

Privacy: Essentially Contested, a Family Resemblance Concept, or a Family of Conceptions?

Kieron O'Hara

*Web and Internet Science Group
Electronics and Computer Science
University of Southampton
Highfield
Southampton SO17 1BJ
United Kingdom
kmo@ecs.soton.ac.uk*

Presented at the 2018 Amsterdam Privacy Conference, 7th October, 2018

Introduction

How to understand the concept of 'privacy' and related ideas has been very controversial for a number of years; Solove famously referred to privacy as a 'concept in disarray' (Solove 2008), while Thomson (1975) argued that "perhaps the most striking thing about the right to privacy is that nobody seems to have any clear idea what it is". This certainly goes beyond matters of semantics. For instance, understanding the value of privacy (Rössler 2005) depends on how it is conceptualised, while the apparent incoherence of privacy has led some to call for its abandonment or downgrading as a normative principle (Jarvis 2011). It also is an important matter of policy, as data becomes an increasingly vital element in the global economy, powering the rise of technology companies, while lawyers and policymakers struggle to punish abuses such as the recent scandal involving Cambridge Analytica's misuse of Facebook data, without shackling innovation or destroying the social value of very large networks.

Different diagnoses of the definitional problem have emerged, in particular that privacy is simply a disparate set of ideas in need of typology, not unification (Prosser 1960, Post 2001, Koops et al 2017), an essentially-contested concept (Gallie 1956, Collier et al 2006, Mulligan et al 2016), or a family resemblance concept (Solove 2008). This paper will examine the multiplicity of interpretations of privacy, as well as these explanations, and will argue that the situation is somewhat less complex than these commentators argue.

The paper consists of four major sections. The first will consider the heterogeneity of privacy concepts, and the second will consider some explanations and responses that this has provoked. The third will suggest an alternative explanation, while the fourth will illustrate how this explanation will help to separate out the different purposes and assumptions of different privacy discourses.

The heterogeneity of privacy concepts

The most common way of demonstrating the heterogeneity of privacy concepts is simply to survey them, and to lay out the dissimilarities. For example, Solove:

Thus privacy is a fundamental right, essential for freedom, democracy, psychological well-being, individuality, and creativity. It is proclaimed

inviolable but decried as detrimental, antisocial, and even pathological. Some claim that privacy is nearing extinction; others argue that the threat to privacy is illusory. It seems as though everybody is talking about 'privacy,' but it is not clear exactly what they are talking about. (Solove 2008, 5).

Or Jarvis, who has less sympathy with privacy as a normative principle:

Do you feel any closer to a definition of privacy? I don't. I see a confused web of worries, changing norms, varying cultural mores, complicated relationships, conflicting motives, vague feelings of danger with sporadic specific evidence of harm, and unclear laws and regulations made all the more complex by context. (Jarvis 2011, 101-102)

It is hard to disagree with Jarvis on this particular point, even if one does not want to endorse his sweeping anti-privacy conclusion. In this paper, I will not query the assumption that privacy is massively, bewilderingly heterogeneous. However, these commentators are not all making the same point. Compare Thomson, quoted earlier, or Prosser:

The law of privacy comprises four distinct kinds of invasion of four different interests of the plaintiff, which are tied together by the common name, but otherwise have almost nothing in common except that each represents an interference with the right of the plaintiff ... 'to be let alone'. (Prosser 1960, 389)

These arguments appear similar to those of Solove and Jarvis, but there is an important difference. Thomson is looking for the *rights* to privacy in law (presumably, even more parochially, in US law). Her response, in a footnote, to a definition of privacy as a type of control is to speculate that "the right to privacy, then, is presumably the right to this kind of control." The conclusion of her paper is that "every right in the right to privacy cluster is also in some other right cluster", and that therefore the right to privacy *in (US?) law* is redundant. Nothing about her argument implies anything about the *concept of* privacy, only that you won't find out what it is by pondering (American?) law books. Similarly, Prosser does not leave the law library. Both authors thereby duplicate the project of Warren and Brandeis (1890), in whose justly classic paper it is assumed that the reader knows what privacy is, and which consists largely of a search for protections within US law. These are specific projects which no doubt have a bearing on philosophical discussions of privacy, but cannot produce definitions themselves. Analogously, there is no right to champagne in US law, but that does not thereby mean there is no such thing as champagne. There is a partial right to die in US law (in some states), but it has little bearing on the concept of death. In general, it is a fallacy to say that, if *the right to X = the right to Y*, then $X = Y$.

If we revisit the quotes from Solove and Jarvis that *are* concerned with defining privacy conceptually, we find other problems. If we adapt Solove's rhetorical paragraph about privacy quoted above to the concept of taxation, for example, the problems with it can be illuminated.

Thus taxation is a fundamental duty, a means to keep the state going, promote citizenship and inclusion, redistribute wealth and decrease inequality, promote certain types of behaviour and disincentivise others, or maintain the social contract. It is decried as state-sponsored theft, oppression and a disincentive to thrift, innovation or self-reliance. Some claim that we should tax wealth, others

income, property, carbon dioxide, consumption, cigarettes, imports, luxuries, land, financial transactions or inheritances. Some want taxes to be progressive, others at a flat rate, and still others want negative income taxes. There is barely any similarity between, say, the administration of a medieval tithe and a value-added tax in a 21st century digitised economy. It seems as though everybody is talking about 'taxation,' but it is not clear exactly what they are talking about.

The first five sentences of this paragraph are all expressions of genuine disputes, but the conclusion is remarkably odd – surely these disagreements about taxation do not imply we don't know what it is. So we shouldn't take the mere fact of heterogeneity of privacy discourse as an indication that the concept is in disarray.

With this caveat in mind, let us consider some responses to the apparent disarray of the concept of privacy.

Responses and explanations

In this section, I will focus on two specific explanations of the problem, that of Solove (2008), that privacy is a family resemblance concept (Wittgenstein 1953), and that of Mulligan and colleagues (2016), that privacy is an essentially contested concept (Gallie 1956). The reason these authors seek such explanations is that they worry that without them, the concept of privacy will prove to be incoherent, and therefore not something that could sensibly feature in meaningful moral discourse. Any normative injunction to protect the privacy of subjects would be incoherent without an alternative, such as a reductionist approach which analysed privacy away in terms of other, better anchored concepts.

Family resemblance

Solove's argument draws on Wittgenstein's notion of *family resemblance* (1953, §§65-71). When we look at instances of particular concepts – Wittgenstein writes of games in the first instance – “you will not see something that is common to *all*, but similarities, relationships, and a whole series of them at that ... a complicated network of similarities overlapping and criss-crossing: sometimes overall similarities, sometimes similarities of detail” (1953, §66). Solove suggests that privacy is such a family resemblance concept – there is nothing in common to all instances of privacy, but they form a kind of chain of similarity and analogy – and hence we should not look for a common denominator or means of generalising across cases of privacy, and the failure to find necessary and/or sufficient conditions does not mean that the concept is incoherent. Solove welcomes the use of privacy as an umbrella term, bringing together interconnected but distinct phenomena (2008, 45).

The target of Wittgenstein's critique was a view of language and definition which he himself had espoused in previous works (1961), and which he had inherited from Frege (1980) and Russell (1956), a view that concepts can be defined in ways that (a) analyse the concept into some combination of more primitive concepts, and (b) describe instances of the concept. So, for example, if we define privacy as control of access, then we (a) understand privacy in terms of these other concepts, and (b) are able to deduce certain facts about any instance of privacy (viz. that someone is in control of some kind of access to himself). On this view, a concept is incoherent in the absence of such a definition, although it is perfectly possible for someone to use a coherent concept without knowing what the definition is. The definition becomes evident only upon careful analysis, which enables us to apprehend what it is that all

instances have in common. Such a universalist view might be held about all language, or could be narrowed down to apply only to certain carefully choreographed contexts, such as science or the law.

Wittgenstein is therefore addressing a somewhat different problem than Solove's. However, does Solove's use of Wittgenstein help? Wittgenstein certainly argues that family resemblance concepts cannot be defined in terms of necessary and sufficient conditions, but he does not suggest that such concepts cannot be *explained*, and indeed "my concept of a game [is] completely expressed in the explanations I could give" (1953, §75). So it is not obvious that Solove is much further forward; we may have removed the pressure for a definition in terms of a common denominator, but we have not determined what form our explanations will take.

Typically, in a family resemblance concept, we explain via pointing, showing examples or paradigms. Different people point to different paradigms, and it obviously follows that there will be tricky borderline cases where there will be differences of opinion, or a vague application (and we often agree about what is vague or tricky). However, it is clear from the relatively common examples that Wittgenstein uses that he does not think that these phenomena are problematic – we can all recognise a game when we see one, and we are almost always in agreement about the usage of the term 'game'. Wittgenstein's point is that this is only odd if we assume that 'game' must be definable in terms of necessary and sufficient conditions, and/or that all games must have something in common.

Much of this is indeed what Solove is keen to resolve, but the notion of family resemblance doesn't quite provide what is needed. As a negative theory, Wittgenstein's idea is simply to show that a conceptual definition is not necessary to ground agreement in the use of a term. Solove's problem is rather the opposite, for the phenomenon to which he points is that we do *not* agree on the use of the term 'privacy'. The family resemblance argument, if we accept it, shows that we don't need a definition to agree in the use of a term; it does not, however, provide any explanation of how a concept about which we do *not* agree can be coherent. What it does show is that there may be agreement without necessary and sufficient conditions.

In other words, if Solove's point that privacy is a concept in disarray is well-taken, then the disarray will still remain even if we agree that privacy is a pluralistic notion. Wittgenstein's point about games is that we agree in how we use the term, not only in clear cases, but also in analogies (the 'Great Game') and we even agree on where the grey areas are (blood sports, gladiatorial contests, or certain types of bullying), and meaning, for Wittgenstein, is essentially connected with the use of terms. Compare the discussion of taxation above – we agree on use, even though there are many debates about it. Yet Solove's problem with privacy is that, on his account, we don't have that kind of agreement, so adopting pluralism isn't going to help with that.

Solove's positive contribution is an ontology of privacy harms (2008, 101-170), a valuable resource, particularly in the context of legal disputes about privacy, which could help determine what tangible harms have been visited upon a plaintiff. However, note that this ontology is independent of the specification of privacy as a family resemblance concept – the concept of privacy could still be incoherent, or defined in some other way, while the harms and the legal protections in the space could be set out as Solove advocates. Hence neither the philosophical argument, nor the ontology of privacy will help remove disarray, if it exists, from the privacy concept.

Essential contestation

A second attempt to rescue the concept of privacy is that of Mulligan and colleagues (2016), in which they argue that disputes are to be expected, indeed welcomed, because privacy is an *essentially contested concept* (Gallie 1956). Such concepts (ECCs) are normative ideas in which dispute is not accidental, and not a sign of conceptual disarray; examples include social justice, democracy and duty. These concepts evoke disagreement not only about marginal cases, but also about paradigms, and are the locus for adversarial debate. In such debate, both sides broadly agree that they are arguing about the same thing, may jointly acknowledge some paradigm examples, and understand, if imperfectly, the position the opponent is taking. Their disagreement is fundamental, producing arguments both offensive and defensive, and yet the dispute is itself productive, helping inform and generate political institutions and practice that optimise social agreement with respect to the concept.

Hence recognition of privacy as an essentially contested concept could be important, as we could then try to structure debate and institutions to optimise social treatment of privacy, as indeed data protection law tries to do by balancing the rights of the individual with the rights of the data controller to process the data. This would at least deflect the critique of those like Jarvis (2011) who wish to sideline the concept altogether.

In his seminal paper, Gallie sets out seven criteria for recognising ECCs, and Mulligan et al (2016) work through these methodically. Many of the arguments are straightforward, and there is certainly a case to be made. However, it is not clear that the fit is perfect.

The first of Gallie's conditions is that an ECC is appraisive, that is it must signify a valued achievement; Mulligan et al are happy that it is. "The best, most true conception of privacy is both hotly contested and normatively desirable. While the results of its use are not always positive, and those who can demand it do not always choose to avail themselves of it, it is always valued, and those who can successfully claim privacy have power over those who cannot" (2016, 5). They also add that ECCs can be descriptive as well as normative, so that the concept of privacy can "both ... describe states that satisfy the conditions for a particular conception of privacy, and ... posit its normative value" (2016, 6).

The problem with this is that it is somewhat circular. Many types of privacy are not particularly desirable (e.g. solitary confinement), and everyone has to balance the joys of privacy with the joys of being visible to one's network. We pick and choose our moments for privacy. Some are thrust upon us, such as a near-universal demand to go to the toilet in privacy, while in other circumstances privacy is prevented, so that any money I earn must be reported to the taxman. Tech firms that use data heavily argue that our privacy is less valuable to us than the services they provide. In well-functioning communities, a thriving civil society means that people are involved in many private associations, which are private to outsiders, but which involve the deliberate sacrifice of privacy to insiders. So we construct our lives by a careful negotiation of privacy with publicity (Altman 1975), which is not typical for an ECC – we do not, for instance, try to balance episodes of justice with alternative episodes of injustice.

As another example, consider unequal application of principles. If we found out, say, that some people were being treated justly and others unjustly in some respect, we

would all agree (a) that that is wrong, and (b) that the only right solution would be to treat everyone justly. If we found that some people had privacy and others not without good reason, for instance if in a Scandinavian country in which most people had their tax returns published, some well-connected people managed to avoid publication, we would also agree that this is wrong. However, there are two equally good remedies: we could either stop publishing tax returns at all, or we could make sure that even the well-connected had their returns published. Yet we would not reason analogously in the case of justice, that a possible right response would be to treat everyone unjustly. The fact that we have a choice in the privacy case implies that privacy does not have the normative force of a paradigmatic ECC.

In many cases, we broadly agree on what privacy is, and disagree on whether we want it in concrete contexts (for example, Westin's surveys [Kumaraguru & Cranor 2005] divide the public into groups depending on their *attitude towards* privacy, not on diverging definitions). Scott McNealy, in his famous intervention, told us to get over the fact that we have zero privacy, not that we still have privacy on a contested definition of it. The disputes relating to ECCs tend to be the other way round – we all agree that we should do our duty in a concrete context (although we may fail to do so), but the disagreement is precisely over what our duty actually is in that context. A British republican who refuses to stand for the national anthem *God Save the Queen* denies that it is his duty to do so; he does not agree with everyone else that this is a duty he prefers, on balance, to neglect.

So while we might agree with Mulligan et al that “the ability to traverse descriptive and normative explains part of the complexity and dynamism of contests over privacy” (2016, 6), we might go on to note the good deal of agreement about the descriptive and disagreement about the normative. And so privacy does not share the characteristic of many ECCs that there is strong agreement about their normativity, but disagreement about descriptions of concrete situations. We tend to agree whether or not we have privacy, and then dispute whether this is good or bad; if privacy was an ECC, we would expect the opposite, that we would agree that privacy was a good thing, but dispute whether or not we have it in a specific context. Both the ECC account and the family resemblance account are challenged by the fact that we actually have a remarkable level of agreement about what constitutes privacy, and our disagreements are about whether or not it is desirable and when.

As with Solove, the analytical framework developed by Mulligan et al (2016, 10-15) is very valuable, and separates out a number of often-ignored distinctions involved in privacy by design, and this paper is not intended to be critical of it. However, classifying privacy as an ECC fails to explain the nature of the disputes about privacy that the framework is intended to resolve.

The state of privacy

Let us revisit the ECC argument. Another of Gallie's criteria is openness, that an ECC “admits of considerable modification in the light of changing circumstances” (1956, 172). This clearly applies to privacy, as Mulligan et al point out (2016, 7-8), often following technological development. To illustrate openness, they consider two models of privacy, a 19th century model of being let alone, (Warren & Brandeis 1890, Prosser 1960), and a 20th century model of data protection. The concept of privacy shifted to “protect what people valued. ... Practitioners in the later context could have

jettisoned the concept. Rather than discard it, they sought to transform it, and to do so by arguing over its meaning” (Mulligan et al 2016, 8).

Without disputing their claim, there are two interesting wrinkles in this account of openness. The first is that the later context doesn't *transform* the concept of privacy. What we see here is addition or accretion, not transformation. In the new context of databases and information processing, a new aspect or conception of privacy appears, that was not evident before the technology was in place, and this aspect gets added to the already bulging file of phenomena under the rubric of privacy. Being let alone continues to be a *bona fide* type of privacy. Similarly, the idea of a photograph snapped in the street (in public) being private became a topic for discussion only when someone invented cameras that you didn't have to stand absolutely still in front of. And this notion of privacy was added to the then-understood roster of aspects. New technologies (or practices, or norms) subtend new aspects of privacy that don't affect older aspects; there is no transformation of the concept in these cases, beyond addition.

Secondly, the examples Mulligan et al discuss, the problems and the responses, all presuppose a specific *legal* context. As noted above, Warren and Brandeis and Prosser do *not* define privacy; rather, taking it as read that the reader knows what privacy is, they search for privacy protections in US law. Their projects are deliberately both geographically and disciplinarily parochial. Similarly, the responses they discuss to the encroachments of data processing on our privacy are all legal responses, mainly in the US, but also including the OECD's guidelines and EU data protection regulation. None of these responses was conceptual; they were all attempts to cajole the state into protecting breaches of privacy that were antecedently widely perceived.

These debates, then, do not concern the concept of privacy; rather, they concern the important question of what kinds of protection are afforded in law, and lay the ground for political debate as to whether the law needs to be changed at certain times in certain places. Following this thought through, we might follow Warren and Brandeis (1890) in positing that the concept of privacy is entirely separate from its treatment in law. 'Privacy' is not a legal concept, although it is regulated in most jurisdictions, in the same way as 'death' is not a legal concept, although it is regulated. Lord Lucan was declared legally dead in 2016, although nobody seriously thinks that he died, if he is indeed dead, at that point. Some invasions of privacy are harmful, and the law can be used, sometimes, to protect people against them. By examining case law, we can get a sense of the breaches of privacy against which citizens might have strong, weak or no protection, but it does not follow from this that they are not all equally breaches of privacy. The law varies from jurisdiction to jurisdiction, but concepts of privacy do not (though they are valued differently). The law also varies over time; at the time of writing, there are moves in the UK to make 'upskirting' a crime, but this clearly does not mean that upskirting will only become a breach of privacy when it is made illegal. We can contrast the non-legal concept 'privacy' with 'data protection', which *is* a legal concept, and has little or no application outside the law.

Separating the legal discussion about how to regulate privacy from the conceptual discussion of what privacy is clears a lot of the mist. We can see Warren and Brandeis, Prosser and Thomson all trying to clarify what protections privacy has in US law, without discussing conceptual matters in detail. Similarly, many of the critiques of existing privacy definitions are grounded in problems other than

conceptual ones, thereby laying down hard-to-satisfy conditions for an acceptable definition.

Solove's survey is perhaps the most comprehensive (2008, 12-38), and instructive in this respect. Let us take three examples. First, he criticises the view that privacy is the restriction of access to the self (2008, 18-21).

Certainly not all access to the self infringes upon privacy, only access relating to specific dimensions of the self or to particular matters and information. ... In the continuum between absolutely no access to the self and total access, the important question is ... what degree of access should we recognize as reasonable? (Solove 2008, 20)

Second, he criticises the view that privacy is a cultural universal (2008, 65-67), quoting Adam Moore.

While privacy may be a cultural universal necessary for the proper functioning of human beings, its form – the actual rules of association and disengagement – is culturally dependent. (Moore 2003, 223)

A third view to be criticised is that privacy denotes some kind of control over personal information (2008, 24-29). Parker, for instance, interprets this broadly, arguing that "control over who can sense us, is the core of privacy" (1974, 280), but Solove replies that "We are frequently seen and heard by others without perceiving this as even the slightest invasion of privacy" (2008, 25).

The intriguing thing about these three examples, and of many more in Solove's comprehensive review, is that *different* grounds are adduced for each of the counterarguments. In the first example, the definition is criticised, not because it fails to specify something that might be considered private, but rather because it does not discuss what level of access is *reasonable*. In the second case, cultural universal theories fail because they do not take account of *cultural variance*. Thirdly, Parker's particular take on control falls because it covers cases that are not *perceived* as privacy-threatening. These three examples, then, imply that an adequate definition must do a lot of work, defining (i) what level of privacy is reasonable, (ii) how it is implemented in particular societies and (iii) also respecting our perceptions. The scope widens as we look further afield in Solove's book, and elsewhere, to include demands to specify the harms of privacy breaches, to give practical guidance, to respect what courts have said, to explain different opinions about value, and so on.

Conceptual overload

A definition that managed to include satisfactory accounts of privacy's reasonable application, cultural variants, individuals' perceptions, and so on, would be heroic indeed. Yet, given this kind of criticism, it is no surprise to find many definitions of privacy packing a lot of material in. Whether the pickiness of the critiques tracked the overloaded definitions, or vice versa, doesn't matter; the fact remains that many commentators try to define privacy in ways that anticipate political positions respecting privacy.

Let us take one example, not to single it out, except that it is often approvingly cited, from (Agre 1998, 7), defining privacy as "the freedom from unreasonable constraints on the construction of one's own identity". Hildebrandt, for instance, specifically approves this definition as rendering explicit the relationship between privacy and identity, bringing in notions of selfhood, implying that the self is to be constructed

and is not a given, showing that privacy is emergent from many interactions, defining privacy both in terms of freedom-from and freedom-to, and also making an explicit association with what is reasonable and unreasonable (Hildebrandt 2015, 80).

The definition certainly packs a lot in, a whole theory of and outlook on the importance of privacy. But is Hildebrandt right in counting this as a strength? Like Napoleon's army with its stretched supply chain, it is vulnerable to attack from a number of directions. Mulligan et al are correct to argue that we will see permanent contestation if all definitions of privacy are this complex. But should they be?

To take some examples, someone from a search engine or a social network site will believe that what they do to the data is reasonable. Hence it is a consequence of the definition, as they understand reasonableness, that someone on their site whose data they are processing, sharing and selling, retains their privacy because nothing unreasonable is happening to them. By putting 'reasonableness' into the definition, a substantive argument between privacy campaigners and technology companies turns into a semantic discussion. The alternative is to have a less rich definition of privacy which makes no claims about reasonableness. In that case, we can all agree that there is a breach of privacy by the technology companies, and can get on with the difficult debate about whether or not the breach is reasonable.

Secondly, the definition connects the concept with that of identity. One does not have to doubt the connection to have misgivings here. Surely most episodes of privacy have little or nothing to do with identity. If I closet myself away to go to the toilet, or watch a football match, or scratch myself in an intimate place, am I really constructing an identity? Identity construction is surely a minority sport, and these quotidian incidents surely involve privacy.

Thirdly, the connection with identity opens up a number of unwelcome complex questions related to that concept, not privacy. For instance, if the purpose of an episode of privacy at t_1 is to construct an identity I_1 at t_2 , who (which identity) instigated the privacy episode? Does that identity have any say over the outcome of the construction?

Fourthly, the definition is hard to put to work in standard contexts. For instance, no-one is in favour of unreasonable constraints upon themselves, while everyone accepts reasonable constraints, so it follows from the definition that everyone is in favour of privacy. Then it becomes hard to explain why so many people are haemorrhaging their data in the general direction of Facebook, and why privacy is viewed differently by so many people.

And fifthly, the definition does not square with the broad observation that, outside of legal, philosophical and political contexts, we can discuss privacy with a large measure of agreement. People without training in these disciplines, who have little feel for the constraints on identity formation, can and do have perfectly reasonable conversations about privacy, which is hard to explain if (Agre 1998) is our text.

No doubt there are various complex appendices that could be added to this definition that will see off these immediate complaints, and the many others that it might attract, in the same way that adding sufficient numbers of epicycles to Ptolemaic circular orbits will reproduce ovoid Keplerian ones. But it is a lot of work simply to establish the nature of what we are talking about. In the context of this conceptual overload, as Solove argues, we will not get agreement, and, as Mulligan et al argue, we will get permanent contestation, and privacy, defined to be normative, will look very like an

ECC, contrary to the arguments in the last section. Is this the correct understanding? Is there an alternative?

What kind of thing is privacy?

An alternative method would be to define privacy neutrally, and then, once the definition is agreed, conduct explicit arguments, for example about what is reasonable and unreasonable. To bake these ideas into the definition means that debate about such issues becomes a tussle over whether or not to accept the definition, without the fulcrum of agreement about terms. This follows the classic strategies of Warren and Brandeis and Prosser, who separated out the definition of privacy from its protections in law, and expected the ordinary person to understand the nature of privacy (though not privacy law) as well as they did.

Dictionary definitions of privacy actively court such neutrality. For instance, the *Oxford English Dictionary* suggests it is “being withdrawn from society or public interest ... avoidance of publicity”. The advantage of such a bland definition is that, once we agree it, we can then get into all the arguments about when it is reasonable to demand to be in such a state and when not, whether such a state is essential to the construction of identity, under what circumstances are such states protected by law, and so on.

If we produce such a definition, then an alternative diagnosis of the problems raised by Solove and Mulligan et al emerges: that rather than the concept being in disarray or dispute, different discourses were being unhelpfully merged, so that people who wanted a political discussion about privacy were talking past psychologists, moral philosophers, cybersecurity experts, lawyers and other practitioners with different points of view. The aim should be to separate these discourses out, rather than to try to resolve all the different and difficult disputes with a single definition.

I do not intend to produce such a definition now, although I would venture to suggest that there would need to be good reason to go far beyond the dictionary definition. The preceding discussion suggests three important conditions of what such a definition should look like.

- First of all, privacy looks like a state, a condition in which we are either in or not in (compare the *OED*). One may be private at time t_1 , and no longer private at time t_2 .
- Secondly, privacy is a relative state – one may be private with respect to some observers and not with respect to others. Privacy is a relation between oneself (or possibly a group – I leave this question open), and potential interlopers.
- Thirdly, as pointed out by both Solove and Mulligan et al, privacy is ineradicably plural and mutable. My argument about Solove certainly does not rule out that privacy is a family resemblance concept, and I would argue that it is. Thus one can be private in some respects and not in others.

Privacy as a family resemblance concept

If privacy is a family resemblance concept, then we will not be able to produce necessary and sufficient conditions for its application, which is what we would expect from the above discussion. We would expect, in Wittgenstein's words, “a complicated network of similarities overlapping and criss-crossing” (1953, §66). I would suggest,

however, that we can point to some important linguistic markers of where privacy interests might appear.

A key phrase in English which usually stakes a claim to privacy is “that’s *my* business” (cf. Hollander 2001). That phrase is usually intended to warn people away, or to register outrage at an intrusion. More generally, important linguistic markers of areas where privacy may come into play are the first person (singular and plural) possessive adjectives, in English ‘my’ and ‘our’: my information, my body, my personal space, our property, my image, my family, my thoughts, our intimate secret, my friends, my history, my income, our ambitions, my shame, our meeting, my bodily functions and the skeleton in my closet are all suggestive of private matters.

As Mulligan et al suggest (2016), new technologies, and indeed new norms, new practices, powers, freedoms, institutions and opportunities, will each create novel situations where we find ourselves describing certain aspects as ‘my X’ or ‘our Y’, and at that point we may find new privacy conceptions evolving (I am using the word ‘conception’ in something like the sense expounded by [Rawls 1971], a freestanding interpretation derived not from any deep theory, but suggested by relevant developments in society and culture). As we work through these conceptions, we will also develop associated notions of what a transgression, or breach of privacy in that sense would look like, and what might be a remedy. Many privacy breaches have potential or actual aspects (compare, a revenge porn scenario where nude pictures are placed online: this is clearly a gross breach of privacy, but we can distinguish a breach with potential, when the pictures have fortunately not been downloaded at all, from a fuller breach, when the pictures have been downloaded and viewed by observers). These conceptions could be described using the framework of (Mulligan et al 2016, 10-15).

So, for example, the development of instant photography created a new conception of privacy – where one was visible in public (and therefore not private in that sense), one might have one’s privacy breached by the capture and distribution of the image. Files and archives created further conceptions of privacy, and automation further still. Digital photography creates privacy issues beyond those of film-based photography. Law or social norms may formalise these issues, so that, for example, a framework evolved to define and protect private property, and data protection evolved to help individuals protect their own information privacy. But the conception stands apart from the protections afforded (or the breaches legitimised) by the state.

Not all privacy conceptions are of interest to everyone. Facial privacy is clearly very important in certain cultures, and so many people wear veils and feel uncomfortable without them (El Guindi 1999). I, as a middle-aged man in Europe, have my facial privacy breached every day. I understand the idea of facial privacy, but it does not bother me at all. This is a type of privacy which I do not value. In general, breaches of privacy will affect people differently across genders, classes, cultures and generations. Nevertheless, a conception of privacy ought to be explicable to any audience. I should be able to go back in time and explain to a medieval person how a photograph would breach her privacy, and she should understand this, although it is not her problem, and she may not fully grasp the consequences. The conception should be no more difficult to explain than the technology that subtends it.

The assumption of privacy as normative may be related to the likelihood that a privacy conception will evolve alongside some kind of perceived problem or harm – these are, after all, the areas where we will discuss these matters most deeply. Yet this

need not always be the case; we can discuss privacy conceptions that affect nobody at all and cause no harm – for example, we might describe the privacy breaches caused by telepaths in a science fiction novel, but there are no telepaths and no privacy breaches resulting from their reading our minds.

The key thing to note about this interpretation of the myriad privacy conceptions that have emerged over time is that they are neutral statements of a relation between oneself (or one's group) and others. They do *not* tell us when privacy is appropriate, valuable, desired, desirable, legal, enforceable, functional, anti-social, vital for democracy, essential for psychological well-being, good, bad or indifferent, and this is deliberate. These are additional factors that are covered in the various non-conceptual privacy discourses.

Privacy discourses

The next question, then, is what privacy discourses there are. I do not want to make a strong claim, but I will set out seven levels of discourse (O'Hara 2016). I call them levels, and number them, because there appear to be stronger connections between adjacent levels than between any random pair. These seven discourses cover much ground, but I do not claim here that they are complete or sufficient.

Seven levels of discourse

The seven levels of privacy discourse are the following (fig.1).

Level 1: conceptual. Here we find discussion of different privacy conceptions, breaches and remedies. New conceptions appear alongside new technologies and other innovations. They will take the form of state descriptions, and be as value-free as possible. Most of this paper so far has been at level 1.

Level 2: actuality. Here we find the measurement of whether or not a person has privacy in a specific (level 1) sense, relative to another or others. Is there, or has there been, a breach, as a matter of fact? Note that this does *not* equate to any of the following: Have my rights been breached? Has the law been broken? Has harm been incurred? Have I noticed anything untoward? Has anything happened that I care about? We locate descriptions of the data environment (Elliot et al 2018) at this level, as well as cybersecurity, which changes the privacy affordances of technology. We also find descriptions of privacy protective (or neglectful) behaviour here too.

Level 3: phenomenology. This often overlooked level concerns what privacy, or its lack, feels like. It includes questions about whether I even notice, or care about, a breach of privacy (cf. my indifference to facial privacy). Note that phenomenology can only be a rough guide to privacy breaches; my taking a shower while someone watches secretly through a spyhole *feels* exactly the same as taking a shower in private. Similarly the same action, taking off one's clothes, feels very different depending whether one does it alone, or in front of one's mother, or in the Starbucks in Terminal 5 at Heathrow Airport. Design, for social networking sites, is important here, making interactions feel private while the data they generate is siphoned off for corporate use. Academic work in this space includes work on design and privacy (Hartzog 2018) and on creepiness (Shklovski et al 2014).

Level 4: preferences. This level covers my desires to be concealed or visible to my network. It also includes questions of what privacy behaviour others want of me, and my preferences about others' behaviour (including, for example, whether they breach

the privacy of my relatives and loved ones). At this level we find discussions of control, consent and economic models of privacy, including the work of Altman on negotiating privacy boundaries (1975), and Acquisti on the privacy paradox (Acquisti et al 2015). There is no reason to expect privacy preferences to be rational.

Level 5: norms. Here we find privacy-related regularities, conventions and expectations. These will vary across cultures, genders, generations and classes. It is where we find such concepts as reasonable expectations, questions about how privacy norms carry over into new situations, and questions about the relation between privacy and other norms (e.g. free speech). Work at level 5 includes Nissenbaum on contextual integrity (2010).

Level 6: regulation. Here we find work to do with law, and other compulsory privacy codes, which the state (or other organisations with authority) expend resources to enforce. Solove's work is mainly focused at this level, which also includes considerations of how privacy is treated across jurisdictions, human rights as encoded in legally-binding codes, and its regulation using other sorts of law, including contract and tort.

Level 7: politics and morality. Here we find justifications of privacy, questions of non-legal rights, issues to do with the value of privacy and questions of power. Work at this level includes Westin on privacy and democracy (1967), Rössler on privacy and autonomy (2005), Agre on identity (1998), and Etzioni on privacy's undermining of security (1999).

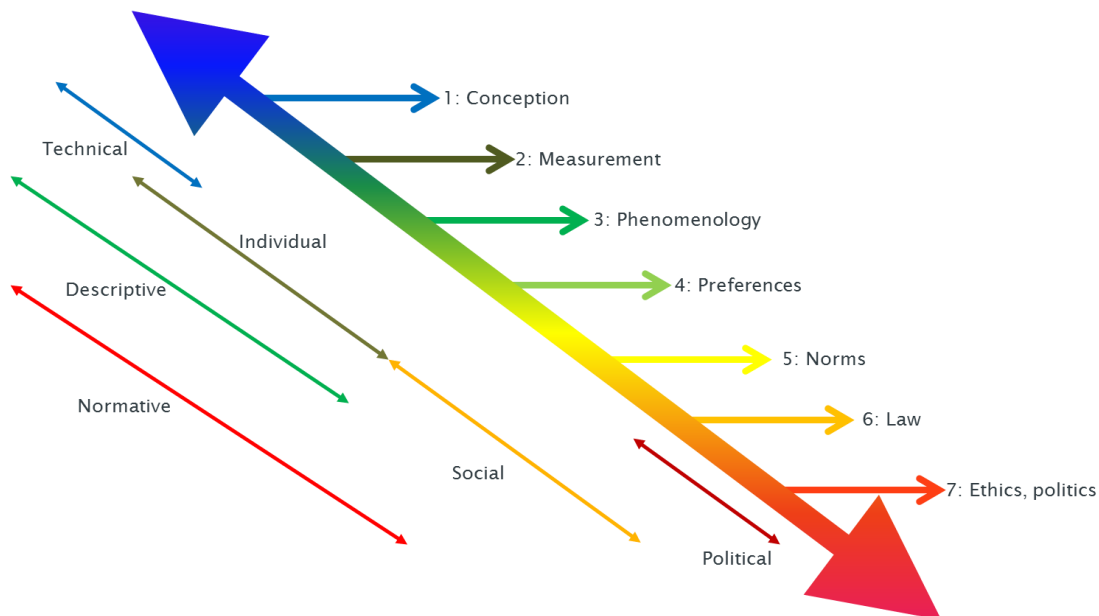


Figure 1: The seven levels of privacy discourse

So, for example, issues relevant to a discussion about (say) nude photographs of X placed online should be separated out as in Table 1.

| Level | Issues |
|-------------------|--|
| Definition | The conception of privacy (and breach) involved, bodily privacy, access to images of private parts. |
| Actuality | Are the photos behind a paywall? What format are they in? How many times have they been viewed or downloaded? |
| Phenomenology | Is X horrified? Excited? Unconcerned? Blissfully unaware that photos have been uploaded? |
| Preferences | Did X consent? Did X even put the pictures up there him/herself? Is this an example of revenge porn, strongly against X's wishes? |
| Norms | How are photos to be treated in this culture? The norm may be that one can only post such photos when one has the explicit consent of the subject. |
| Regulation | Who has the rights to the image? |
| Morality/politics | Should we outlaw pornography? Revenge porn? Is this a freedom of speech issue? Is there a public interest in publication? Is it bullying? |

Table 1: Different levels of a privacy discussion

We should not confuse these levels, for instance arguing that, because X consented (level 4) there is no breach of privacy (level 2). There *is* a breach of privacy, but in this case it happens to be one about which X is unconcerned.

The framework allows us to see how definitions of privacy can become overloaded with extraneous material from other levels. Recall Solove's criticisms of privacy theories above, which straddle the levels. Statements about when privacy is reasonable are properly level 5; those about cultural mechanisms for privacy are also level 5; issues about perception belong at level 3. None of these belongs at the definitional level. We also see that the work of Warren and Brandeis and Prosser, succeeds because they focus entirely at level 6, and pay no attention to definitional (level 1) questions. Agre's definition in terms of identity is largely placed at level 7, although the mention of unreasonableness is an intrusion from level 5 (there is no level 1 content in the definition).

Interaction between the levels

Clearly it is neither possible nor desirable to keep the levels separate. They will often overlap, or intrude upon each other. A privacy breach feels bad (level 3) because a norm has been broken (level 5). Streaking is exciting (level 3) because it breaks social norms (level 5). My preferences (level 4) may be unreasonable (level 5). Privacy law (level 6) seeks to deter or remedy harms (level 4). I want to share information (level 4) because my social network makes it feel good (level 3). Contextual integrity (level 5) is ethically important (level 7). Data sharing habits (level 5) are incompatible with democracy (level 7). Face recognition technology (level 2) makes my laptop seem much more secure (level 3). Privacy defined as a type of intimacy (level 1) is important for the autonomy of the individual (level 7). A new conception of privacy (level 1) will often emerge because of new harms (level 4), perceptions of harm (level 3), or changes to the privacy context (level 2).

Nevertheless, the discussions should remain distinct, we should be clear when we are operating at the different levels, and we should not draw too many conclusions across the levels. For instance, just because my preferences are unreasonable, that does not mean that they are not my preferences. EU law is meant to support social norms, but that does not mean that we can read off the norms from the legal texts. And if privacy law legitimises a set of practices, that does not mean that those practices do not breach privacy.

Conclusion

To conclude, the disarray to be found in privacy debate is often, if not always, the result of the confusion of levels, rather than from any problem with the concept itself. Privacy may be an essentially contested concept, and certainly looks like a family resemblance concept, but locating it within these philosophical classes does not resolve the tensions that Solove (2008) and Mulligan et al (2016) wish to relieve. The heterogeneity of views of privacy comes down to something more like the parallel coexistence of different conceptions of privacy in Rawls' sense (1971). These conceptions coexist at least partly because different definitions of privacy have been drawn up in the context of pre-existing views of its value (or otherwise), so that they typically have baked into them partisan views of the value that privacy provides. Substantive debates about privacy's value then easily degenerate into semantic disputes.

The varying interpretations of privacy can therefore be put down in part to essential issues with the multiplicity of its application, as many commentators have argued, but also in part to its perception as an appraisive concept that "signifies or accredits some kind of valued achievement" (Gallie 1956). This paper argues that the concept of privacy is less like the appraisive concepts to which Gallie refers (such as democracy or social justice), and rather more like a concept about which we generally agree, but *upon* which we may still fundamentally disagree (such as tax or the weather).

I have set out 7 separate but interacting levels of privacy discourse. I do not claim that these are the last word – for example, level 7 may be better broken up. However, they do separate out many of the key factors, and force us to recognise distinctions between psychology and sociology, the individual and the collective, conventions and constraints, the normative and the descriptive, and the personal and the political. Many problems with privacy definitions come down to illicit elision between such ideas. The framework I set out in this paper, alongside other valuable resources like Solove's taxonomy of harms (2008) and Mulligan et al's analytic framework (2016), will hopefully help prevent such confusing solecisms.

References

- Alessandro Acquisti, Laura Brandimarte & George Loewenstein (2015). 'Privacy and human behavior in the age of information', *Science*, 347(6221), 509-514, <https://doi.org.10.1126/science.aaa1465>.
- Philip E. Agre (1998). 'Introduction', in Philip E. Agre & Marc Rotenberg (eds.), *Technology and Privacy: the New Landscape*, Cambridge, MA: M.I.T. Press, 1-28.
- Irwin Altman (1975). *The Environment and Social Behavior: Privacy, Personal Space, Territory, Crowding*, Monterey: Brooks/Cole.

David Collier, Fernando Daniel Hidalgo & Andra Olivia Maciuceanu (2006). 'Essentially contested concepts: debates and applications', *Journal of Political Ideologies*, 11(3), 211-246.

Mark Elliot, Kieron O'Hara, Charles Raab, Christine M. O'Keeffe, Elaine Mackey, Chris Dibben, Heather Gowans, Kingsley Purdam & Karen McCullagh (2018). 'Functional anonymisation: personal data and the data environment', *Computer Law and Security Review*, 34(2), 204-221, <https://doi.org/10.1016/j.clsr.2018.02.001>.

Fadwa El Guindi (1999). *Veil: Modesty, Privacy and Resistance*, Oxford: Berg.

Amitai Etzioni (1999). *The Limits of Privacy*, New York: Basic Books.

Gottlob Frege (1980). 'On concept and object', in Peter Geach & Max Black (eds.), *Translations From the Philosophical Writings of Gottlob Frege*, 3rd edition, Oxford: Blackwell, 42-55.

W.B. Gallie (1956). 'Essentially contested concepts', *Proceedings of the Aristotelian Society*, 56, 167-198.

Woodrow Hartzog (2018). *Privacy's Blueprint: The Battle to Control the Design of New Technologies*, Cambridge MA: Harvard University Press.

Mireille Hildebrandt (2015). *Smart Technologies and the End(s) of Law*, Cheltenham: Edward Elgar.

John Hollander (2001). 'The language of privacy', *Social Research*, 68(1), 5-28.

Jeff Jarvis (2011). *Public Parts: How Sharing in the Digital Age Improves the Way We Work and Live*, New York: Simon & Schuster.

Bert-Jaap Koops, Bryce Clayton Newell, Tjerk Timan, Ivan Škorvánek, Tomislav Chokrevski, & Maša Galič (2017). 'A typology of privacy', *University of Pennsylvania Journal of International Law*, 38(2), 483-575.

Ponnurangam Kumaraguru & Lorrie Faith Cranor (2005). *Privacy Indexes: A Survey of Westin's Studies*, Institute for Software Research International, report CMU-ISRI-5-138, <http://repository.cmu.edu/cgi/viewcontent.cgi?article=1857&context=isr>.

Adam D. Moore (2003). 'Privacy: its meaning and value', *American Philosophical Quarterly*, 40(3), 215-227.

Deirdre K. Mulligan, Colin Koopman & Nick Doty (2016). 'Privacy is an essentially contested concept: a multi-dimensional analytic for mapping privacy', *Philosophical Transactions of the Royal Society A* 374:20160118.

Helen Nissenbaum (2010). *Privacy in Context: Technology, Policy and the Integrity of Social Life*, Stanford: Stanford University Press.

Kieron O'Hara (2016). 'The seven veils of privacy', *IEEE Internet Computing*, 20(2), 86-91, <https://doi.org/10.1109/MIC.2016.34>.

Richard B. Parker (1974). 'A definition of privacy', *Rutgers Law Review*, 27, 275-296.

Robert C. Post (2001). 'Three concepts of privacy', *Georgetown Law Journal*, 89, 2087-2098.

William L. Prosser (1960). 'Privacy', *California Law Review*, 48, 383-423.

John Rawls (1971). *A Theory of Justice*, Oxford: Oxford University Press.

Beate Rössler (2005). *The Value of Privacy*, Cambridge: Polity Press.

Bertrand Russell (1956). 'The philosophy of logical atomism', in R.C. Marsh (ed.), *Logic and Knowledge: Essays 1901-1950*, London: Allen & Unwin, 175-281.

Irina Shklovski, Scott D. Mainwaring, Halla Hrund Skúladóttir & Höskuldur Borgthorsson (2014). 'Leakiness and creepiness in app space: perceptions of privacy and mobile app use', in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '14)*, New York: ACM, 2347-2356, <https://doi.org/10.1145/2556288.2557421>.

Daniel J. Solove (2008). *Understanding Privacy*, Cambridge MA: Harvard University press.

Judith Jarvis Thomson (1975). 'The right to privacy', *Philosophy and Public Affairs*, 4(4), 295-314.

Samuel D. Warren & Louis D. Brandeis (1890). 'The right to privacy', *Harvard Law Review*, 4, 193-220.

Alan Westin (1967). *Privacy and Freedom*, New York: Simon & Schuster.

Ludwig Wittgenstein (1953). *Philosophical Investigations*, Oxford: Blackwells.

Ludwig Wittgenstein (1961). *Tractatus Logico-Philosophicus*, London: Routledge & Kegan Paul.