The Particularized Judgment Account of Privacy

Alan Rubel
Assistant Professor
School of Library and Information Studies
Program in Legal Studies
University of Wisconsin, Madison
arubel@wisc.edu

Abstract: Questions of privacy have become particularly salient in recent years due, in part, to information-gathering initiatives precipitated by the 2001 World Trade Center attacks, increasing power of surveillance and computing technologies, and massive data collection about individuals for commercial purposes. While privacy is not new to the philosophical and legal literature, there is much to say about the nature and value of privacy. My focus here is on the nature of informational privacy. I argue that the predominant accounts of privacy are unsatisfactory and offer an alternative: for a person to have informational privacy is for there to be limits on the particularized judgments that others are able to reasonably make about that person.

Keywords: privacy, surveillance, information ethics, particularized judgments

Introduction

Privacy questions have become particularly salient in recent years due, in part, to information-gathering initiatives precipitated by the 2001 World Trade Center attacks, increasing power of surveillance and computing technologies, and massive data collection about individuals for commercial purposes. While privacy is not new to the philosophical and legal literature, there is much left to say about the nature and value of privacy. My task here concerns the nature of informational privacy. I argue that the predominant accounts of privacy are unsatisfactory, and I offer an alternative, that for a person to have informational privacy is for there to be limits on the particularized judgments that others are able to reasonably make about that person.

Please cite to published version: *Res Publica* 17(3): 275-290 (2011).

Background

'Privacy' is conventionally used in many contexts, though it is controversial where it is appropriately applied. Certainly privacy concerns both information and observations. But there is significant dispute as to whether privacy applies to decisions (e.g. having an abortion).[1] My focus is the nature of informational privacy; I do not address the relation between informational privacy and persons' carrying out decisions without interference. Note, too, that my target here is the descriptive sense of privacy—when we can properly say that a person has privacy, regardless of whether she ought to have privacy. There is also 'private' in the normative sense—when one should be able to prevent others from gathering information. The normative sense of 'private' is best understood in terms of when something ought to be descriptively private or when one ought to be able to decide whether something will be descriptively private. Though my task here is a descriptive analysis, I address implications of the account for privacy's value below.

Access Definitions

Privacy is generally understood in terms of either access or control. A central feature of access views is that one's privacy is conceptually independent from her power to maintain it. Privacy has to do with others' actual access, or ability to access, a person. William Parent, for example, argues that privacy is 'the condition of not having undocumented personal knowledge about one possessed by others' (Parent 1983, p.269). Parent understands 'personal knowledge' to include facts (i.e. not falsehoods or opinions) that people would not want revealed. Ruth Gavison proffers a broader view of access: 'an individual enjoys *perfect* privacy when he is completely inaccessible to others' (Gavison 1984, pp.349-350, italics in original). Perfect privacy demands that no information is known about an individual, no attention is paid to that individual, and no one has physical access to that individual. Similarly, Anita Allen maintains that saying 'a person possesses

[1] See DeCew (1997, pp.46-60); Allen (1988, pp.82-122); Solove (2002, pp.1116-1118).

or enjoys privacy is to say that, in some respect and to some extent, the person (or the person's mental state, or information about the person) is beyond the range of others' five senses and any devices that can enhance, reveal, trace, or record human conduct, thought, belief, or emotion.' (Allen 1988, p.15) On these accounts, 'access' is best understood as cognitive, rather than physical (Powers 1996).

While it is plausible that privacy involves access, such views are incomplete. First, a straightforward formula in which greater access corresponds to decreased privacy would fail to account for all the salient features of privacy. Consider informal agreements among neighbors: one might agree to mutual access to each others' yards to prevent break-ins or trespassing. In a neighborhood with few incidents of these crimes, such an agreement would mean that the agreeing parties have less privacy than before, but it's difficult to see why one should care about net privacy at all. What is important is that there are two impacts on privacy: the agreement, which decreases privacy, and the effect upon certain crimes (assume that it marginally decreases them), which increases privacy. One might consider many other facts to determine whether such a trade-off is worthwhile (e.g. voluntariness, knowing the agreeing neighbors) but net change alone doesn't seem to matter. This first problem is minor and can be remedied by being more specific about the parties involved.

A second problem concerns what precisely access consists of. Consider the vast amount of data gathered by large Internet retailers or credit card companies. When a person makes purchases online the retailer obtains certain information. Where a retailer has this data and uses it along with data from millions of other consumers to analyze P's buying habits, predict P's behaviors, and market more effectively to P, it seems clear that P has less privacy. However, just what is it that the retailer is accessing? Are predictions about future buying habits 'facts,' as required by Parent and Allen? Perhaps, but there should be some explanation as to why.

Please cite to published version: *Res Publica* 17(3): 275-290 (2011).

A third question concerns falsehoods. Exposure to falsehoods about a person doesn't lead to actual knowledge and cannot constitute access to facts about a person. Likewise, exposure to a falsehood about a person cannot constitute cognitive access to that person. Nonetheless, the dissemination of falsehoods would sometimes appear to decrease privacy. Suppose that a healthcare provider confuses medical records such that P's name is attached to the medical history of another. If the provider releases that record, it would seem that P's privacy has diminished. P has a legitimate complaint against the medical provider and that complaint is grounded in a diminution of her privacy.

One might argue instead that P has a legitimate complaint, but one based upon the fact that the medical provider has disseminated false or damaging information rather than the fact that P's privacy has diminished. That would be a mistake. Suppose that the information released reflects positively on P by omitting a condition that P wanted concealed. In that case, P's complaint would not be that the record released is false or damaging because P would prefer the release of a false record to the release of the true record, and the release of the false record actually benefits P. Rather, P's complaint is that the record was released, period—and that release diminishes her privacy.

Control Definitions

The clearest articulation of a control-based view is Charles Fried's:

> Privacy is not simply an absence of information about us in the minds of others; rather it is the *control* we have over information about ourselves.
>
> To refer for instance to the privacy of a lonely man on a desert island would be to engage in irony. The person who enjoys privacy is able to grant or deny access to others.(Fried 1984, pp.209-210, italics in original)[2]

[2] Similarly, Elizabeth Beardsley understands the 'conceptual core' of privacy as the right to selective disclosure, such that having privacy is having the ability to selectively disclose information Beardsley (1971). See also Moore (2010); Gross (1971); Van Den Haag (1971); Westin (1967, p.7).

An advantage of control views is that they track the moral importance of control of information. A number of commentators argue that we value privacy precisely because it allows us to disclose information selectively (Inness 1992, pp.41-55; Rachels 1975). However, from the fact that we value privacy because it allows selective information disclosure (e.g. for contract purposes), it doesn't follow that privacy *just is* the ability to selectively disclose. It may be that many important cases of restricted access also allow for disclosure at our discretion. Thus, we may come to value that access restriction.

Another advantage of control views is that they can explain privacy decrease in cases of fortuitous non-disclosure. The lovelorn fan of a movie star might wish for nothing more than the star to read his heartfelt letters, even though the letters go unopened. One might maintain that the letters cannot remain private insofar as someone else possesses them, despite no one accessing them (either physically or cognitively). Similarly, it is difficult to maintain that one has privacy regarding one's naked body on a nude beach, even if others fail to see it. The control view explains why: one lacks privacy because others noticing or not is beyond one's control. Surely, though, the fan retains some greater degree of privacy in the letters than he would if the star did read them, though his degree of control is the same in each case.

Notice also that there are cases in which people have no control over information, but have privacy with respect to that information nonetheless. For example, when a physician leaves an examination room to give a patient privacy to discuss treatment options with family, the patient neither exercises nor maintains control over access to the room. That privacy is at the will of the physician because she might not have left in the first place and might return at any moment. There are social controls in place, but there need not be. The physician might have left for other reasons, fortuitously conferring the family some privacy.

Please cite to published version: *Res Publica* 17(3): 275-290 (2011).

The Particularized Judgment Account

Three-part Relation

The remainder of the paper provides an account of informational privacy that avoids the problems with the views discussed above. The account is in two parts: first, privacy must be understood as a three-part relation between a person or persons (P), an object or domain of privacy (O), and some other person or persons (Q). 'P_{OQ}' denotes such a relation. To say 'P has privacy' will always be incomplete; however, we can clarify 'P has privacy' by first specifying an object or domain (O), (e.g. health, finances), and second, by specifying others (Q) who stand in some relation to P and O (e.g. everyone, the state). Thus, to say that health information is private is to say of a person (P), with respect to her health (O), that some other person(s) (Q) lack knowledge of or access to information regarding P and O, that P maintains control of O with respect to Q, or the like, depending on the substantive account of privacy.

This claim should be uncontroversial: it is compatible with the views of privacy discussed above and would not conflict with any substantive account of privacy. Nonetheless, it is crucial. Failing to specify domains and other persons in discussing privacy obscures ways in which privacy may diminish. For example, if P sends an email to Q, but mistakenly sends it to his entire address book as well, it is clear that P's privacy with respect to the email (O) has diminished with respect to everyone in his address book (Q). We might be tempted to say that the email is therefore not private, but that would be too quick. If some of the recipients forward the email or post it online, it would seem that P's privacy regarding the email with respect to the rest of the online world has diminished. Thus, failing to specify the entire P_{OQ} relation can obscure ways in which privacy may diminish.

Further, understanding privacy as a three-part relation allows us to assess the effects of privacy in otherwise ambiguous cases. Above I used the example of neighbors monitoring each

Please cite to published version: *Res Publica* 17(3): 275-290 (2011).

others' homes to prevent break-ins and noted that it is not particularly useful to ask about the program's effects on privacy per se. It is more useful to examine particular privacy relationships: P's privacy about his actions (O) with respect to his neighbors (Q) and P's privacy about the contents of his home or his actions (O) with respect to potential trespassers (Q). A related point is that analyzing particular P_{OQ} relations will allow for better analysis regarding the moral relevance of privacy. Whether privacy is valuable, or whether persons have claims to privacy, will likely turn on the P_{OQ} relation. Thus, whether one has a claim to privacy with respect to one's health will depend upon whether it is privacy with respect to one's doctor, spouse, neighbor, and so forth.

Particularized Judgments

The second part of the particularized judgment view is more substantive. The central claim is this: P's having informational privacy about O with respect to Q means that Q's ability to make reasonable particularized judgments about P regarding O is limited.

The first step in specifying the particularized judgment view is an explanation of what information matters for privacy. One requirement is that the information pertain to P. This seems obvious, but it is important nonetheless, and it merely falls out of our understanding privacy as a three-part relation. Any instance of P having privacy about O with respect to Q will only make sense in terms of information that pertains to P: if O is just some aspect of the world (e.g. astronomy), then there is no sense in which P has privacy about O with respect to Q, even if Q knows nothing of O, and there is no sense in which P has diminished privacy about O with respect to Q where Q learns about O.

It is not the case, however, that all types of information pertaining to P bear upon P's privacy. General biological facts may be applicable to P, and learning a biological fact may allow Q to make a reasonable judgment about P: learning that humans are vertebrates allows Q to make a reasonable judgment about P's physiology (O), namely that P has a spine. However, what matters is

the degree to which information differentiates P from others. While knowledge of a biological fact will not alone bear upon P's privacy, knowledge of some aspect of P's physiology that is not shared by others does. Learning that P has a Y chromosome differentiates P from other humans and allows Q to make a reasonable judgment about P with respect to gender (O), namely, that P is male. The judgment that P is male is not a very particularized judgment, and absent other information about P, Q's ability to make particularized judgments about P within other domains remains limited.

In other words, privacy admits of degree, and P may have greater or lesser privacy about O with respect to Q. One way in which privacy may be greater or lesser is the degree to which Q can differentiate P from other persons with respect to O. Compare Q's learning that P has a Y chromosome with Q's learning that P has a chromosomal abnormality such as Down's syndrome. Because having such an abnormality distinguishes P from others with respect to his genetic condition (O) to a greater degree than having a Y chromosome, Q's learning of the abnormality allows Q to reasonably make a more particularized judgment about P and O than Q's learning of P's Y chromosome.

In addition to being particularized based on the degree to which information differentiates P from others, judgments may be more or less particularized in terms of grain. Consider P's privacy about his finances with respect to Q. P might have privacy about his finances with respect to Q in that Q knows nothing about P's income and net worth. Alternatively, Q might be able to conclude from P's actions that P is quite wealthy, decreasing P's privacy about his finances with respect to Q. A further possibility is that Q has very detailed information about P's income, investments, and so forth. Thus, P_{OQ1} (where Q has no good information about P's finances) > P_{OQ2} (where Q only has information about P's spending habits) > P_{OQ3} (where Q is P's financial advisor and has detailed information about P's income and assets). This is a difference in privacy based on the grain of judgments, not the degree to which a judgment differentiates P from others. Compare the possibility

that Q learns that P is the one-thousandth richest person in the world. Such a fact differentiates P from everyone else with respect to financial status, but it doesn't permit Q to make as fine-grained judgment about P with respect to finances as when Q has detailed information about P's income, assets, and so forth.

Reasonableness: Judgments, Falsehoods

The next issue concerns what quality of information matters in assessing privacy. Factual data about a person is potentially private, but is the revelation of facts the only thing that matters with respect to privacy? Above I stated that falsehoods could decrease privacy and that the failure to account for falsehoods is a weakness of access accounts. Here I address further how judgments and falsehoods should factor into privacy. Whether they diminish privacy turns on the extent to which such information allows one to make reasonable particularized judgments about P regarding O.

Consider first others' judgments of P—'judgments' being propositional beliefs, supported by evidence (however spurious), that require some degree of inference. The distinction I have in mind differentiates (1) between information about a series of symptoms and information regarding whether one has a syndrome and (2) between data and conclusions supported by that data.[3] Whether P experiences fatigue, sleeplessness, nausea, and so forth is clearly information that bears upon P's privacy. By learning that P experiences these things, Q may reasonably make more particularized judgments about P regarding health. Now suppose that Q learns that P's symptoms indicate that P has chronic fatigue syndrome (CFS). The question is whether P's privacy regarding health with respect to Q decreases between time t_0, when Q knows of P's symptoms, and time t_1, when Q learns that those symptoms indicate that P has CFS.

[3] Note *any* such belief will require inference, if only about the reliability of our belief-forming mechanisms; the difference I refer to here is one of degree. Nonetheless, some beliefs require more inference than others, and here I wish to pick out those beliefs that require a substantial degree of inference. The point I wish to make is just that privacy may continue to decrease as judgments accrue; the whole picture of P's privacy is not complete once observations are made.

Please cite to published version: *Res Publica* 17(3): 275-290 (2011).

On this account, P's privacy about her health with respect to Q decreases. Q's learning of the relation between P's symptoms and CFS increases Q's ability to make reasonable, particularized judgments about P with respect to P's health. One such judgment is that P has CFS. In turn, the judgment that P has CFS allows Q to make further reasonable particularized judgments about P's health. For example, Q can make reasonable judgments about P's prognosis. Likewise, where P's physician makes the judgment that P has CFS, Q's learning of that diagnosis decreases P's privacy about his health with respect to Q. Regardless of whether Q thinks P has CFS because she inferred it based on her understanding of disease or whether she thinks it based on her trust in a doctor's abilities, the proposition that P has CFS allows Q to reasonably make more particularized judgments about P's health.

Suppose, however, that the judgment that P has CFS is incorrect, either because P's physician misdiagnosed P, or because it turns out that CFS does not really exist (perhaps instead it is an amalgam of other conditions). In that case, the proposition 'P has CFS' would be false. When Q hears that P has CFS, does P's privacy about his health with respect to Q diminish? As noted above, access views preclude falsehoods from bearing upon privacy. Accordingly, one might argue that false information doesn't diminish privacy. Such an objection is unsustainable.[4]

First, that view would require us to revise judgments of privacy loss over time. Suppose that a doctor diagnoses P with some condition, that the diagnosis is consistent with current medical understanding, and that P's medical record is released shortly after diagnosis. It would appear at the

[4] One could argue that falsehoods cannot decrease privacy because only true propositions constitute information. See Floridi (2004, p.197). Q's hearing a falsehood about P imparts no information, and hence P's informational privacy cannot decrease.But information need not be true. James Fetzer maintains that for a proposition to constitute information, it need only be well-formed and meaningful Fetzer (2004, pp.224-225). Even if it were correct that information must be true, we would still need a concept to apply to well-formed and meaningful propositions that are either untrue or whose truth is unknown (Fetzer suggests 'nformation,'), and we would have to investigate whether such propositions are relevant to privacy. Concluding that privacy depends only upon truths on the grounds that information is by definition true would assume the answer to the question at hand—namely, whether falsehoods are relevant to personal privacy.

Please cite to published version: *Res Publica* 17(3): 275-290 (2011).

time of the disclosure that P's privacy about her medical condition has diminished with respect to the public. What happens, however, if decades later a more-sophisticated understanding reveals that the condition did not really exist? The view that only factual information decreases privacy would entail that P's privacy had not actually been diminished and that P could have no claim to privacy regarding the record. We would have to revise the conclusion that releasing a medical record constituted a privacy loss for P. The better account is that release of the record diminishes P's privacy at the time of disclosure, regardless of evolving medical science.

Additionally, disclosure of true propositions can decrease privacy to the same degree as the revelation of false propositions. Consider the following: P, a juvenile, confesses to an assault that she did not commit. P's social-worker, S, informs two people about P. First, S tells Q that P has committed the assault; second, S tells R that P confessed to committing the assault. Q has learned something false, but R has learned something true. Certainly P's privacy about her criminality diminishes with respect to R. The contention that privacy cannot diminish where a person has received false information, however, would require us to conclude that P's privacy has diminished *only* with respect to R. That's implausible. P's privacy about her criminality has diminished with respect to both Q and R. The particularized judgment account explains why: hearing that P committed an assault allows Q to make a reasonable particularized judgment about P regarding criminality, just as hearing that P confessed to committing an assault allows R to reasonably make a particularized judgment about P regarding criminality.

Not just any falsehood diminishes privacy on the particularized judgment account; the degree to which privacy is diminished turns on the reasonableness of the inferences that one may draw, not on the truth or falsity of what a person learns. So, Q may learn something false (P has CFS) and may make reasonable particularized judgments about P regarding health on that basis. But there are, of course, falsehoods that Q could hear that would not allow Q to make reasonable

particularized judgments. Suppose that P's doctor has diagnosed P with 'the vapors.' If Q hears that P has the vapors, Q's ability to make reasonable particularized judgments about P's health remains unchanged, given the vagueness and dubiousness of the condition. Note, too, that even if Q had learned something true, viz., that the doctor made the diagnosis, Q's ability to make reasonable particularized judgments about P's health would also have remained unchanged. At best, Q could make a reasonable particularized judgment about P's doctor regarding competence. It is not a proposition's truth or falsity that makes it relevant for privacy, but the reasonable particularized judgments that it allows others to make.

What constitutes a reasonable judgment in this context? A judgment is reasonable to the extent that the information underwriting the judgment constitutes evidence of the judgment relative to a set of background beliefs that is itself reasonable.[5] In other words, a reasonable judgment need not be true, for an inference could be based upon misinformation or a mistaken set of background beliefs, and reasonable given the misinformation or background beliefs. Thus, reasonable judgment is not the same as knowledge. Given a mistaken set of background beliefs, some proposition A might allow one to make a reasonable inference that B is true, whereas in fact, A's being true doesn't make it more likely that B is true.

Suppose that in the mid-19th century, P's doctor measures P's skull and concludes that P has a 'primitive' forehead. According to medical science of the time, this would indicate that P is mentally deficient and has criminal tendencies.[6] Q pilfers the record stating that P has a primitive forehead, allowing Q to make a judgment about P with respect to his mental faculties. Of course, Q would know nothing about P's mental faculties. The judgment is reasonable in light of the

[5] See Christensen (1997) for a discussion of the conceptual difficulties regarding what it means for evidence to confirm a hypothesis relative to a set of background propositions.

[6] See Pick (1989, pp.109-135); Young (1970, pp.11-14). Pick and Young describe attempts of phrenologists seeking to establish connections between features of persons' skulls and their mental faculties or traits.

Please cite to published version: *Res Publica* 17(3): 275-290 (2011).

background beliefs of the time even though it is untrue. Thus, P's privacy regarding his mental faculties decreased with respect to Q based on reasonable judgments, not on Q's knowledge or upon Q's inferring something true.

Because reasonable judgments are limited to those in which information constitutes evidence supporting those judgments in light of reasonable background beliefs, privacy doesn't decrease just because information may cause some Q or other to draw outlandish conclusions about P regarding O. So, where Q sees a picture of P (who appears human) and concludes that P is from Mars, it doesn't follow that the picture decreased P's privacy about his birth planet with respect to Q. Similarly, we can differentiate privacy decreases associated with different types of inference. Consider Q's making an educated guess about some fact regarding P based on weak evidence; that entails some degree of privacy loss, but such an inference is less reasonable than one based on stronger evidence and rigorous inquiry.

Questions remain about what constitutes a reasonable set of background beliefs and when an inference is reasonable. There is vagueness, but that's not problematic. Reasonableness is a matter of degree. The more widely held background beliefs are and the greater degree to which information supports an inference relative to those background beliefs, the more reasonable those inferences would seem to be. In turn, the more reasonable Q's particularized judgments about P and O are, the more P's privacy decreases. Thus, revealing the measurements of P's skull would decrease P's privacy more when people generally believed that P's dimensions indicate diminished mental faculties than it would when people generally believe that P's measurements indicate nothing whatsoever about P's abilities. Still, a particularly knowledgeable Q might lack the mistaken set of background beliefs regarding skull measurements, and Q's learning of P's dimensions would not decrease P's privacy regarding his faculties with respect to Q.

Please cite to published version: *Res Publica* 17(3): 275-290 (2011).

Moreover, the vagueness can be accounted for by being mindful of the Qs involved. Suppose that P is a high-school student who has a habit that her classmates take to be indicative of being gay (by way of adolescent misconception), but which in fact does not so-indicate. The particularized judgment account recognizes that P has less privacy about her sexual orientation (O) with respect to her classmates (Q) after they notice the habit than before. It seems wrong to say that P has the same degree of privacy about her sexual orientation with respect to Q when Q believes that P is gay as when Q has no beliefs on the matter at all. However, with respect to the rest of the population, which doesn't have the same misconception, P's privacy about her sexual orientation would not decrease.

Other Factors: Likelihood, Potential, and Anonymity

Having addressed the type of information that is relevant in addressing privacy loss, in this section I address factors relevant in determining when one's ability to make reasonable particularized judgments increases. The first of these is likelihood. As Q gathers information about P and O, her ability to make particularized judgments generally increases. Put in terms of the account at work here, the limits of Q's ability to make particularized judgments about P and O are in part a function of background likelihoods. Some information would lead Q to make a certain inference if she would only look at it a certain way or if she had certain skills, perceptiveness, or knowledge. For example, P's constricted pupils, heavy eyelids, and languor could support Q's inference that P has indulged his heroin habit. But if Q knows nothing of P's habit or about signs of heroin use or if Q just doesn't expect P to be a user, then Q will likely not make the inference. Nonetheless, P's privacy about his use with respect to Q decreases as Q observes P's dopey actions, but it decreases P's privacy less than it would decrease should R (who has a well-trained eye for signs of heroin use) observe P's actions. Likewise, R's ability to make reasonable, particularized judgments about P regarding drug use is greater than S's, who has the same background knowledge, but fails to make the inference

because she has more important things to think about. The reason is simply that R is more likely to make the inferences.

Another relevant factor is potential to garner information. From the fact that Q has not been exposed to information about P regarding O, it doesn't follow that P has privacy about O with respect to Q. Suppose that P and Q share a computer with a browser that logs all websites visited and that P does not modify that log. Q, who is computer savvy, could easily trace all of P's Web browsing, but doesn't. The fact that Q could trace P's browsing seems to decrease P's privacy. That is, P has less privacy with respect to his browsing (O) where Q could easily open a log of that browsing. P could argue that he needs software to clean his history on the grounds that he wants his privacy. That need would not be entirely mitigated by Q's guarantee that he would not trace P's browsing habits. An advocate of a control-based view might argue that P's lack of control over the information is the reason P has less privacy, but this need not be the case. Even if we suppose that P could simply erase the browser's log (but doesn't), it would seem that P's privacy still decreases.

Further, the potential to access information is distinct from the likelihood of accessing information. Suppose that Q and R can each open P's browsing records. Q, however, is utterly uninterested doing so, while R is mildly curious. R is therefore more likely than Q to look at the records, and it seems that P's privacy about his habits is greater with respect to Q than with respect to R.

The focus on likelihoods and potential inferences addresses some of the advantages of control views. One of these, noted above, is that control views can explain why one has decreased privacy even in cases of fortuitous non-disclosure. Hence, where P posts his thoughts on his blog for all to see, his privacy regarding them decreases even if his thoughts are not interesting enough for anyone to read. This account explains the decrease in privacy not as lack of control, but as greater likelihood of others reading.

Please cite to published version: *Res Publica* 17(3): 275-290 (2011).

So, P's privacy is a function of both Q's not actually making particularized judgments and the likelihood (or potential) of Q's doing so. This is important in addressing sophisticated analysis of pre-existing data. Consider Internet retailers and social networks. These have a great deal of user information, but at early stages their ability to analyze the information is limited. They have fewer users, less sophisticated relational databases, and less advanced analytic tools than they do later. But as such entities invest more and create larger datasets and better databases, their ability to make very particularized judgments about users increases. If P shares some information early on, the site has for the moment a limited ability to make particularized judgments. Even if P stops using the site, it will be able to make better inferences about P in a variety of domains (e.g. his tastes, political views, etc.) based on greater information about other people and better analytic tools. Because of this potential, P's privacy decreases in those domains when he initially shares the information. That is, one factor in the retailer's ability to make particularized judgments about P is possessing some basic data about P's purchases. The limits on the retailer's ability to make particularized judgments decreases at the point P shares his data (and decreases further when the retailer develops better tools).

Likelihood and potential to make reasonable, particularized judgments are a function of both information and Qs' capacity to make use of information. This is a key feature to this account. P might willingly disclose information to Q, relinquishing privacy regarding that information. But to the extent that Q is incapable of drawing inferences P retains some degree of privacy. This helps us understand how disclosure of identical information to two different entities could result in different degree of privacy loss. In case one, P discloses information about her tastes to a small online retailer; in case two, P disclose the same information to a larger online entity that compiles data from a wide variety of sources. In each case, P's disclosure is the same, and each Q's access to that information is the same. Yet in case two, P's privacy decreases more, solely in virtue of the larger

entity's great ability to analyze that information and make more reasonable, particularized judgments about P.

With this discussion in mind, we can address cases of potential privacy loss where information is not readily identifiable with a particular person. Consider the release of medical records with names and addresses redacted and data mining of electronic communications in which all communications are analyzed, but only those fitting certain criteria are individuated. One might argue that such cases do not involve loss of privacy as some have suggested.[7]

On the view offered here, increased likelihood of, or potential for, particularized judgments helps explain why these cases do involve diminished privacy. To begin, there is a surprisingly small amount of information that is not identifiable. Information that is not attributable to an identifiable person when viewed in isolation is often identifiable when data sets are linked.[8] The potential for re-identification of purportedly anonymized information is large and far greater than supposed a short time ago (Ohm 2009). Collection of anonymous data about P at time t_0 removes an important limit on some Q making particularized inferences about P (and identifiable as P) at some later time t_1. So, the potential for re-identifying information helps explain how collection of anonymous or de-identified information can diminish a person's privacy on the particularized judgment account.

However, the potential for re-identification is not the only way in which information not readily identifiable with a specific person can decrease privacy. Anonymous information can

[7] Richard Posner argues that no privacy violation occurs in data mining until a person has actually viewed information collected. Posner (2006, pp.96-97). Recently Matthew Tokson has argued that no privacy loss occurs if information is sorted electronically. Tokson (2011).

[8] Computer scientist LaTanya Sweeney was able to use summaries of hospital visits with explicit identifying information redacted in combination with voting records to identify patients, including the governor of Massachusetts. Analyzing 1990 census data, Sweeney determined that three pieces of information (postal code, birth date, and sex) uniquely identify 87 percent of people in the United States; city, birth date, and sex can uniquely identify 53 percent; county, birth date, and sex can uniquely identify 18 percent. Sweeney (2000). In a widely publicized case, America Online released data regarding user searches, with 'personal' information redacted and unique numbers assigned to each. Because the queries often related to users' lives, they were far from anonymous, and *New York Times* reporters tracked one such user down for an interview. Barbaro & Zeller Jr. (2006). See also Narayanan & Shmatikov (2008).

decrease a person's privacy independently of its potential to be re-identified. Suppose that someone steals P's journal, in which P recounts his ennui, and posts the pages online. The journal contains nothing that would allow one to recognize P. Surely this diminishes P's privacy. The information contained in it is about P, and it allows others to make reasonable particularized judgments about P regarding his ennui. The readers of the journal do not believe that P has ennui, merely that the writer of the journal has ennui. Nonetheless, the inference is about P. Likewise, if someone posts a nude picture of P online, with P's face obscured, P's privacy regarding his body decreases when Q views the picture. Q's judgments in response to the picture are about P, even if Q doesn't know that the picture is of P.[9] Of course, the weight of a privacy violation may at times turn on whether the relevant Qs learn that the information is about, or the picture is of, P.

Access and Control Accounts

It is useful to consider how the particularized judgment account is distinct from access and control views. To begin, it shares with access accounts that one's privacy in some domain is independent of whether one can actually exercise control over information in that domain. So, the fact that I am oblivious to and, hence, unable to control who learns about my deepest subconscious motivations, does not entail that I have no privacy regarding those motivations. The account is also distinct from control accounts insofar as it allows that one can retain privacy fortuitously, as when an observer has enough information about P to make an inference, but fails to as a matter of sheer luck. P's clothes, mannerisms, and speech might be overwhelming evidence of his background, but if Q is too distracted to realize it, P retains some privacy regarding his background with respect to Q, despite having no control over Q's thinking.

[9] Put another way, 'making a particularized judgment about' is referentially transparent. We can substitute 'P' and 'the author of the journal' without altering the truth value of 'Q made a particularized judgment about P, viz., that P has ennui.' To use a familiar example, Q might make a particularized judgment about Cicero, viz., that he is a gifted writer. Q has likewise made such a judgment about Tully, even though Q does not know Tully and Cicero are the same person. Of course, Q does not *believe that* Tully is a gifted writer.

The particularized judgment account is distinct from access views by accommodating the potential for falsehoods to diminish privacy. Moreover, it offers an explanation of how increasingly sophisticated ability to analyze data can diminish privacy, even where an individual provides access to no new data, as the cases of online retailers and social networks demonstrates. In contrast with access views, the particularized judgment account shares with control views a sensitivity to the ways in which potential inferences or exposure diminish privacy. Unlike control views, it focuses on how removing limitations on others' abilities to make particularized inferences about a person includes removing impediments such that the likelihood of such inferences increases.

Implications

Although this account is descriptive, it has several implications for how we understand privacy's value and the conditions under which we have claims to privacy. One of the most important is that because privacy decrease is multifaceted, it would appear unlikely that we can provide a single account of privacy's value. Various commentators have tried to articulate a single value for privacy, including privacy as instrumentally valuable, as a necessary condition for moral autonomy, as constitutive part of respect for persons, and as a political value (Rubel 2007, pp.925-935). However, on the account offered here, privacy can decrease in myriad ways, and the degree to which one has privacy is a function of many factors. There is little reason, therefore, to think that privacy has a single type of value. Privacy regarding one's voting habits with respect to state actors may be an important political value, whereas privacy regarding one's shopping habits with respect to marketers may be instrumentally valuable. Indeed, many instances of privacy loss are likely of no moral concern.

Another implication of the account is that, contrary to much current regulation, privacy issues do not end when a person shares information with another person or entity. Because an initial disclosure may be compared with other data sets and subjected to more sophisticated analysis,

privacy diminution may occur that goes far beyond that initial disclosure and far beyond what one could reasonably expect when sharing information. Such further decreases are not fully explained by access or control views. If privacy is indeed valuable, we will have to account for decreases that go beyond initial disclosures. Moreover, if we think that we can relinquish privacy claims by agreements (e.g. agreements with entities collecting information about Internet use), we will have to consider whether agreeing to use of information extends to any inferences from such information.

This account makes general claims about privacy problematic and provides tools for more nuanced explanations of the privacy implications of technologies and practices. Consider novel surveillance technologies such as traffic cameras that detect and automatically issue citations for driving violations. Drivers whose actions are captured by traffic cameras have diminished privacy regarding their driving practices with respect to the government, but to the extent that cameras displace police stops, privacy may increase in morally salient respects. Much information gathered in traffic stops is far beyond what is gathered by cameras: the cleanliness and contents of one's car, how one reacts under the stress of a police encounter, and so forth. Depending on the person looking, such information can allow for fairly particularized judgments about the driver. These are potentially of greater importance than the information collected by traffic cameras, and might give rise to privacy claims where detection of driving violations might not.

Objections

I've defended the view that P's having informational privacy about O with respect to Q means that Q's ability to make reasonable particularized judgments about P regarding O is limited. I have advanced an expansive understanding of the type of information that bears upon privacy and of the ways in which one's ability to make reasonable, particularized judgments increases. But if privacy may decrease in so many different ways, and so many different things count in determining whether P has privacy, one might object that almost *anything* can bear upon whether P has privacy.

Please cite to published version: *Res Publica* 17(3): 275-290 (2011).

The response to the objection requires keeping several things in mind. One is the particular domain of privacy in question. P's walking down the street where Q can observe decreases P's privacy about his whereabouts with respect to Q. It is hard to see how it could be otherwise. The other possibilities are that P *does* have privacy about his location with respect to Q or that P's location is not the sort of thing about which P could have privacy in the first place. Neither of these is plausible. Thus, even utterly ordinary occurrences do seem to bear upon privacy. Another factor to remember is that decreases in privacy may be very small. Merely being observed walking down the street may decrease one's privacy, but it doesn't decrease it very much. Moreover, observing someone walking along need not matter much. On my view, privacy loss happens persistently, but only some such losses matter morally. P's loss of privacy about his whereabouts with respect to Q as he strolls down the street is a loss that would seem not to matter morally.

It is unsurprising that so many things can bear upon privacy. As noted, even highly generalized propositions may count with respect to one's privacy. Suppose a well-designed study shows that all movie-watchers have fantasies about being movie stars. This fact allows Q to infer that P fantasizes about being a movie star and, as a result, it would seem that P has less privacy about his fantasies (O) with respect to Q. Note, however, that this works in the other direction as well. Consider the research of Alfred Kinsey, which showed that there is enormous variation in people's sex lives. It would seem that this research would make it harder for people to make reasonable particularized inferences about others' sex lives. Suppose that, just before Kinsey's research was published, Q believed that most all people in long-term relationships (including his neighbor, P) had sex in a certain way, with a certain frequency. Suppose, however, that after Kinsey's research was published, Q came to believe that people in long-term relationships had enormous variation in their sexual habits (including P). It seems that the research actually *increases*

P's privacy, for enormous variation makes it harder to make reasonable particularized judgments about others' sex lives.

Another objection is that the particularized judgment account is really just a species of access account. Q's ability to make reasonable particularized judgments about P regarding O is an ability to access information about P. The first response to this objection is that the view offered is not about access insofar as falsehoods cannot provide access to P—access views focus on access to facts about P or cognitive access to P, neither of which occurs upon the dissemination of falsehoods.

There's a second response: the claim that falsehoods may decrease a person's privacy is controversial, and some will reject it. Rejecting the claim regarding falsehoods, however, doesn't undermine an analysis of privacy in terms of reasonable particularized judgments. One could maintain that P's privacy about O with respect to Q decreases when Q's ability to make reasonable, particularized, and truthful judgments about P and O increases. Though this appears closer to an access view, it is distinct in that it incorporates P's ability into the analysis. Two different observers (e.g. a sophisticated social network and a small retailer) can have access to the same data about P, but P would have less privacy with respect to the one with more sophisticated capabilities. Finally, even to the extent that this account relies on access, it offers an advantage in specifying the conditions under which access to P increases and focusing on inferences as the relevant factor in assessing privacy.

Conclusion

My task has been to provide a descriptive account of informational privacy. I've argued that the predominant accounts—focusing on access and control—are problematic. Instead, I've offered an account that does two things. First, it demands that privacy be understood as a three-part relation among a person or persons (P), a domain or object of privacy (O), and another person or persons (Q). Second, it focuses on inferences as the salient aspect of privacy: P has privacy about O with

respect to Q when Q's ability to make reasonable particularized judgments about P and O is limited. Questions remain, including questions about privacy's value. The particularized judgment view suggests that there is no single value that might give rise to privacy claims and that, in any case, decreases are ubiquitous. Nonetheless, the account herein provides better tools for analyzing privacy protections, diminutions, and claims.

Acknowledgments

I very much appreciate the many helpful comments I've received from Claudia Card, Robert Streiffer, Russ Shafer Landau, Harry Brighouse, Victoria Nourse, Fred Harrington, Madison Powers, and Tom Beauchamp; audiences at the University of Wisconsin, University at Albany, and Georgetown University; and the paper's anonymous referees.

Works Cited

Allen, Anita. 1988. *Uneasy access: privacy for women in a free society*. Totowa, N.J.: Rowman & Littlefield.

Barbaro, Michael and Tom Zeller Jr. August 9, 2006. A face is exposed for AOL searcher no. 4417749. *New York Times*.

Beardsley, Elizabeth. 1971. Privacy, autonomy, and selective disclosure. In *NOMOS XIII: privacy*, eds. J. Roland Pennock & John W. Chapman, 65-70. New York: Atherton Press.

Christensen, David. 1997. What is relative confirmation? *Noûs* 31: 370-384.

DeCew, Judith. 1997. *In pursuit of privacy: law, ethics, and the rise of technology*, Ithaca, N.Y.: Cornell University Press.

Van Den Haag, Ernest. 1971. On privacy. In *NOMOS XIII: privacy*, eds. J. Roland Pennock & John W. Chapman, 149-168. New York: Atherton Press.

Fetzer, James. 2004. Information: does it have to be true? *Minds and Machines* 14: 223-229.

Floridi, Luciano. 2004. Outline of a theory of strongly semantic information. *Minds and Machines* 14: 197-221.

Fried, Charles. 1984. Privacy [a moral analysis]. In *Philosophical dimensions of privacy*, ed. Ferdinand Schoeman, 203-222. Cambridge: Cambridge University Press.

Please cite to published version: *Res Publica* 17(3): 275-290 (2011).

Gavison, Ruth. 1984. Privacy and the limits of law. In *Philosophical dimensions of privacy*, ed. Ferdinand Schoeman, 346-402.

Gross, Hyman. 1971. Privacy and autonomy. In *NOMOS XIII: privacy*, eds. J. Roland Pennock & John W. Chapman, 169-181. New York: Atherton Press.

Inness, Julie. 1992. *Privacy, intimacy, and isolation*. New York: Oxford University Press.

Moore, Adam. 2010. *Privacy rights: moral and legal foundations*. University Park, Pa.: Pennsylvania State University Press.

Narayanan, Arvind and Vitaly Shmatikov. 2008. Robust de-anonymization of large sparse datasets. *Proceedings of the 2008 IEEE Symposium on Security and Privacy*: 111-122.

Ohm, Paul. 2009. Broken promises of privacy: responding to the surprising failure of anonymization. *UCLA Law Review* 57: 1701-1777.

Parent, W.A. 1983. Privacy, morality, and the law. *Philosophy and Public Affairs* 12: 269-288.

Pick, Daniel. 1989. *Faces of degeneration: a European disorder, c.1848-c.1918*. Cambridge: Cambridge University Press.

Posner, Richard . 2006. *Not a suicide pact: the constitution in a time of national emergency*. Oxford: Oxford University Press.

Powers, Madison. 1996. A cognitive access definition of privacy. *Law and Philosophy*, 15: 369-386.

Rachels, James. 1975. Why privacy is important. *Philosophy and Public Affairs* 2: 323-333.

Rubel, Alan. 2007. Claims to privacy and the distributed value view. *San Diego Law Review* 44: 921-956.

Solove, Daniel. 2002. Conceptualizing Privacy. *California Law Review* 90: 1087-1155.

Sweeney, Latanya. 2000. Uniqueness of simple demographics in the U.S. population. *Laboratory for International Data Privacy*, Working Paper LIDAP-WP4.

Tokson, Matthew. 2011. Automation and the fourth amendment. *Iowa Law Review* 96: 581-647.

Westin, Alan. 1967. *Privacy and freedom*. New York: Atheneum.

Young, Robert. 1970. *Mind, brain and adaptation in the nineteenth century: cerebral localization and its biological context from Gall to Ferrier*. Oxford: Clarendon Press.

Please cite to published version: *Res Publica* 17(3): 275-290 (2011).