December 3, 2020

How participants see privacy

Nick Doty
UC Berkeley, School of Information

Status of This Document *This is a section of a chapter on findings from a published dissertation: Enacting Privacy in Internet Standards.*

How participants see privacy

In my conversations with standard-setting participants, I asked about their own views on privacy: how they considered it as part of their work, in their lives personally and for users of the Internet. Given the important decisions these engineers, protocol designers, lawyers and executives make regarding online privacy, it's important to understand their own perspectives. It also served as a useful introduction into more specific questions about Do Not Track or experiences with particular technical standards that might have privacy impacts.

As described previously,[1] privacy and security are values of special relevance to the Internet and the Web. But though they are distinctly relevant, security and especially privacy are still complex and contested ideas and their application to the Internet or to software engineering in general is not settled.

Here I will show the range of views of privacy as a concept to assist in understanding how those mental models affect decisions about privacy on the Internet. Next, I look at some common touchstones that drive motivating examples for technical standard-setting participants, including particular sensitive datatypes and their implications. One touchstone in particular – how one thinks about privacy for one's own children – helps illuminate the ways that participants think about privacy for others. Finally, tied to these different concepts of privacy and

[1] See Privacy and Security: Values for the Internet.

thinking about privacy for others, participants speak about the actual work of privacy in technical and legal settings.

Views on privacy differ

It's common to talk about privacy, a notoriously complex, challenging and challenged concept, through some narrower property, goal or sense. With the people I talked to, it often seemed that someone would start talking about privacy-as-something – and even for people who explicitly recognized variations in their own views on privacy and variation in the views others hold (see below), it would frequently be useful to talk about one particular sense or part of privacy at a time. Privacy-as-x can include a wide variety of concepts: the privacy analytic from Mulligan and Koopman identifies 14 distinct dimensions for classifying claims of privacy (Mulligan, Koopman, and Doty 2016), which I'll refer to regularly.

Privacy-as-confidentiality An easily accessible example is privacy-as-confidentiality: a sense of privacy as keeping certain facts secret from others. This is brought up occasionally by interviewees, often as a contrasting concept, to say that others used to see privacy in this narrow way, but they realize that it's more than that.

Here privacy is described as protection against the threat of violating confidentiality:

> privacy would be something related [to security] but a little bit different. It'd be more of like [...] trying to discover something that the user thought was hidden but is really not.

But confidentiality can also be invoked as a historical contrast:

> So 20 years ago when we thought about privacy it was really secrecy, right. It was, "Don't tell anybody anything about me"

Researchers have also noted this distinction, that privacy-as-confidentiality may have been an early attempt at privacy to engineering, as driven by applying cryptographic functionality from security engineering (Danezis and Gürses 2010).

Privacy-as-control More commonly accepted or invoked was some sense of privacy-as-control, what I would categorize as either informational self-determination (A. F. Westin 1967) or a sense of user understanding and capability to effect a choice about information. In the privacy analytic, control over personal information is the *object* of privacy, it's what privacy gives you.

From a very software engineering perspective, this gets described as permission, consent and control over software:

> when you look at a lot of web APIs you have to make sure that you're keeping the user first, and I think that's the mantra that we tend to talk about, making sure that the person using our products is in a position where they can make knowledgeable decisions about what they want to do with software.

> To me, the most useful [definition of privacy] is the right of an individual to control what happened to his own information. And it means that I may decide voluntarily to give certain information to another party. [. . .] I may decide voluntarily I'm willing to give it to this party in exchange for whatever benefits I derive from it. I tell Google my current location, it gives me a map of the area, the stores around me or something. I know why they want to benefit from it but I'm getting some benefit too, so I will do it.

From a business perspective, control may be closely tied to transparency and the availability of controls. We might see these either as distinctive views of the concept of privacy (e.g. that privacy is the ability to see what information is collected and what controls are available) or as combining both what privacy is and what functionality has to be implemented for the value to be maintained.

> Privacy for us was primarily the interaction with a consumer and how information was either collected or what controls were provided to them and what disclosures, transparency came along with that.

It's clear that the different models of privacy have overlaps and connections. Transparency frequently comes up in the context of privacy-as-control because how can someone meaningfully exercise control if they don't know what's happening?

> We're very clear about the information we get through [data source], what we do with it, what we don't do with it. We say what we don't do with it. And so, people can make that choice.

Privacy-as-protecting-data Related to privacy-as-control but with less focus on the user interaction are senses of privacy regarding data handling or how data about a person is used after it's collected; these might be defined as the *target* of privacy, the personal data that is being protected (Mulligan, Koopman, and Doty 2016). That could be as simple as not publishing log files or more complex enterprise privacy management systems.

> So usually, for me, privacy means you have personal data: it could be IP addresses, it could be email, it could be whatever. And privacy research is about how to best handle this data, protect the data, make sure that the data is used according to consent.

> I think there were certain types of activities that [Company] felt like it would like to be able to do if it took reasonable steps to protect data. And there were people at [Company], and many of them product or engineering people, who were very, very cautious. They would call up all the time, can I do this? I've put a flag on this. I have this data separate over here. There were a lot of people really taking care, but within the context of taking care and pseudonymizing data there were also tremendous business pressures.

These senses of privacy in data handling often have some sense of responsibility, stewardship or appropriateness about how data is stored or used. From a European regulation perspective, these concepts might be more familiar as much privacy-related law is specifically about data protection (typically, ensuring conditions for processing of personal data), rather than privacy rights.

Privacy-as-context-sensitivity While less commonly raised than these concepts of control or protecting data, other conceptions of privacy were significantly identified. Related to personal information but touching more on social norms and appropriateness would be respect for context. Most interviewees are not specifically referring to Nissenbaum's theory of contextual integrity (2004), but may be influenced by it; there are multiple sources of "context" as a source for privacy in engineering, including ubiquitous computing (Benthall, Gürses, and Nissenbaum 2017).

How one tech employee described context and appropriateness:

> Another aspect I think that we miss is that we have personae in real life, what I do at work and what I do in my hobbies and what I do

> at home or what I do in voluntary work and for other people, these are all distinct aspects of myself. If I volunteer, I'm going to pick an obnoxious example, at a clinic for sexually abused children, right, I do not need adverts about sexually abused children following me at work, and so this happy way that the online services just mashes all together, you know if my job at work is doing video quality assessment of online videos and some of them are pornography that doesn't mean I'm interested in pornography at home. So that's another aspect of online privacy that I think we completely missed, that is it appropriate now, is this contextually appropriate, and that's privacy again.

In this quote and a few other conversations, there's an identification that some information is specific to a particular situation or part of life and not appropriate to come up elsewhere, what we might call the collapsing of contexts as in danah boyd's work (2008).

Distinct but related are some ideas of privacy-as-relevance, that your privacy can be violated by "too much information" or information that you didn't want to come across about family or colleagues.

Privacy as freedom from intrusion While still related to information about people, privacy-as-relevance or privacy-as-context-appropriateness connect to the privacy concept as freedom from intrusion. To connect to the privacy literature, we would typically look further back, to Warren and Brandeis and 'being let alone' (1890), a particular interest in the late 19th century when photography and newspapers were technologies changing the basic assumptions about intrusions into our daily activities.

Two participants particularly highlight this idea of being left alone in the context of targeted advertising, expressing feelings of annoyance.

> And as long as they're left with the opinion that users don't care they'll do whatever flashy thing makes the most awesome user experience where, you know, you buy shoes online and they deliver special shoelaces to you in the next day because they think you're awesome and they think that you want that. Some people do, <laughs> you know. I don't. <laughs> I want you to just go away. I don't want to have any interaction with these people. I just want the thing that I ordered, you know.

> I used to think that personalized advertising would be an improvement over general advertising, but actually I find it hugely annoying and intrusive, and it's stupid in many cases, you know, I wasn't looking at that for myself, I was looking at that because my friend Nick was in my office and he said, "Maybe we could find a product online," I was looking for my son, I was looking for him, or I've already bought the damn thing, have you not noticed I've already bought the damn thing, and so the way it follows you around, it's sort of like having a terrier, it's constantly going, yap, yap, yap, behind you all the time, nipping at your heels, it's just infuriating.

Analogies to the Do Not Call program in the US have been familiar in Do Not Track discussions (where the Do Not Track name comes from), despite rather large differences in design and implementation. Also, Do Not Call is more narrowly targeted to privacy in the sense of intrusion (telemarketers ringing your landline during dinner), although intrusion (in addition to concepts of control over information) is also sometimes identified as relevant to online advertising.

Intrusion (for example, the physical intrusion of stalking) can also be a frustration to maintaining one's own autonomy, a value identified as protected by privacy.[2]

> And I kind of never felt that autonomy even as an adult because I was then growing up with the internet, and so as an example, I went to my boss and said [. . .] "Okay. You're putting our work schedules up where anybody can see them, and I have somebody showing up at my place of work before I get there," and it was "Well, too bad. We're not going to change what we do," and they were online, and that's just how it was.

Recognizing the variety of views of privacy The variety of senses of privacy that get discussed also reflect varying levels of concern among participants about their own privacy. And that variety among the population is something the participants themselves recognize about Internet users, which has important implications for the design of Web technologies.

[2] Again following the privacy analytic, autonomy may be a *justification* for privacy, a reason that privacy is needed.

For one baseline characteristic metric of privacy concern, I surveyed interviewees based on the Privacy Segmentation Index,[3] which divides people into the categories of privacy fundamentalists, privacy pragmatists and the privacy unconcerned. That index has been used to show general trends in the public: that many (and a growing number) are pragmatic about privacy, while people who are unconcerned about privacy shrinks as a proportion (perhaps because of increased awareness) and privacy fundamentalism is a growing minority (Kumaraguru and Cranor 2005). Among my interviewees, only a single one was classified as unconcerned, the majority (16) were pragmatists, and a substantial minority (8) were fundamentalists. (It often was not obvious to me, even among people I know professionally, what category an interviewee would fall into.) This generally reflects the trend in the public index, but our group of technical experts, privacy lawyers and advertising industry employees are especially unlikely to be unconcerned or unaware of privacy.

And while personal and professional perspectives on privacy certainly vary among my interviewees, participants also recognize or conclude explicitly that views of privacy differ among different professions, cultures, age groups and especially among the body of non-expert users of the Internet and the Web.

> I think most people are somewhere in the middle and they have, you know, different things that they post that want to go to different audiences or they want everything to go to a somewhat larger group of friends. But I don't think there's a one-size-fits-all. I mean, I think about it much more in terms of letting people understand what's going on, letting them make choices that are right for them, rather than us deciding, you know, everything has to be public or everything has to be secret.

That views of privacy differ substantially among users is one core reason to pursue user choice or preference expression mechanisms at all. Without such a difference, added infrastructure to enable choices and communicate preferences would be unnecessary intrusion: if tracking of online behavior is harmful or always unwanted, then blocking it is more efficient and beneficial than letting users make a choice about it; if tracking of online behavior isn't a genuine privacy concern, then letting users choose not to be tracked wouldn't provide any advantages. This conclusion is a key motivation behind Do Not Track and other expressive privacy

[3] Also described as the Core Privacy Orientation, see A. Westin (2001).

features: regarding the different paradigms possible,[4] DNT provides the end user with a variety of choices that are then communicated on to participating parties, rather than relying on a singular view of privacy interests.

A more complete list of the privacy-as- concepts identified in my corpus is included as an appendix.

Touchstones for privacy and impacts on others

How do participants in technical standard-setting talk about privacy in their own lives or in designing online services? Rather than falling back on abstract, philosophical language, it was very common for people I talked with to jump to particular motivating examples, whether it was specifics about their own life or hypotheticals. While the range of those touchstones was wide, particular topics were often repeated, especially sensitive datatypes (regarding location, sexual orientation or health) and family members, especially their own children.

Sensitive datatypes and salient privacy topics One direct way to conceive of privacy and explain its importance is to focus on the particular *target* of privacy, on what it is that we think privacy is meant to protect (Mulligan, Koopman, and Doty 2016). While it was common for participants in technical standard-setting to refer to views of privacy as control over information, they also identified the particular datatypes over which control were important, either to them or to the users they think about.

Several participants in IETF and W3C technical standard-setting referred to the privacy implications of geolocation functionality – that a device or online service can determine (with sometimes uncanny precision) where you're currently physically located. Location has particular salience for privacy because of three distinct properties of location data:[5]

1. it reveals other information (health conditions, employment, social connections, etc.) about people, based on where they go;

[4] See, previously, Do Not Track, a "handoff".

[5] Alas, as an impatient scholar, I've been presenting this three-part framework since 2010 without formally publishing it. See slides (Doty 2010) and related report (Doty, Mulligan, and Wilde 2010).

2. it's often uniquely identifying;
3. it facilitates physical intrusion.

One engineer discussing the Geolocation API directly touches on (at least) two of those factors:

> We don't want to give any information out that we don't absolutely have to. Location is a very sensitive one where if you travel from your house to work every single day, the service provider is gonna have a pretty good idea of where you live. In fact, if a service provider sees you going to a 7- Eleven instead of a Peet's Coffee they can make decisions about your lifestyle and what economic status you're at

However, part of why geolocation privacy in particular is raised so frequently when talking about technical standards is that at both IETF and W3C, defining APIs for communicating precise geolocation information was one of the first experiences with mobile device sensors, and the debate and architectural models would become the basis for many subsequent technologies (camera, microphone, light sensors, accelerometers, fingerprint readers, and on and on). Interactive user permissions on the Web started with Geolocation, and there were (relatively) heated debates over sticky policies (user's being able to specify machine-readable permission about use and retention) between IETF and W3C.

Other sensitive datatypes cited include health information, or particular categories of health that seem especially sensitive. As someone in the ad industry described it, advertising based on certain sensitive topics themselves seemed bad for societal outcomes:[6]

> it's a little problematic because there's no definition of "sensitive" [. . .] but what I was mostly concerned about, and it ties back to the other one about chilling effects – mental health, for example. Companies create very sensitive interest profiles on mental health in ways that I

[6] The particular limitation here is on the *use* of these sensitive categories rather than their *collection*, so it might be that the target of privacy is not specifically the data itself, but harms related to targeted messages about people within those sensitive categories. However, the sensitivity of use may also be related to potentially disclosing a sensitive health condition to others based on the presence of targeted advertising on that person's device, in which case we might say that the *target* is the personal data about health conditions and the *from-whom* is friends, family or people who might share a device with the individual. Having clear orthogonal dimensions can make it easier to tease out these differences.

> personally didn't think was a great thing for industry or society, and we decided that's sensitive, right?

While it's acknowledged that "sensitivity" of information is difficult to describe (perhaps in much the same ways that "privacy" is), a connection is made to chilling effects – that knowing that sensitive information about you is collected and used might discourage you from learning or discussing those topics that are sensitive to you. Sexual orientation was raised by multiple participants as a sensitive datatype regarding interpersonal relationships, but also in the context of a fear of inhibiting discussion or chilling young people from learning more about sexuality.

Privacy impacts for others Privacy is a sensitive, personal, subjective, contested value, which motivated my asking standard-setting participants – people who debate and design protocols that implicate online privacy for Internet users – for their personal views on privacy. But the participants in these interviews, and the participants in technical standard-bodies worldwide, and the employees of tech companies that build online services, are in many ways not similar to or representative of the population of users of the Internet. Based on demographic categories but also based on technical savvy or knowledge, the developers of Internet protocols and software are quite distinct from the median end user.[7] It is perhaps as important then to consider what designers think about privacy for other people as they think of privacy for themselves – the *subject* of privacy in the analytic mapping (Mulligan, Koopman, and Doty 2016). While I included a prompt in my interview guide to uncover ideas about user thoughts on privacy, it often came up unprompted, in three ways:

1. distinguishing that the speaker was not concerned about their own privacy, or that the speaker recognized they were more concerned about their own privacy than others might be;
2. noting the lack of understanding by users of the Web about how technologies that affect online privacy work; and,
3. identifying family members as a particular and compelling case of concern for the privacy of other people.

Why might these interviewees not be concerned about their own privacy despite their knowledge and work in a privacy-relevant field? For one, because the

[7] See Who participates and why it matters.

participants in the technical and legal fields tend to have many advantages and privileges of class, race, educational background and (relatively) stable governance.[8]

> That's just my personal interest. Because certainly those photos [of drinking in college] would have existed, and probably do exist in a Polaroid somewhere. But there's not a lot of downside there. I personally am not terribly worried about government data collection about me. I understand why people are. And I tend to be more trusting of certainly the U.S. government from – not because I think that they're adept at protecting privacy or data, I just don't think that they're nefarious, and I don't have much – I don't really have anything to hide. And so that doesn't really worry me. So I think if they can be subject to similar baseline requirements like data security, then it doesn't worry me that much.

This form of explanation – the lack of risk ("downside") and the lack of "anything to hide"[9] – emphasize how the lack of concern about personal privacy in these certain threat models is contingent, and the speaker repeatedly interleaves the explicit idea that these are personal calculations and will be different for others.

Some interviewees are also less concerned about keeping things private specifically because their own work is done in public or might involve some publicity. That can range from people who consider themselves public figures to engineers who just do more work online:

> But for example, so I'm in a gym and we have lots of events and so on there and when they send out emails, oftentimes they'll send an email and they'll have like a long CC list and I always react, "That's not really cool because some of these people might not have wanted their email address shared." I personally don't care, I mean, my email address is super easy to find and this is a pretty common way to react,

[8] These advantages and stabilities are described further in directly considering the ethical implications of "studying up" around this population.

[9] Writing on "nothing to hide" as a fallacious argument is widespread and I wouldn't be sure who to cite on the rhetorical topic. I don't take this individual's passing remark as an endorsement of a "nothing to hide" argument against privacy as a value of importance and I don't include it as a criticism of that perspective. Indeed, one of the primary reasons that nothing-to-hide is a poor argument against supporting privacy – that privacy is a value for protecting people who may have less power or protecting society so that people can take unpopular positions – is demonstrated by an individual distinguishing their personal fears from others'.

> I don't personally care about a lot of these things but I am very aware I think about when people's private information is shared.

Despite the relative privilege and advantages that people I spoke with share, some also identify themselves as in some cases likely to be more concerned about their own privacy than others.

> So part of why I don't use Facebook and Uber and LinkedIn is because of their track record with what they do with information, and there's a real cost to your life, right? I was in [City] on Monday, and it took me about a day to realize that's a town that no longer really has a functioning taxicab system. Apparently it was a weak system to begin with, and it got just decimated by Uber and Lyft, and it was so bad that I downloaded Lyft Monday night and used it to get around town on Tuesday. They were my first and second Lyft rides ever, and this is after three or four or five years of everyone in the world telling me that, "You can't function in human society without these apps." So that's one example [of things that might seem unnecessarily paranoid to others]. I mean, that's justified paranoia, but that's one example.
>
> [...]
>
> Well, two billion Facebook users can't be wrong, right? So I'm not trying to make any super-nuanced points about empirical research I've seen. I think I'm just reflecting on an increasing feeling that my choices are out of lockstep with just about everyone I know personally and also with what I read in the press about what the world is doing.

"paranoia" can be a term suggested to describe being outside of a community's social norms, rather than its formal denotation about irrationality. And in contrast to the privilege distinctions, these different evaluations of privacy can be among people who are similarly situated ("everyone I know personally").

One theme that gets at that kind of distinction – where others might not be concerned about themselves as subjects of privacy while others identify it as a concern – is user understanding, or more often the lack thereof, about Web technologies and their privacy implications. Some of these assessments are quite blunt:

> So I had done usability research, and I understood that people were by and large clueless about where their data was going

Users don't know what companies collect this information about them:

> I was fired up about it. I still am. The notion that a company I've never heard of has a list of websites I've gone to is not awesome, and I think folks – actually, I think there's plenty of science showing folks don't like it and would like to be able to limit it, and so I was concerned about it.

Users only understand when triggered by a particular event[10] and user attention and understanding are hard to persist over time:

> most of the time, of course, unless something happens like that, triggering, what the fuck, you know, how do you know that I know these 40 people, unless there's a triggering event like that, of course, most users don't notice, and it takes a disaster for them to notice, and, of course, we don't want to run the industry such that we run until we hit the iceberg and then we panic, we'd rather not hit the iceberg in the first place, thank you very much, but I have a fear that we're going to hit the iceberg.
>
> That's the problem with this stuff. Unless you are constantly reminded of it, you forget about it, right? That's the general mass of people on the web. They get pissed that Facebook put some complex thing to read that they know is not improving their privacy but taking it away. They get pissed for a day – whoosh – and they're right back in their normal life. They don't change. They don't jump out. The problem is the threat is not – what's the word – acute, right? It's gradual, and so it's going to get you later in life kind of like before. It's not something you react to in the present tense.

And that users' lack of understanding or awareness or ability to control may be an intentional design outcome:

> you can articulate that you care, but you have so much going on that it's really hard for you to take steps, which is why I would hope that

[10] There's a separate code in my dataset on "exogenous events" which I initially anticipated to be about Snowden revelations, which do come up in that sense, but Cambridge Analytica is also frequently cited.

> the government would address the more significant harms, because people can't possibly understand, and that's intentional. I mean, that is absolutely intentional. Industry doesn't want them to understand. It's confusing. [. . .] you understand, asymmetric information: you can't grasp it. Even as a parent now, I deal with parents all the time. [. . .] Parents have no concept of what's going on.

In these quotes, interviewees connect the lack of understanding – because it's "confusing" or "gradual" – to the lack of taking action to prevent privacy harms (in these cases, typically collection of information about them). There is an implicit response here to the well-known "privacy paradox" – if users are concerned about their own privacy, why don't they alter their behavior more often to better protect it? Experts identify a lack of understanding in users, which provides an explanation of the lack of action.

Privacy for one's children Most surprising for me[11] in these interviews was how frequently people I spoke with cited their own children in describing how they thought about privacy in their own lives. In part this may be paternalism in its original sense, that parents make decisions for their children because children may not have the awareness, understanding or knowledge to decide about information about themselves. But interviewees also recognize the lack of autonomy that children may have when parents are making choices on sharing information about them. It seemed that there was often more salience to the protection of children, the risks for their future lives and their ability to decide, than for the (privileged) parent themselves.

> personally I think I am always aware of privacy-related issues when using the Web, right, in different contexts. So, for instance, if I were going to share . . . I think everybody has rules about how they share data and how they share things on social networks, for instance. I don't generally share pictures of my kids or use their names when I'm writing stuff on Facebook, for instance. That is a personal sort of set of rules that I've hit upon. I know other people don't abide by those, and it kind of is a good example I think of how people have different views about privacy when they're using social applications in particular. To me the privacy issue isn't so much my privacy. It's

[11] Nota bene: I am not a parent.

> about that if I'm sharing information about my kids they're not old enough yet to be able to make that decision in an informed way, and I don't feel like I can make that decision for them, so therefore I'm not sharing information about them.

Children and family may also be cited as a contrast, where you might not care about limiting your own public image but of course wouldn't make the same decisions for children, again with the connection to making one's own decisions:

> By no means am I a private person. [. . .] However, of course, there are things that I don't wish to share with the world, or maybe I wish to share them certain audiences, but not others. My family is not as eager to be, you know, super-visible, so, I keep them from– I don't share a lot of pictures about my kids. My wife never wants to be shared or tagged. So, I look at a goal– and I don't look at privacy– and I argue that for most people privacy is not an absolute goal. We want autonomy. We want freedom to make decisions.

Or a contrast in terms of generations and how younger people might not appreciate the risks of sharing information:

> my personal view of privacy, it's gonna make me sound like an old man. I worry that people younger than me don't realize how dangerous putting something up on Instagram is or putting something up on Facebook is, and I think they're probably – in society there are probably lots of examples of, "oops, I shouldn't have shared that" and some of the ramifications.

Considering one's child's privacy can also have an impact on how one thinks about their own privacy, out of the same basic concept of protectiveness and importance:

> I just think, you know, it's probably social pressure. My wife puts pictures up of our kids, and so my kids are online, so why would I not put myself – I mean I'm certainly not as important as my children, right, so I think that may have had a large part to it.

There is a universal quality here about a parent's responsibility for, protection of and respecting the future choices of children.

The work of privacy

A distinct way to talk about privacy is to talk about the work that “doing privacy” consists in.

Privacy-as-compliance Many interviewees (especially lawyers and less often people in engineering) discussed privacy in their work as privacy-as-compliance: less about the value itself and more about ensuring compliance with a privacy law or with a set policy. Many privacy teams in tech companies report up to the general counsel rather than through the product part of the organization. Or the privacy team is the “keeper of the policy structure” including laws and other negotiated agreements. This can have a few distinct senses though (and interviewees will often refer to more than one): where the goal of privacy work is to comply with privacy regulation; where privacy work is about maintaining internal or external accountability that policies and practices are being upheld; or, where privacy work is managing risks, which could be security breaches, or more downstream, the unwanted news coverage or regulation that privacy issues could spur.

> So privacy is a great example where sovereign entities have laws and regulations in place on the topic, but those laws and regulations tend not to be written in such a way that they’re immediately obvious how one would implement those things. And frequently in order to verify whether or not people are meeting those regulations, there’s a desire to see certification in some form of compliance that might be ascribed to those behaviors. And so in order to do that you have to have some kinds of controls that you put in place, as well as criteria by which those controls would be executed. And so we focus on coming to international agreements on those topics relative to large-scale regulatory requirements, or to establish foundational concepts in emerging areas, where we know that that type of activity is likely to happen.
>
> All of those [businesses] have completely different perspectives on this concept of privacy. Some people think of it as compliance to a strict regulation, EU Safe Harbour or COPPA. Some people think of it as compliance to best practices [. . .].

"You have to sort of make it up as you go along" In addition to privacy work as a legal effort to maintain compliance with some external law or requirement, there is a distinct effort in the legal work of making internal, organizational policy to apply to some technical or business practice.

> there are areas of grey, right, where we don't have an established policy, we're looking at doing something new or novel, and therefore we can provide guidance to the organization to say this is our policy area, we don't have a policy, let's say, in your specific area, but here's where we would say the risk profile is for this particular area, and then we would give our recommendation on where they want to go. In those scenarios it's more of an assistive role in the organization.

And there is frustration with the de facto perspective of privacy-as-compliance in the professional sphere:

> [Privacy is] about the ethical and responsible use of data about people. I don't view my job as compliance, which is the problem with privacy today.

Some identify privacy as less focused on legal compliance and more on policy development, in contrast to other legal work:

> I just thought it was interesting and in flux, and it was clear that there weren't rules of the road yet in the US, so that's basically what I thought is that this is really interesting. And I was in a meeting where [other privacy lawyer] said a year or two ago, "So with GDPR, are we just going to become like other lawyers where we just follow the law?" And I was like, "Oh, my god. How boring would that be?" Privacy is not like that. <laughs> You have to sort of make it up as you go along – at least that's been the case in the past.

The "in flux" nature is attributed in part to the relevant youth of privacy in law and regulation, at least in the Internet context. But as a result, it makes the work of doing privacy as making it up, which might include lobbying, or interpreting new law, or arguing for policy approaches, as opposed to systems that just ensure compliance with more well-established regulation. That requirement for ongoing interpretation under broad or ambiguous regulation has been credited with empowering the field of privacy and bringing outside groups in to debate the privacy impacts of corporate actions (Bamberger and Mulligan 2015).

On the more technical side, there may also be a sense that the work of privacy is about discovery rather than simple implementation. Richmond Wong (2019) studies the field of human-computer interaction and explores how design practices can be used to explore, critique and present alternatives to privacy problems, in contrast to the perspective of privacy being a single fixed concept (like control) with design and engineering as putting that concept into practice.

Whether the work of privacy should be about contesting a particular concept of privacy is an open question. The argument that privacy is essentially-contested recommends that the "progressive competition" over the value is a beneficial feature that makes privacy more useful as a concept (Mulligan, Koopman, and Doty 2016). But it's notable that in some cases the value or purpose of privacy might be obscured in how it's discussed or considered.

The tension between whether privacy is settled elsewhere (like through formal regulations) and then implemented vs. being contested in the same place that it's being realized recalls the tension between separation and integration in how ethical concerns more generally should be a part of engineering practice.[12] It also connects to competing notions of organization-centered vs individual-centered views of multistakeholder process.[13]

What to conclude from these diverse views of privacy

Various privacy-as-control views are well-understood and common among this population of privacy experts and engineers developing technical standards that contribute to flows of information. That's no great surprise, but it should inform our understanding of the controls and mitigations that are likely to be considered in that setting. Different views of privacy, different threat models and concerns, may not get the same protection from additional transparency or data handling controls. How well will these views of privacy and corresponding expertise and developed tools and practices accommodate distinct privacy concerns: around fairness or online harassment, say? Or, as others have pointed out (Kostova, Gürses, and Troncoso 2020), how will views of privacy as control and control mechanisms work as software architectures change?

[12] See Separation vs. integration" in the earlier chapter on The Ethics of Engineering.

[13] See the section of this chapter on Individuals vs. organizations.

Recognizing different views of privacy means more than anticipating gaps during the engineering process. For compliance with privacy law and regulation, legal counsel are considering how to comply for Internet services that cross jurisdictional lines; for attracting customers from different countries and cultures, product designers are considering how to appeal to people with different cultural attitudes towards privacy. As privacy continues to be contested, there is an impulse to accommodate that ongoing debate with architectural designs that support public policy values without first settling all questions about their exact scope.

Understanding, effective capability and power are explicitly identified factors that respond to the motivating question about responsibility within the socio-technical system. Recall the vignette of "An ad that follows you"[14] where it isn't clear who is responsible or what you the user could do differently.

While tempting, we don't need to conclude that because experts in Internet protocols, online advertising and privacy draw a connection about the privacy interests of their children that privacy experts are advocating for a policy position of online paternalism. Nor should we conclude that paternalism is the proper or most effective approach we should pursue in looking at how to design for privacy among a non-representative group of end users.

Some conclusions we can draw from the significance of parenting as a theme, though. First, experts and designers of Internet protocols and online services may be attuned to thinking about the privacy interests of people different from themselves: many recognize the variation in preferences, levels of understanding and values about different conceptions of privacy. Second, there are mental models readily at hand for considering the impacts to people who are less expert or less capable of making their own decisions – people are familiar with the privacy of other people and people who can't decide for themselves from their intimate lived experience in raising children. In addition to exploring inclusive processes, participatory design approaches and user research grounding, we can also identify that thinking about the impact on differently-situated others is an existing practice in the technical field of Internet privacy.

Finally, competing views of privacy are complemented by competing views of privacy work. When privacy is enacted in developing technical standards, is that the work of debating the concept of privacy and the normative questions of what we should protect or how responsibility should be distributed? Or is the work a more technical matter of reifying policy that has been decided elsewhere into concrete form?

[14] See Do Not Track, a "handoff" in the earlier chapter on Privacy and Security.

References

Bamberger, Kenneth A., and Deirdre K. Mulligan. 2015. *Privacy on the Ground: Driving Corporate Behavior in the United States and Europe*. MIT Press.

Benthall, Sebastian, Seda Gürses, and Helen Nissenbaum. 2017. "Contextual Integrity Through the Lens of Computer Science." *Foundations and Trends in Privacy and Security* 2 (1).

boyd, danah. 2008. "Taken Out of Context: American Teen Sociality in Networked Publics." SSRN Scholarly Paper ID 1344756. Rochester, NY: Social Science Research Network. https://doi.org/10.2139/ssrn.1344756.

Danezis, George, and Seda Gürses. 2010. "A Critical Review of 10 Years of Privacy Technology." *Proceedings of Surveillance Cultures: A Global Surveillance Society*. http://www.researchgate.net/publication/228538295_A_critical_review_of_10_years_of_Privacy_Technology/.

Doty, Nick. 2010. "Geolocation, Privacy and the Web." UC Berkeley TRUST seminar, September. https://npdoty.name/slides/location-privacy-web2.pdf.

Doty, Nick, Deirdre K Mulligan, and Erik Wilde. 2010. "Privacy Issues of the W3c Geolocation API." *arXiv:1003.1775*, March. http://arxiv.org/abs/1003.1775.

Kostova, Blagovesta, Seda Gürses, and Carmela Troncoso. 2020. "Privacy Engineering Meets Software Engineering. On the Challenges of Engineering Privacy ByDesign." *arXiv:2007.08613 [Cs]*, July. http://arxiv.org/abs/2007.08613.

Kumaraguru, Ponnurangam, and Lorrie Faith Cranor. 2005. "Privacy Indexes: A Survey of Westin's Studies."

Mulligan, Deirdre K., Colin Koopman, and Nick Doty. 2016. "Privacy Is an Essentially Contested Concept: A Multi-Dimensional Analytic for Mapping Privacy." *Phil. Trans. R. Soc. A* 374 (2083): 20160118. https://doi.org/10.1098/rsta.2016.0118.

Nissenbaum, Helen. 2004. "Privacy as Contextual Integrity." *Washington Law Review* 79 (1): 101–39. http://heinonlinebackup.com/hol-cgi-bin/get_pdf.cgi?handle=hein.journals/washlr79&section=16.

Warren, Samuel D., and Louis D. Brandeis. 1890. "The Right to Privacy." *Harvard Law Review*, 193–220.

Westin, A. 2001. "Privacy on & Off the Internet: What Consumers Want." Technical report, Tech. Report for Privacy & American Business. Hackensack, NJ: Privacy & American Business.

Westin, A.F. 1967. *Privacy and Freedom*. New York: Atheneum.

Wong, Richmond Y., and Deirdre K. Mulligan. 2019. "Bringing Design to the Privacy Table: Broadening 'Design' in 'Privacy by Design' Through the Lens of

HCI.” In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, 1–17.