Episteme, 10, 2 (2013) 179–192 © Cambridge University Press
doi:10.1017/epi.2013.15

EPISTEMOLOGICAL DIMENSIONS OF INFORMATIONAL PRIVACY

KLEMENS KAPPEL
kappel@hum.ku.dk

ABSTRACT

It seems obvious that informational privacy has an epistemological component; privacy or lack of privacy concerns certain kinds of epistemic relations between a cogniser and sensitive pieces of information. One striking feature of the fairly substantial philosophical literature on informational privacy is that the nature of this epistemological component of privacy is only sparsely discussed. The main aim of this paper is to shed some light on the epistemological component of informational privacy.

1. INTRODUCTION

As is familiar, there are different types privacy concerns (Tavani 2007). One type relates to physical access to some restricted zone, say to one's body or immediate physical vicinity, or one's home, car or private garden. A second privacy worry concerns a mental, social or legal sphere within which one should be free to make one's own decisions and deliberations, say within one's own life, or within the social institution of marriage or family. Yet a third privacy concern has to do with sensitive information, typically information about one's personal matters, such as one's medical record, religious beliefs or sexual preferences.

It should be clear that these three kinds of privacy concerns are rather different. If some stranger stands too close to me, he may be said to have violated my privacy in some sense. But this is because his mere proximity will reduce the set of actions I can perform, and because his presence might interfere with my mental life by taking up mental space, forcing me to think about his presence. Perhaps, by standing very close to me, he also gets access to certain bits of information about me that he couldn't easily have gotten in other ways, say about the character of my bodily odour. This may be a problem of informational privacy. However, the main problem is more likely that I now get an unsolicited sense of what he smells like. Invasions of physical space may lead to invasions of informational privacy, but clearly they need not. Conversely, one can access sensitive information about an individual without physically approaching this individual. And clearly, one can invade privacy by restricting autonomy, and do so without accessing personal information, say by passing various forms of legislation. Conversely, one can access sensitive information without limiting autonomy or freedom.

We can separate different, more specific concerns about information. One has to do with epistemic access to sensitive information. But other important issues regard the distribution of such information, and the editorial rights with respect to this information, or

https://doi.org/10.1017/epi.2013.15 Published online by Cambridge University Press

the representation of the information. Who holds the rights to delete or edit sensitive information about one held in various files or archives? While these questions are important, in this paper, I will only be concerned with questions regarding *epistemic access* to sensitive information, and I will refer to this as *informational privacy*.

It seems obvious that informational privacy has an epistemological component that consists of certain kinds of epistemic relations between a cogniser and some sensitive piece of information. One striking feature of the fairly substantial philosophical literature on informational privacy is that the nature of this epistemological component of privacy is only sparsely discussed. The main aim of this paper is to try to shed some light on the epistemological component of informational privacy. In the field of epistemology, we develop these complicated theories and fine-grained distinctions regarding matters epistemic. What, if anything, can this theorising tell us about the nature of informational privacy, or the nature of infringements of informational privacy, or our moral and legal concerns about maintaining informational privacy? There are other issues about privacy that should interest epistemologists, especially issues about the wider socio-epistemic consequences of policies and institutions upholding various privacy regimes. While these are interesting questions, I will not have anything to say about them.

The outline of the paper is as follows. It turns out that what we should say about the epistemological component depends in various ways of other non-epistemological features of a theory of privacy. So, in section 2, I propose what I call a generic epistemic theory of informational privacy. In section 3, I turn to the main question concerning the epistemological dimensions of privacy, what sort of epistemic access matters for privacy. Finally, in section 4 I offer a summary and some concluding remarks.

2. A GENERIC EPISTEMIC THEORY OF INFORMATIONAL PRIVACY

I want to start by simply stating what I will call a generic epistemic theory of informational privacy. First, I want to suggest that we benefit from disambiguating talk about privacy intrusions, violations of privacy and privacy invasions and similar expressions. These expressions are typically ambiguous between what I will call *privacy diminishments*, *privacy wrongs* and *privacy harms*.

Let me discuss each of these in turn, starting with the notion of a privacy diminishment. The basic idea is quite simple: If (but only if) Adam gets epistemic access to some facts regarding Bill that are sensitive given Bill's circumstances, then Bill's privacy is diminished with respect to Adam. We can generalise this as follows:

Privacy diminishment
For an individual A, a domain of facts P that is sensitive given A's context C: A's privacy regarding facts belonging to P is diminished with respect to B just in case B gets the right kind of epistemic access to these facts.

Note a couple of things. First, privacy diminishments are assumed to range over facts, or true propositions. So, a privacy diminishment is in that sense a *factive* notion. Second, privacy diminishments are a matter of degree. One has more or less privacy depending on the range of facts in P that has been accessed, and one's degree of privacy with respect to some domain of facts is relative to specific other agents. One can increase privacy with

https://doi.org/10.1017/epi.2013.15 Published online by Cambridge University Press

respect to one agent, while decreasing it with respect to another. Third, the notion of privacy diminishment depends on a prior notion of facts or truths belonging to a protected zone, or a domain as sensitive facts, as I say. Here and throughout I use 'sensitive facts' as a technical term that merely functions as a name for some class of facts that is in desert of some sort of privacy protection. It is not meant to be implied that individuals are particularly vulnerable to these facts being known by others, or that sensitive facts concern particularly delicate parts of someone's life.

Fourth, I want suggest that the boundaries of the protected zone are contextually dependent in quite subtle ways. So, whether some fact counts as sensitive depends on context. If I am at home, facts about the colour of my hair may belong to the protected zone. This is not so if I ride on a public bus. In that case it is not that I have *waived* some standing right that others don't get to see the colour of my hair. Rather, when I am in public space, such facts are not in the protected zone at all. Other types of facts remain protected, even when I am in public space, say facts about my religious denomination or sexual orientation.

Consider a woman currently undergoing chemotherapy. Suppose she wants to hide from public view what she takes to be the not very aesthetically pleasing remains of what used to be her beautiful hair. She therefore uses a scarf that completely shields her hair from public view. Her context is different from mine. Facts about the visual impression of her hair resulting from chemotherapy are, I want to suggest, part of the protected zone, even when she is riding the bus. These facts are sensitive in the technical sense I have in mind. It is not that I and everyone else on the bus have waived our rights not to have our hair seen, but this woman has not. Rather, though we are all on the same bus, our contexts are different. So, suppose I use my special camera to inspect her hair. This would be an invasion of her privacy; it would lead to a diminishment of her privacy, and would also involve a privacy wrong, a notion to be discussed in a moment. The woman could, of course, decide to go on the bus without her scarf; in that case, she would be granting others a sort of permission to see her hair.

Fifth, the notion of a privacy diminishment is *path-independent*. If I get to know about some sensitive fact, then privacy with respect to that fact is diminished; it doesn't matter how this knowledge was obtained. More about this in section 3.

Consider now what I will call a *privacy wrong*, the sort of moral wrong that pertains to typical cases of wrongful privacy diminishments. Again, the basic idea is simple. If Bill gets epistemic access to facts about Adam that are sensitive given Adam's context, and if this access constitutes a violation of the relevant norms, then this is a privacy wrong. Unless, of course, Adam or – someone else with the right kind of authority to do so – has granted Bill the permission to get access to these facts in these particular ways. Generalising this a bit, we get the following:

Privacy wrong
For an individual A, a domain of facts P that are sensitive given A's context C: B commits a privacy wrong if B gets a form of epistemic access to facts in P that are conventionally regarded as wrongful, unless B has been granted this form of epistemic access by someone in authority to do so.

As is already clear, this separates privacy diminishments from privacy wrongs. So, one's privacy might be diminished, and yet there is no privacy wrong, because the epistemic

https://doi.org/10.1017/epi.2013.15 Published online by Cambridge University Press

access in question is not illegitimate, or because someone with the right kind of authority has permitted such access.

A few comments are in order. First, privacy wrongs are defined over forms of epistemic access. I turn to this crucial question in section 4. Second, it is assumed that someone may be in authority to grant access to sensitive facts. Usually, of course, I have the authority to permit access to facts about me. If I intentionally reveal sensitive facts by granting explicit permission to access them, or by voluntarily putting those facts at display in public places, and others thereby get to know these facts, then no privacy wrong has been committed.

But the individuals that sensitive facts concern need not have the authority to grant access to them. Under Danish law, for example, information about one's medical record is sensitive in the sense that employers do not have legal access to it for the purpose of making hiring decisions. Now, individual employees may have a strong interest in disclosing their medical information – if it testifies to one's impeccable health, of course. But in a hiring situation, one cannot grant a prospective employer the legal right to study one's medical record. Similarly, in a job interview, a female job applicant cannot be asked to reveal information about whether she is pregnant or not, or whether she plans to have children in the future. If she nonetheless reveals this information, it would be illegal for the employer to use it in a hiring decision. The moral idea underlying these provisions is that one counteracts certain forms of inequalities by restricting access to certain sorts of information – then one might think that individuals do not even have the moral authority to permit this access to prospective employers.[1]

Finally, it is convenient to have the notion of a privacy harm, which I will stipulate as follows:

Privacy harm
The harm that one may suffer as a result of diminishments of one's privacy, or as a result of a privacy wrong being perpetrated towards one.

Note the causal constraint that goes into the definition of a privacy harm. Assume that we have an independent notion of harm according to which a harm suffered by someone involves a diminishment of well-being, or a setback of interest in some way, or some other independent way in which one's life goes worse. Then clearly there can be privacy diminishments and privacy wrongs without there being privacy harms. But there cannot be privacy harms without there being either a privacy diminishment or a privacy wrong. This is so in virtue of causal constraint that goes into the definition of a privacy harm. Clearly, however, there can be harms that are *just like* privacy harms, except that they involve access to facts that are not in the protected zone, or types of epistemic access to sensitive facts that do not conventionally constitute a privacy wrong, or access not to facts, but to falsehoods.

Sensitive facts?

The generic epistemic theory raises two very natural questions. The definition of privacy diminishments depends on a prior determination of the class of sensitive facts. How is that

1 Thanks to associate professor Janne Rothmar for help with the legal issues.

https://doi.org/10.1017/epi.2013.15 Published online by Cambridge University Press

class delimited? Second, what types of epistemic access to sensitive facts matter? The former question has been subject to some debate in the literature on privacy, but the second question has not. I want to suggest that the conventionalist nature of plausible answers to the first question places certain limits on the sort of precision and detail we can expect regarding the second question. I start with a few remarks about the first question, and then turn to the second question in section 3.

Usually we think of sensitive facts as including facts about our religious denomination, health, sexual preferences and certain parts of our lifestyle. Sometimes one's political views are regarded as sensitive, as are facts about one's financial affairs. Often one's true views about colleagues and friends would be protected. What makes these, rather than some other class of facts, enter the protected zone?

I want to suggest that it is more or less a matter of social convention what is in the protected zone, and what is not. As is already indicated, these conventions feature a complicated form of contextual dependency, and surely they are also vague in many respects.

Some support for this conventionalist view may be obtained from considering various alternatives. One could imagine a sort of primitivist view according to which there are intrinsically private facts. On such a view, some features of our life and existence just deserve informational protection *per se*, irrespective of individual or social consequences of protecting this information. Such a view seems to have few if any defenders in the philosophical literature.

Parker defines privacy as 'control over when and by whom the various parts of us can be sensed by others', where 'parts of us' includes 'objects very closely associated with us' (Parker 1973: 281). This suggests a view according to which sensitive facts can be defined in terms of physical proximity. Clearly, this is inadequate. No matter how close I am to you or your belongings, I may be unable to see or otherwise sense what your deepest religious views are. But if I can manage to get a copy your private email correspondence, I may be able to find out, even if I never get physically close to you or any of your possessions.

Most commentators assume that privacy concerns are grounded in a variety of practical interests. Adapted to our current question, the idea would be that whether some domain of facts should count as sensitive depends our interests, and on the general consequences of permitting various types of access to these facts, as well as the consequences of attempting to uphold a convention placing such facts in a restricted zone. I want to endorse a conventionalist version of this sort of view.

Let's briefly examine a couple of the most prevalent of such ideas. Adam Moore agues that privacy rights are necessary for human flourishing or well-being (Moore 2003: 220ff.; see also Moore 2008). What he really seems to say, however, is that periods of isolation or seclusion are necessary for well-being and flourishing, and apparently, this is supported by numerous animal experiments as well as anthropological findings. Clearly, however, this does not directly concern informational privacy as defined here. One can be secluded or isolated in a physical sense (there are no others in one's physical proximity), or in a social sense (one is not involved in social relations or exchanges), and yet a wide range of sensitive facts about one might be known by others. Or, conversely, one can be physically surrounded by people and socially entrenched with others, and yet a range of one's sensitive information is kept private. So, our practical interest in informational privacy is distinct from whatever interest we have in physical and social seclusion, though of course physical and social seclusion may be a means of protecting privacy.

https://doi.org/10.1017/epi.2013.15 Published online by Cambridge University Press

James Rachels suggests that 'privacy is necessary if we are to maintain the variety of social relationships with other people that we want to have and that is why it is important to us' (Rachels 1975: 326). But though Rachels in places does talk about informational privacy, this is probably not what he first and foremost has in mind. Rather, his suggestion is that 'our ability to control who has access to us, and who knows what about us, allows us to maintain the variety of relationships with other people that we want to have ... [is] ... one of the most important reasons why we value privacy' (Rachels 1975: 329). So, it is the power to control physical and social access to us that basically matters, according to Rachels. But, again, this does not seem to have a direct bearing on informational privacy. I may have the power to keep you away from me and not to get socially involved with me, unless I so decide. And yet for all that, you could know all sorts of sensitive facts about me.

It seems unlikely that there would be one unified type of practical interest that grounds privacy concerns. It rather seems that there is a variety of different reasons that might be relevant, and here are the more obvious candidates. (i) We have a sense of embarrassment which means that we experience shame or discomfort when certain facts about us are shared, whereas distribution of other facts leave us indifferent. Our sense of embarrassment is no doubt very culturally malleable. (ii) In many interactions, there are clear strategic reasons, including risk of abuse or blackmail, for why we should not want certain information to fall into the hands of others.

(iii) In his brilliant essay 'Concealment and Exposure' (Nagel 1998), Nagel points out that many forms of informational privacy are important for facilitating smooth collaboration and joint deliberation between individuals who might disagree about many issues, or whose interpersonal relations might suffer or be destroyed beyond recognition were all personal views not only known by others, but also acknowledged in between then. This is different from the concern about strategic interest. It may be in my strategic interest that everyone is familiar with my low regard for my colleague (say because I believe that this will intimidate him, or impose a social cost to others who might otherwise collaborate with him, or have high regard for him). But Nagel's point is that we have conventions governing access to and dissemination of such information, and these conventions serve the general interest of smooth collaboration.

A distinct concern has to do with (iv) reasons of equality and fairness, an example of which I gave earlier. Taking medical information out of the circuit on the labour market, for example, is one of the ways in which one can attempt to even out brute bad luck in the natural lottery.

Of course, I don't mean to suggest that these practical reasons are exclusive or exhaustive. Also, I don't assume that each of them grounds only privacy concerns. Obviously, we might for example have strategic interests in protecting information, though this information does not count as sensitive in any way relevant for privacy concerns. Think for example of the politician who does not want the facts about his corruption to become widely known. Such facts are, in my terminology, typically not in the protected zone, so it would not be a privacy diminishment if others get access to them.

Yet, I want to suggest that privacy concerns trace back to such more basic moral concerns. It will be a matter of controversy what they boil down to, but for the purpose of this paper, I just assume that it is a matter promoting individual well-being and autonomy, and promoting equality and fairness.

Basically, I want to suggest that these concerns shape and justify (to the extent that they are justified) two sets of privacy conventions: *content conventions* that determine the

https://doi.org/10.1017/epi.2013.15 Published online by Cambridge University Press

content of the protected zone in various contexts, and *access conventions* that determine which types of epistemic access constitute a privacy wrong (more about this later). There is no underlying distinct moral issue about informational privacy. Informational privacy is a higher-order moral issue – an issue that depends on a range underlying moral factors, as well as on contingent features of the world. In part, this is why the philosophical questions that arise about informational privacy are difficult and confusing.

3. WHAT KIND OF EPISTEMIC ACCESS MATTERS FOR INFORMATIONAL PRIVACY?

According to the generic epistemic theory of informational privacy, privacy depends inversely on epistemic access. This immediately raises the question: what sort of epistemic access matters for privacy? This question in turn covers at least three more specific questions. (i) Epistemic access to facts is a modal notion, and the question is how to get the modals right? (ii) A second question concerns which *epistemic states* the relevant kind of access should be taken to involve. (iii) Finally, there is a question concerning which *epistemic pathways* the relevant kinds of access involve. I discuss these questions in turn.

Epistemic access: control and limitation theories of privacy

According to the generic epistemic theory, privacy depends inversely on epistemic access. There are, however, two other prominent theories of privacy on the market, the control theory and the limitation theory. We may interpret these theories as offering a broad and in certain respects demanding interpretation of the modality involved in privacy.[2] When adapted to other parts of the framework proposed here, they say roughly the following (for an overview and further references see Matheson 2007; Tavani 2007):

The control theory of privacy
Some individual A is private with respect to a domain of facts P relative to B, just in case B's epistemic access to P is under A's *control*.

The limitation theory of privacy
Some individual A is private with respect to some (sensitive) domain of facts P relative to some other agent B, just in case B's epistemic access to P is *restricted*.

Restricted access here typically refers to physical restrictions. So, B's access to certain facts about me is restricted if I am in my house, while B not, and I am therefore out of B's view, or if the facts are noted in my diary which is locked up in a safe that B cannot enter.

The control theory and the limitation theory both face a similar challenge. Whether some fact is private or not does not depend at all on whether it has been epistemically accessed by others. Rather, a fact's being private simply turns on whether certain barriers

2 Thanks to Don Fallis for suggesting that one can treat the control and limitation theories as a part of the modal question.

https://doi.org/10.1017/epi.2013.15 Published online by Cambridge University Press

to epistemic access are in place. So, according to control and limitation theories, barriers to access, or rights or capacities to manage such barriers (control), is part of what *constitutes* some fact's being private. On the generic epistemic theory, these would be means of protecting privacy, rather than constituents of privacy.[3]

This seems counterintuitive for a number of reasons. Consider the control theory. Suppose that it is out of my hands (as is indeed the case) whether my financial record is accessed by someone who shouldn't see it: whether this happens is very much up to various banks and credit institutions, and to the tax authorities, or course. According to the control theory of privacy, my privacy is diminished or absent with respect to these matters, simply because I lack control, even if no part of this sensitive information has ever been leaked.[4] Intuitively it seems that the control theory gets the modalities wrong.

The limitation theory faces a parallel problem. Suppose someone forgets his diary containing very sensitive information on a bench in a public park. I find it, and am in a position to take it and read it. There are no barriers to access, and the author has lost any control. But still, privacy is infringed only if I read the diary, not if I hand it back to the owner without having read it, as I of course should. It is not, of course, that restrictions and control over epistemic access are unimportant for privacy concerns. The only point is that it seems odd to define privacy diminishments in terms of these, as control and limitation theories of privacy tend to do, rather than directly in terms of epistemic access.[5]

Maybe, however, this is a merely verbal issue. Restriction and control on the one hand, and epistemic access on the other, are two independent parameters cutting across each other, yielding four different types of cases. Obviously, there are cases where sensitive information is subject to restriction and control, and has not been epistemically accessed. But one might argue that there are also practically significant real-life cases in which sensitive information *is* subject to restriction and control, but has nonetheless been epistemically accessed without proper authorisation. And surely there are cases in which sensitive information is not subject to restriction and control, yet for some reason it happens not to have been accessed.

When characterised in terms the two independent parameters – restriction/control and epistemic access – all views would also agree about how to categorise each such case. And judging by the mere content of the theories, there might be no differences regarding the moral import of the different sort of cases. Everyone could agree that epistemic access to sensitive information should be limited, and should be subject to control of certain agents. Everyone could agree that when such information is unprotected, or is not subject to the right kind of control, it poses a risk. Everyone could agree that what matters is not only that unauthorised access to sensitive information does not occur. A worthy but more ambitious aim is that institutions, legislation and social conventions provide control, and

3 In effect this was what Thomson pointed out in her early paper, see Thomson 1975: 304, n. 1. See also the discussion in Matheson 2007.

4 Moore replies to this challenge by suggesting that these objections have force 'only if we construe "control" in a general way. If defined narrowly – say over each instance – her case [Thomson's case] falls apart' (Moore 2003: 218). I have to confess that I do not understand what Moore means here.

5 Moore and Tavani suggest a combination of the control theory and the limitation theory (RALC), see Tavani 2007, 11–12. For what it is worth, it seems to me that their combined theory inherits the problems of control and limitation theories, but I cannot defend this claim here.

https://doi.org/10.1017/epi.2013.15 Published online by Cambridge University Press

ensure that unauthorised epistemic access *cannot easily* occur.[6] Likewise, everyone could agree about the moral difference it makes whether some item of sensitive information has actually been accessed or not, just as there need be no disagreement about the conditions under which certain kinds of epistemic access to sensitive information are morally objectionable, and when it is not. What the two general classes of privacy theories disagree about might then merely be the use of labels like 'private', or 'privacy diminishment'. If all this is true, it seems that we basically face a terminological disagreement between epistemic theories on the one hand, and control and limitation theories on the other.[7]

So, let's assume that privacy depends inversely on epistemic access *simpliciter*, not on barriers to, or control over, such access. This takes us to a second modal question. When, in epistemology, we say that some subject S has epistemic access to a true proposition *p*, we usually mean either that S knows (or justifiably believes) that *p*, or we mean that S is in a position such that S can easily acquire knowledge (or justified belief) that *p*, say by directing her perceptual attention to the fact that *p*, though S may not yet have done so. So the expression 'epistemic access' usually contains a modal in need of interpretation. Which modal force of 'epistemic access' would be most appropriate in a definition of informational privacy? In large part, this is a matter of settling on a convenient manner of speaking. I suggest that we say that, if someone could very easily acquire knowledge of some sensitive fact that *p*, but has not done so, then this fact is still private, though it obviously on the verge of not being private much longer. Of course, not much hinges on this decision, as long as one keeps the terminology clear.

Epistemic states

The next question concerns epistemic states. Privacy diminishments are an inverse function of epistemic states regarding sensitive facts, but I left open which epistemic states matter. We can easily produce a list of candidate states. It is natural to suggest that knowledge is an epistemic state that matters. But an equally plausible candidate would be justified true belief, where the level of justification is substantial, but does not quite suffice for knowledge. Yet another candidate would be barely justified true belief, by which I mean a propositional state that is sustained with as little epistemic justification as is possible for this state to be a belief-state, rather than some other type of state. (See also Blaauw 2012 and Fallis 2012 for different views.)

6 Note an illuminating analogy to the distinction between the epistemic theory of privacy and the limitation and control theories. The analogy is to the distinction between a liberal and and republican notion of liberty. The liberal holds that liberty consists in non-infererence in one's choices. The republican proposes a notion of liberty according to which freedom is non-domination. I enjoy non-domination when no one is in a position to interfere arbitrarily with the choices that I am in a position to make (see Pettit 1997). So, the epistemic theory of privacy stands to the limitation and control theories as the liberal theory of liberty stands to the republican theory.

7 David Matheson objected to this in discussion, and insisted that there would normally be a non-verbal disagreement between the two views regarding the conditions under which a privacy diminishment occurs. I find this difficult to see. The two types of theories may agree about every underlying feature of the cases, including all the morally significant ones. In that case they would merely disagree about whether to apply the label 'privacy diminishment'. When the proponent of the control or limitation theory insists that someone's privacy has diminished in a particular case, he means something slightly different from what a proponent of the generic epistemic theory would mean using the same words about the case.

https://doi.org/10.1017/epi.2013.15 Published online by Cambridge University Press

Apart from this spectrum, there are familiar deviant cases to consider. Take the case of a true belief held on the basis of *misleading* evidence. Can such epistemic states constitute privacy invasions in the same way as knowing a sensitive fact can, or at least in similar ways? The list may go on, of course, and we may consider many other epistemological distinctions and scenarios, say high stakes/low stakes scenarios, clairvoyance, misleading but unprocessed information in the environment, and so on. And there are issues about how to represent the doxastic states involved. Should we think in terms of the tripartite schema, according to which one can either fully believe *p*, fully believe non-*p*, or neither believe *p* nor non-*p*? Or should we go for something more fine-grained, say representing degrees of belief by an real number between [0;1]?

Or consider this case: I know some sensitive medical fact about some individual, identified to me only by his social security number. Unbeknownst to me, this social security number uniquely identifies my next-door neighbour, Adam. Adam despises me, and he would hate me knowing these facts about him. Has Adam's privacy with respect to me been diminished by the knowledge about him that I have acquired? Such cases, of course, point to a general issue about the mode of presentation of the sensitive facts over which privacy diminishment is defined.

For each of the well-known possibilities and scenarios so familiar from the epistemological literature, we might ask how they affect privacy diminishments. I am, however, sceptical whether these questions admit of precise and categorical answers. In large part, this has to do with the conventionalist nature of privacy concerns that I outlined above. I suggested that content conventions determine sets of sensitive facts, and access conventions determine modes of access that constitute privacy wrongs. The conventions, in turn, are shaped by underlying concerns about protecting and promoting individual wellbeing and autonomy, and protecting and promoting various forms of equality and fairness. Conventions, however, are generally blunt instruments. They are framed to handle paradigmatic cases. This means that they are likely to be vague or indeterminate regarding esoteric possibilities, and those mentioned above are exactly such esoteric possibilities. So, if the story involving content and access conventions is on the right track, we shouldn't be surprised if these epistemological distinctions are in a sense too fine-grained for the purpose of understanding the epistemological component of privacy.

Insofar as there are answers permitting some level of precision, it is difficult to believe that they could sustain categorical views. By categorical views I mean views asserting, for example, that the epistemic state of knowing, and only that, matters for privacy. According to such views, justified but yet not known true belief has no effect whatsoever on privacy diminishment. This is hard to believe. If knowledge affects privacy, as it surely does, then so do epistemic states that are in many respects just like knowledge, among them highly justified true beliefs.

Knowledge, of course, is a paradigmatic epistemic state that our privacy conventions should be expected to regulate. But there is a another state that is worth calling attention to, and this is what Nagel (1998) is concerned with:

Public acknowledgement
A proposition P is publicly acknowledged by individuals A and B just in case A and B are both committed to treat P as true for the purpose of joint decision making and collaboration, and A and B both know this.

https://doi.org/10.1017/epi.2013.15 Published online by Cambridge University Press

This is meant to be a very rough indication of a notion of public acknowledgement, and I don't want to consider various ways of stating it more precisely. Clearly, as defined, public acknowledgement is different from knowledge and belief, including collective knowledge and belief. A proposition can be publicly acknowledged, though some or even all of the individuals involved in this public acknowledgement don't believe it, or even know it to be false. Conversely, some fact may be known or sincerely believed by everyone involved, and yet not publicly acknowledged, though again this would often be an unstable situation.

Here is a case to suggest why public acknowledgement is important for privacy concerns. I happen to know some very intimate and sensitive fact *p* about you. You know that I have this knowledge. However, I signal that I want to keep it out of our interactions, say by clearly pretending that I am agnostic about *p*, or that I believe non-*p*. I can do so by acting in ways that would seem to suggest that I don't know whether *p*, or don't believe that *p*, even when you know pretty well that I know that *p*. You trust me on this, so in your dealings with me, you are confident that I will not take advantage of what you know that I know. The crucial point, of course, is that I refrain from publicly acknowledging the fact that *p*, and that I signal that I don't want to do so. As a result, my knowledge of a sensitive fact about you does not harm you in the way that it could otherwise have.

Epistemic pathways

As noted, we need a distinction between the epistemic states that turn a sensitive fact non-private, and the particular epistemic pathway by which someone has acquired this epistemic state (cf. Johnson 1989). To see this, consider another case. Suppose it is a sensitive fact about my next-door neighbour Adam that he supports a neo-Nazi political party. Consider now a variety of ways in which I can get to know about this fact: I may drill a hole in the wall separating our houses, I may tap Adam's phone or spy on his mail correspondence. Surely, these would all be privacy wrongs.

But suppose that I learn about Adam's membership in the party in a very different way: an acquaintance of mine, Bertrand is also member of the neo-Nazi party. Bertrand confides the fact of his membership to me. Another friend of mine, Charles, doesn't know that Adam and Bertrand are both members of the neo-Nazi party. But Charles happens to know that Adam and Bertrand are members of the *same* party. Putting these bits of information together, I now know that Adam is a member of the neo-Nazi party. Though Adam's privacy regarding this fact is now diminished with respect to me, it no longer seems that I have perpetrated a privacy wrong. The reason, of course, is that I used a legitimate epistemic pathway.

Consider then a variation of the case. I try to find out about Adam's political opinions by breaking into his house and reading his mail. I find the letter I have been looking for, and read it. However, the crucial sentence confirming his party membership turns out to be written in German, and printed in gothic letters. So, though the sentence I look at in fact welcomes Adam as a new member of the glorious party, I fail to comprehend the sentence, and therefore form no beliefs about his membership. I notice no other facts about my neighbour (assume that I am completely obsessed with determining the nature of his political views).

Again, intuitively, this would seem to be a privacy wrong, and again this is because it makes use of an illegitimate epistemic pathway to access protected information. But actually, I did not get epistemic access to the crucial fact; although I visually inspected a

https://doi.org/10.1017/epi.2013.15 Published online by Cambridge University Press

sentence expressing the crucial proposition, I didn't manage to form the epistemic state I aimed for. Thus, Adam's privacy didn't decrease. So, while privacy diminishments depend on epistemic states, privacy wrongs depend on epistemic pathways.

Consider finally this case. My neighbour Adam is actually an anti-fascist activist trying under false pretences to get admitted to a neo-Nazi party because he wants to expose their reprehensible activities. The letter I read, after having broken into his house, states that Adam is not, never was and never will be considered for membership in this group (he has been found out!). However, because of my poor grasp of German, and because of the gothic letters, I miss a couple of crucial negations. I am now confident that Adam is a full-blown honorary lifetime member of the neo-Nazi party.

This not a privacy diminishment as defined above, since this notion was defined over facts or true propositions. Yet, it surely seems that I have perpetrated a privacy wrong on Adam, and this is because I made use of an illegitimate epistemic pathway. It seems irrelevant for the moral wrongdoing that the fact I was after turned out not to be there, and that the epistemic state I form is a false belief. Depending on what I choose to do – now that I am confident that Adam is a convinced neo-Nazi – Adam may suffer a harm that is in all respects similar to a privacy harm, except that it derives from a falsehood. In this respect, the case is similar to the case about public acknowledgement above, in which a false but publicly acknowledged belief may cause a type of harm that is indistinguishable from privacy harms, except that it involves a falsehood.

4. SUMMARY AND THREE EPISTEMIC DEPENDENCIES

We can summarise some of these considerations as follows. I have suggested three categories by which we can disambiguate talk about privacy concerns: privacy diminishments, privacy wrongs and privacy harms. There can be privacy diminishments without privacy wrongs or privacy harms, and privacy wrongs without privacy diminishments. Privacy harms are defined as harms that arise from either privacy diminishments or privacy wrongs. So, there are two different forms of privacy harms. As a result there can be privacy harms without privacy wrongs, and privacy harms without privacy diminishments, but not without either. There can be harms that are entirely similar to privacy harms, yet they depend on neither privacy diminishments, nor privacy wrongs.

So, the notions of privacy diminishments, privacy wrongs and privacy harms can come apart in various ways. The underlying reason for this is that the notions of privacy diminishments, privacy wrongs and privacy harms are variously subject to three forms of *epistemic dependencies*. Privacy diminishments are *truth-dependent*. In addition, privacy diminishments are *state-dependent*: they require certain epistemic states, such as knowledge or justified belief, to occur. Privacy wrongs, however, are *path-dependent*, they depend on epistemic pathways, certain ways of trying to acquire epistemic states about sensitive facts. Certain pathways are illegitimate, while other ways of gaining access to the very same facts are not. Privacy wrongs, however, are neither truth-dependent nor are they state-dependent. Privacy harms inherit either truth-dependency or path-dependency because of the causal link to privacy diminishments, or privacy harms. But apart from this, privacy harms do not require any particular epistemic state to occur.

What epistemic states serve to make facts non-private? Knowledge is one epistemic state that fits in; the same seems to go for substantially justified true beliefs. I haven's

https://doi.org/10.1017/epi.2013.15 Published online by Cambridge University Press

said much about what epistemic pathways are required for a privacy wrong, but surely this is highly contextual matter. The conventionalist nature of privacy concerns should make us suspicious about hoping for too much precision.

Some may be inclined to want more alignment in the conceptual scheme, and in part this is what the debate about the epistemic dimensions of privacy is about. There are two classes of cases that come to mind:

1. Failed attempts to invade privacy may lead to false beliefs or to unjustified true beliefs, or to no belief at all. We might want to say that, in such cases, there is not only a privacy wrong, but also a privacy diminishment. If we do so, we should revise the truth-dependency and the state-dependency of privacy diminishments.
2. Beliefs about sensitive matters may be the cause of harm that is very similar to privacy harms, no matter how the beliefs are acquired, and even when the beliefs are false. These harms can even arise as a result of propositions being publicly acknowledged while most agents privately regard them as false. We might want to say that in such cases there are genuine privacy harms. If we do so, we should either revise the notion of a privacy harm such that it no longer causally depends on either privacy wrongs or privacy diminishments. We would then get rid of the truth-dependency and the path-dependency.

Alternatively, we might revise the notion of a privacy diminishment, such that it no longer is truth-dependent. We could say, for example, that when a sensitive proposition is publicly acknowledged, it is no longer private, no matter whether it is true or not. This would remove the truth-dependency of privacy diminishments. We could also revise the notion of a privacy wrong to remove its path-dependency. So, for example, participating in the public acknowledgement of some sensitive information about someone else could be viewed as one way of perpetrating a privacy wrong, even if one does not make use of any specific epistemic pathway to access the facts in question.

To some extent, these choices are terminological decisions to be made. As indicated, our decisions will influence the epistemic properties of the notions used to characterise our privacy concerns, and this is why they might be of relevance to the general question of the epistemological dimensions of privacy. But while it may be practically important to characterise the epistemic states and the epistemic pathways that we worry about in privacy concerns, it may not be too important how we fit these pieces into a conceptual scheme for characterising privacy concerns, as long as we are clear about what we are doing.

Presumably, content and access conventions are continuously subject to social negotiation. What we most fundamentally want is that this transformation is informed by the moral goals that rationalise the privacy conventions. As a response to changing technological and social circumstances, privacy conventions should be revised so as to best promote the aims of well-being, autonomy, equality and fairness. For example, vastly increased computerised data-mining capacities may imply that apparently insignificant and probabilistic individual bits of information about individuals turn into a privacy hazard, and we might want our privacy conventions to change accordingly.[8]

8 An earlier version of this paper was presented at Epistemological Problems of Privacy and Secrecy Conference, Delft University of Technology, June 2012. I am grateful to participants for comments made in conversations and during presentations, especially Martijn Blauw, Don Fallis, Sandy

https://doi.org/10.1017/epi.2013.15 Published online by Cambridge University Press

REFERENCES

Blaauw, Martijn. 2012. 'The Epistemology of Privacy.' Epistemological Problems of Privacy and Secrecy Conference, Delft.

Fallis, Don. 2012. 'Privacy and Epistemic Access.' Epistemological Problems of Privacy and Secrecy Conference, Delft.

Johnson, Jeffery L. 1989. 'Privacy and the Judgement of Others.' *Journal of Value Inquiry*, 23: 157–68.

Matheson, David. 2007. 'Unknowableness and Informational Privacy.' *Journal of Philosophical Research*, 32: 251–67.

Moore, Adam D. 2003. 'Privacy: Its Meaning and Value.' *American Philosophical Quarterly*, 40(3): 215–27.

——. 2008. 'Defining Privacy.' *Journal of Social Philosophy*, 39(3): 411–28.

Nagel, Thomas. 1998. 'Concealment and Exposure.' *Philosophy and Public Affairs*, 27(1): 3–30.

Parker, Richard. B. 1973. 'A Definition of Privacy.' *Rutgers Law Review*, 275: 275–96.

Pettit, Phillip. 1997. *Republicanism. A Theory of Freedom and Government*. Oxford: OUP.

Rachels, James. 1975. 'Why Privacy is Important.' *Philosophy and Public Affairs*, 4(4): 323–33.

Tavani, Herman T. 2007. 'Philosophical Theories of Privacy: Implication for and Adequate Online Privacy Policy.' *Metaphilosophy*, 38(1): 1–22.

Thomson, Judith Jarvis. 1975. 'The Right to Privacy.' *Philosophy and Public Affairs*, 4(4): 295–314.

KLEMENS KAPPEL is Associate Professor at the University of Copenhagen, Department of Media, Cognition and Communication, Division of Philosophy. His main research area is in social epistemology and related questions in political philosophy. He is currently directing a research project called 'The Special Role of Science in Liberal Democracy'.

Goldberg, Alvin Goldman Axel Gelfert, John Hardwig, David Matheson, Pierre le Morvan, Boaz Miller, Philip Nickel and Rik Peels. Thanks also to an anonymous referee from this journal.

https://doi.org/10.1017/epi.2013.15 Published online by Cambridge University Press