with the result being that distributive justice "tends to preclude thinking about what people are doing, according to what institutionalized rules, how their doings are structured by institutionalized relations that constitute their positions, and how the combined effect of their doings has recursive effects on their lives." Distributive justice cannot conceive of the way that both distributions and structures shape actions except by supposing that actions are constrained only by distributions of goods. Young finds this especially problematic when thinking of power as capable of distribution, as "power is a relation rather than a thing" mediated by structure of agency and actions that obscures especially the ways in which social structures and systems "exclude people from participating in determining their actions or the conditions of their actions" through processes and relationships rather than possession of some abstract concept of power (1990, 24–33).

One sees this most clearly in Greene's otherwise strong argument for political obscurity. Generally, Greene's chief concern is the ability to hold and act on dissenting political views:

> Political obscurity refers to the state of one's political preferences being shrouded or otherwise difficult to discern or distinguish by others. A person enjoys political obscurity when she can go about her day as she so chooses without others perceiving or otherwise determining the nature of her political views. The politically obscure person is *able to control and manage the extent of disassociatio*n from the political views she holds (or once held) or political actions taken in the present and in the past. (2013, 373, emphasis added)

But when Greene translates this concept into a right that can be distributed (as quoted in the previous section, “the fundamental right to exist without one's political preferences being continuously recorded”), the language of action, and in fact the actors involved, disappear. Rather than controlling and managing, the person simply “exist[s]”; the others actively trying to know and influence her political views become a passive, impersonal occurrence that happens to a person.

This, too, is a fundamental inadequacy of a distributive justice-based theory of privacy. A right that is possessed gives no consideration to what one might do with such a right; it is merely distributed and its recipients wished the best of luck with it. Its moral status is not affected if the right should prove inadequate to the purposes for which its recipients intend to use it. In the case of information privacy, consumers’ concerns about privacy are not simply about what data enterprises hold about them, but about how it will be used to further interests that may well conflict with the intentions and plans of the consumers. The memorable vignette from the Target case is of the father who first confronts the store manager about marketing goods for pregnant women and newborns to his teenaged daughter, only to later apologize, saying, “It turns out there’s been some activities in my house I haven’t been completely aware of. She’s due in August. I owe you an apology (Hill 2012).” Target excluded both the father and the daughter from participating in decisions about how they would to respond to a major life event. The wrong is not (or at least not simply) knowing that the daughter was pregnant, or even knowing that she was sexually active, but rather that Target, not the woman, determined the circumstances under which her family would find out about it. That is not simply a question of a right to prevent certain flows of information, but of the basic justice of manipulative marketing from a company who had determined that childbirth was an excellent opportunity to shift someone’s buying habits.

These two criticisms come together to understand why Gamergate cannot be understood as simply a matter of asymmetric privacy. Gamergate was not simply a heated

dispute over any kind of principled matter (especially not ”ethics in journalism”); the threats directed toward Quinn, Wu, and Sarkeesian were not just insults meant to hurt their feelings. Gamergate depends critically on the structure of gender relations in video game culture. It began when Quinn’s ex-boyfriend, in a lengthy diatribe, accused her of being sexually unfaithful, it peaked when someone still unidentified threatened to massacre “the craven little whore” who has poisoned the *men* of America. Doxing them was a threat intended to silence them and maintain a system of structural power that favors men, one that was, due to the legally enshrined hyper-masculine culture of Utah that treats carrying a gun as the *sine qua non* of manhood, successful in at least Sarkeesian’s case. These three women’s information privacy was violated not simply in that personal information was distributed improperly but because information about them was used as a tool of domination and oppression. A justice-driven theory of information privacy will be inadequate if it cannot engage such issues.

There are promising approaches to privacy that are compatible with ideas of structural justice. Approaches treating privacy as a form of obscurity have been quite sensitive to structural issues. Greene’s analysis does not rely on her rights-driven formulation, and in fact offers excellent analysis of the structural features of both information technology and petitioning processes to understand the injustices that she identifies:

> What if the real (and much more difficult to document) harm befell those who did not—or would not—sign the petition? What if the harm in releasing petition names is not to activists being mooned or shouted at as they advocate publicly for their cause? What if the real privacy victim is a mother of two, passing a petition circulator entering the grocery store, fearful that signing a petition—even for a cause in which she very much believes—might create a lifelong indelible association with that cause on her Internet record? (2013, 370)

She points, for example, to the publication of signers of a Maryland petition to bring an anti-same sex marriage measure to the ballot by an LGBT newspaper that resulted in one university terminating its chief diversity officer, who had signed the petition, as an example of how a lack of political obscurity undermines the capacity to act politically. Information technologies that undermine political obscurity by making possible frictionless gathering of information about people's political beliefs have a chilling effect on political participation and action.

Similarly, the broader idea of obscurity as a protection against state surveillance that Hartzog and Selinger (2015) argue for is rooted very much in an analysis of processes and relationships of power rather than distributive concepts. They define obscurity as making information hard, but not impossible, to find or interpret, arguing that this quality makes information "safe." Explicitly, they seem to mean safe from exposure: "when information is difficult to acquire or burdensome to interpret, the only people who will be inclined to do the detective work are those who deem the expense an acceptable cost." But implicitly, obscurity preserves the safety of its subject by "mak[ing] it hard (or harder), but not impossible, to discover irrelevant, inadequate, and embarrassing details" that would limit one's ability "to manage the accessibility and comprehension of social exchanges by outsiders, the loss of which can be quite harmful." It makes the actions of citizens less intelligible, which one might interpret as acting as a counter to the state's interests in legibility (Scott 1998). Obscurity thus functions as a structural rather than a legal right, protecting interests because social structures prevent their violation.

Neither Hartzog and Selinger nor Greene defend obscurity as a distributive right, nor do they frame it as a question of justice in information transfer. Indeed, Hartzog and Selinger argue quite explicitly that they do not aim to bar such transfers but simply to make them more difficult in order to shape relationships among actors and facilitate certain kinds of action for certain actors. And they can show how the structural conditions in which information is made and distributed affect one's ability to both develop one's capacities as a person and to participate in

determining one's actions. One might draw a similar conclusion about Daniel Solove's problems-based approach to privacy: "A privacy invasion interferes with the integrity of certain activities and even destroys or inhibits some activities (2008, 8)." His imagery of the contrast between Orwell and Kafka reminds one that the harm of surveillance is not just the threat of discipline from Big Brother but also the feeling of being powerless against systems that shape one's life.

5 CONCLUSION

Young's critique, and the compatibility of some of the more successful approaches to privacy, suggest the virtue of a structural justice approach to information privacy. Information privacy is rife with unanswered questions of distributive justice. Privacy at least contributes to achieving just distributions of social goods. Information flow models of privacy assume answers to questions about justice in acquisition and transfer that may be indefensible on their own or incompatible with each other. Rights approaches often assume rather than articulate a justification of privacy itself. And there are considerable implicit differences between the two in what is to be distributed. It should not be seen as practical to address questions of information privacy without considering these questions. And Young's critique of the distributive paradigm reveals deeper problems with understanding the question of justice in information privacy. Information privacy is as much a matter of social structure as it is of distributing material or moral goods, and the focus on distribution obscures the ways in which information privacy violations challenge the ability to participate in determining one's actions. These critiques suggest that a more productive line of inquiry would be to pursue information justice as a matter of primarily structural rather than distributive justice.

6 REFERENCES

Agamben, Giorgio. 2005. *State of Exception*. Chicago: University of Chicago Press.

American Library Association. 2002. “Privacy: An Interpretation of the Library Bill of Rights.” HTML file. June 19. http://www.ala.org/Template.cfm?Section=interpretations&Template=/ContentManagement/ContentDisplay.cfm&ContentID=132904.

Apple Inc. 2015. “iTunes Store - Terms and Conditions.” October 21. http://www.apple.com/legal/internet-services/itunes/us/terms.html.

Britz, Johannes, Anthony Hoffmann, Shana Ponelis, Michael Zimmer, and Peter Lor. 2012. “On Considering the Application of Amartya Sen’s Capability Approach to an Information-Based Rights Framework.” *Information Development*, August. doi:10.1177/0266666912454025.

Brock, Gillian. 2013. “Needs and Distributive Justice.” In *The Routledge Companion to Social and Political Philosophy*, 444–55. New York: Routledge.

Campbell, John Edward, and Matt Carlson. 2002. “Panopticon.com: Online Surveillance and the Commodification of Privacy.” *Journal of Broadcasting & Electronic Media* 46 (4): 586–606. doi:10.1207/s15506878jobem4604_6.

Carter, Richard. 2015. “Privacy Is Dead, Harvard Professors Tell Davos Forum.” *Yahoo! Tech*. January 22. https://www.yahoo.com/tech/privacy-dead-harvard-professors-tell-davos-forum-144634491.html.

Dewey, Caitlin. 2014. “The Only Guide to Gamergate You Will Ever Need to Read.” *The Washington Post*, October 14, sec. The Intersect. https://www.washingtonpost.com/news/the-intersect/wp/2014/10/14/the-only-guide-to-gamergate-you-will-ever-need-to-read/.

Donovan, Kevin. 2012. “Seeing Like a Slum: Towards Open, Deliberative Development.” SSRN Scholarly Paper ID 2045556. Rochester, NY: Social Science Research Network. http://papers.ssrn.com/abstract=2045556.

Friedman, Milton. 2002. *Capitalism and Freedom*. 40th anniversary ed. Chicago: University of Chicago Press.

Gaus, Gerald F., and Fred D’Agostino, eds. 2013. *The Routledge Companion to Social and Political Philosophy*. Routledge Philosophy Companions. New York: Routledge.

Gerner, Jürgen. 2000. “Singular and Plural Anaphors of Indefinite Personal Pronouns in Spoken British English.” In *Corpora Galore: Analyses and Techniques in Describing English: Papers from the Nineteenth International Conference on English Language Research on Computerised Corpora (ICAME 1998)*, edited by John M. Kirk. Amsterdam; Atlanta, GA: Rodopi.

Goodin, Robert E., and Philip Pettit, eds. 1993. *A Companion to Contemporary Political Philosophy*. Blackwell Companions to Philosophy. Oxford, UK ; Cambridge, Mass: Blackwell.

Greene, Rebecca. 2013. “Petitions, Privacy, and Political Obscurity.” *Temple Law Review* 85: 367–411.

Gurstein, Michael. 2011. “Open Data: Empowering the Empowered or Effective Data Use for Everyone?” *First Monday* 16 (2). http://firstmonday.org/htbin/cgiwrap/bin/ojs/index.php/fm/article/view/3316/2764.

Hartzog, Woodrow, and Evan Selinger. 2015. “Surveillance as Loss of Obscurity.” *Washington and Lee Law Review* 73 (3): 1343–87.

Hayek, Friedrich A. von. 1994. *The Road to Serfdom*. 50th anniversary ed. / with a new introd. by Milton Friedman. Chicago: University of Chicago Press.

Hill, Kashmir. 2012. *How Target Figured Out A Teen Girl Was Pregnant Before Her Father Did*. http://www.forbes.com/sites/kashmirhill/2012/02/16/how-target-figured-out-a-teen-girl-was-pregnant-before-her-father-did/.

Johnson, Bobbie. 2010. “Privacy No Longer a Social Norm, Says Facebook Founder.” *The Guardian*. January 10. http://www.theguardian.com/technology/2010/jan/11/facebook-privacy.

Johnson, Jeffrey Alan. 2014a. “The Ethics of Big Data in Higher Education.” *International Review of Information Ethics* 21 (July): 3–10.

———. 2014b. “From Open Data to Information Justice.” *Ethics and Information Technology* 16 (4): 263–74. doi:10.1007/s10676-014-9351-8.

———. 2015. “Information Systems and the Translation of Transgender.” *TSQ: Transgender Studies Quarterly* 2 (1): 160–65. doi:10.1215/23289252-2848940.

Kolm, Serge-Christopher. 1993. “Distributive Justice.” In *A Companion to Contemporary Political Philosophy*, 438–61. Oxford, UK ; Cambridge, Mass: Blackwell.

Locke, John. 1980. *Second Treatise of Government*. Edited by C. B. Macpherson. 1st ed. Indianapolis, Ind: Hackett Pub. Co.

McDonald, Soraya Nadia. 2014. “‘Gamergate’: Feminist Video Game Critic Anita Sarkeesian Cancels Utah Lecture after Threat.” *The Washington Post*, October 15, sec. Morning Mix. https://www.washingtonpost.com/news/morning-mix/wp/2014/10/15/gamergate-feminist-video-game-critic-anita-sarkeesian-cancels-utah-lecture-after-threat-citing-police-inability-to-prevent-concealed-weapons-at-event/.

Neugebauer, Cimaron. 2014. “Terror Threat against Feminist Anita Sarkeesian at USU.” *Standard Examiner*, October 15. http://www.standard.net/Police/2014/10/14/Utah-State-University-student-threatens-act-of-terror-if-feminist.

Nissenbaum, Helen. 2010. *Privacy in Context: Technology, Policy, and the Integrity of Social Life*. Staford, California: Stanford Law Books.

Nozick, Robert. 2013. *Anarchy, State, and Utopia*. New York: Basic Books, a member of the Perseus Books Group.

Plato. 1991. *The Republic of Plato*. Translated by Allan Bloom. 2nd ed. New York: Basic Books.

Raman, Bhuvaneswari. 2012. "The Rhetoric of Transparency and Its Reality: Transparent Territories, Opaque Power and Empowerment." *The Journal of Community Informatics* 8 (2). http://ci-journal.net/index.php/ciej/article/view/866/909.

Rawls, John. 2005. *A Theory of Justice*. Original ed. Cambridge, Mass: Belknap Press.

Robeyns, Ingrid. 2013. "The Capability Approach (And Social Justice)." In *The Routledge Companion to Social and Political Philosophy*, 446–66. New York: Routledge.

Schlosberg, David. 2004. "Reconceiving Environmental Justice: Global Movements And Political Theories." *Environmental Politics* 13 (3): 517–40. doi:10.1080/0964401042000229025.

Scott, James C. 1998. *Seeing like a State: How Certain Schemes to Improve the Human Condition Have Failed*. New Haven: Yale University Press.

Slee, Tom. 2012. "Seeing Like a Geek." *Crooked Timber*. June 25. http://crookedtimber.org/2012/06/25/seeing-like-a-geek/.

Smith, H. Jeff, Tamara Dinev, and Heng Xu. 2011. "Information Privacy Research: An Interdisciplinary Review." *MIS Quarterly* 35 (4): 989–1015.

Solove, Daniel J. 2008. *Understanding Privacy*. Cambridge, Mass: Harvard University Press.

Terrell, Josh, Andrew Kofink, Justin Middleton, Clarissa Rainear, Emerson Murphy-Hill, and Chris Parnin. 2016. "Gender Bias in Open Source: Pull Request Acceptance of Women versus Men." doi:10.7287/peerj.preprints.1733v1.

Trauzettel-Klosinski, Susanne, and Klaus Dietz. 2012. "Standardized Assessment of Reading Performance: The New International Reading Speed Texts IReST." *Investigative Opthalmology & Visual Science* 53 (9): 5452. doi:10.1167/iovs.11-8284.

U.S. Department of Health & Human Services Office for Civil Rights. 2013. "Your Health Infiormation Privacy Rights." Portable Document Format file. February 7.

http://www.hhs.gov/sites/default/files/ocr/privacy/hipaa/understanding/consumers/consumer_rights.pdf.

van Wel, Lita, and Lambèr Royakkers. 2004. “Ethical Issues in Web Data Mining.” *Ethics and Information Technology* 6 (2): 129–40. doi:10.1023/B:ETIN.0000047476.05912.3d.

Young, Iris Marion. 1990. *Justice and the Politics of Difference*. Princeton, N.J: Princeton University Press.