Encoding Privacy: Knowledge Production in Data Protection Compliance Work

Rohan Grover
Annenberg School for Communication and Journalism
University of Southern California
rohan.grover@usc.edu

GigaNet Annual Symposium 2022

Draft as of October 20, 2022. Do not circulate or cite without permission.

There is a gap between data privacy in law and in practice. On one hand, data protection regulations such as the GDPR and CCPA are seen as key strategies to protect individuals from potential harms from data collection and surveillance, but on the other hand prior research has found that user-facing features don't necessarily match expectations. The stakes for understanding and reconciling this gap are increasingly high as more jurisdictions around the world adopt data protection regulations, often modeled after the GDPR and CCPA, that make assumptions about the feasibility of compliance in order to turn privacy into an individual responsibility in the form of "privacy self-management" (Hull, 2015).

Rather than assuming that this gap can be closed with greater transparency or more rigorous enforcement, this paper asks: how do developers conceive of both their *responsibility* and their *ability* to achieve compliance with data protection regulations such as GDPR and CCPA? Based on interviews with 14 technical workers in different geopolitical and organizational contexts, this paper finds that developers are uniquely capable of shaping data protection work by creatively interpreting specific clauses in order to uphold "the spirit of the law"; yet they often perceive other actors—especially users—to be arbiters of compliance; all while expressing doubt that certitute in compliance is possible due to dependencies upon other actors and systems.

Thus, this paper argues that data protection compliance work should be seen as a form of knowledge production in order to shift analytic focus from the exclusive study of the social context of compliance to its content; that is, the knowability, facticity, and achievability of compliance. Three implications are highlighted: opening up comparisons between developer teams and scientific laboratories as sites of knowledge production; raising the urgency of decolonizing privacy studies; and calling for scrutiny of developers' social identities as opposed to presuming rationality in interpreting privacy expectations. These implications provide a framework for exploring data privacy compliance as a social process of knowledge production uniquely yet not fully available to a population of developers whose work has high stakes for the very concept of privacy itself.

Keywords: privacy, data protection, knowledge production, GDPR, CCPA

1. Introduction

In September 2019, a software development team for a user-facing app with approximately 500,000 daily active users worked with a product manager to implement new data protection and opt-out features to comply with the California Consumer Privacy Act (CCPA). Initially, they were excited to shape privacy safeguards for their users. However, there were many unresolved questions: What constitutes "selling" data? Should they roll this out to all users or only in California? What's the discovery strategy? Should they audit opt-outs to confirm compliance or just expect it to work? As they pored over the legal guidance from the privacy lawyers, they grew increasingly frustrated by shifting guidance, poor communication, and insufficient resources to instrument the features that they thought were necessary.

Soon, it became clear that the paramount objective was to implement something—*anything*, really—within just a few weeks, before the law would take effect in early 2020, to show that the company made a good faith effort at compliance. Under pressure and frustrated by poor communication and corporate bureaucracy, the development team revised their scope, sacrificing features and usability to meet their deadline. What they launched looked very different from what they had planned just a few weeks earlier, and they had low confidence that they were actually compliant with CCPA.

This uncertainty of the status of data privacy compliance is not unique to that development team. In April 2022, a leaked internal document from Facebook stated:

> We do not have an adequate level of control and explainability over how our systems use data, and thus we can't confidently make controlled policy changes or external commitments such as 'we will not use X data for Y purpose.' And yet, this is exactly what regulators expect us to do… (Franceschi-Bicchierai, 2022)

According to Vice's reporting, a former employee added:

> Facebook has a general idea of how many bits of data are stored in its data centers. The where [the data] goes part is, broadly speaking, a complete shitshow… It gives them the excuse for keeping that much private data simply because at their scale and with their business model and infrastructure design they can plausibly claim that they don't know what they have. (Franceschi-Bicchierai, 2022)

Prioritizing deployment over quality or skipping documentation is routine in software development, but the quality and stakes for data protection work are different. There appears to be an *inherent indeterminacy* that interacts with, and often impedes, compliance, and this matters because of the accelerating expansion of regulatory action around the globe in recent years. Thus, this study asks: How do developers conceive of their responsibility and their ability to achieve compliance with data protection regulations such as GDPR and CCPA? It pursues this question by building on previous research that approaches developer practices in general, and privacy work in particular, as a social process informed by developers' attitudes and experiences.

Prior research has demonstrated that there is a gap between privacy in law and in practice. On one hand, data protection regulations such as the GDPR and CCPA are seen as key strategies to protect individuals from the harms of digital data collection, including violating an expectation of privacy, collecting and selling personal information without consent, and threatening rights to free speech and association. However, a plethora of research (largely by HCI and computer science researchers) has shown that the GDPR and CCPA fail to deliver the

data protections they promise because they are subject to what Waldman (2019) calls *legal endogeneity*: symbols of compliance—such as trainings, audits, and documentation—standing in for real privacy protections. This argument has been supported by previous findings that user-facing configuration options don't necessarily match data protection practices implemented by companies. The stakes for understanding and reconciling this gap are increasingly high as more countries and jurisdictions around the world adopt data protection regulations, often modeled after the GDPR and CCPA.

Thus, privacy laws cannot be understood only at the moment when they reach the public (i.e., by studying interfaces and technical configuration options) or by examining their intentions alone (i.e., by studying policy discourse). Instead, they also need to be examined at the site of their enactment: in software developer teams who translate policy discourse to code. Therefore, this paper builds on prior research about privacy work in software development by synthesizing themes from 14 interviews with developers and adjacent technical workers, such as data, product, design, and operations. These findings are organized around four guiding questions: How do developers approach compliance work? How is compliance work situated within organizations? How do developers make decisions about compliance work? What do developers think about the lasting effects of GDPR and CCPA?

Based on these findings, this paper argues that data protection compliance work can be understood as a form of knowledge production: both the public and policymakers are unaware of the exact ways in which regulations are manifested in user-facing interfaces, so developers can be seen as experts who create and hold exclusive knowledge about what was actually enacted on a particular website or application. This conceptualization of *privacy knowledge* enables analytic distinction between the *context* in which privacy is constructed and the very *content* of privacy knowledge. In other words, it recasts privacy as not only a principle but a specific kind of expertise about the actual relationship between privacy expectations and privacy outcomes in a particular application. For example, certain developer teams are exclusively aware of the specific categories of data subject to user configuration in a given app, the advertisers and partners from whom that data is withheld, and the level of confidence, scrutiny, and due diligence applied to various components of data protection regulations. Once those decisions are enacted in code, the *actual* relationship between decisions and code is a category of knowledge exclusively available to developers with a unique ability to ascertain—or construct—it.

In summary, this paper argues that software developers should be understood not only as the "human factor" in sociotechnical systems, but instead as actors who hold and produce expert knowledge in the privacy domain. Understanding this role is important because privacy knowledge is by nature inaccessible to both policymakers and the public, while also encoding privacy norms for users. This paper highlights three implications: it opens up comparisons between developer teams and scientific laboratories as sites of knowledge production; raises the urgency of decolonizing privacy studies; and calls for scrutiny of developers' social identities as opposed to presuming rationality in interpreting privacy expectations. These implications provide a framework for exploring privacy as not only a socially constructed principle but also a form of knowledge uniquely available to a population of developers that lacks the structure of a scientific community despite being conferred a significant role in enacting privacy self-management as a technique of subjectification.

2. Studying Privacy in Software Development

This study builds on previous research that has investigated how developers approach privacy. This literature begins with the premise that privacy values can be embedded in technical design and development, a framework called Privacy by Design (Cavoukian, 2009). Inspired by this framework, many researchers have explored the decisions developers make when designing and development software, especially when implementing privacy features.

One challenge with Privacy by Design is "translating the general abstract notion and the meaning of informational privacy (or, in its European term, data protection) into concrete guidelines" (Hadar et al., 2018, p. 260). This has opened up a line of research about how developers approach privacy work in software development because they are delegated responsibility to interpret and design privacy frameworks (Hadar et al., 2018; Tahaei & Vaniea, 2019). In particular, Greene and Shilton (2018) have argued that Privacy by Design "positions developers and mobile companies as ethical agents" and then, accordingly, "platforms emerge as de facto regulators" (p. 1643). This effectively promotes platforms as well, such as iOS and Android, as not only sites of privacy mediation but also privacy co-regulators themselves, especially because iOS and Android developers maintain different conceptualizations of privacy.

Therefore, several studies have explored how developers conceptualize and approach their relationship to privacy work. For example, Tahaei, Jenkins, et al. (2021) interviewed computer science students and found similar attitudes toward privacy and security as professional developers, including "hacker and attack mindsets", inexperience with security instrumentation, and a tendency to trust other developers' solutions without applying scrutiny. Therefore, many developers turn to online forums for advice with approaching privacy work. For example, Shilton and Greene (2019) found that forums are important spaces where developers deliberate ethics and values, and they characterized the stories and justifications developers evoke to explain their opinions. Additionally, Tahaei et al. (2020) analyzed privacy-related questions on Stack Overflow using both topic modeling and qualitative analysis, and found that developers seeking answers to questions about privacy derived their answers from platforms (i.e., Apple and Google) or individual opinions from other uncredentialed developers. However, Tahaei et al. (2022) have found that responses to questions on online forums can be biased toward particular privacy paradigms.

Collectively, these studies indicate that developers are generally uncertain about how to tackle challenging privacy questions, and that they rely either on online forums or platforms to define privacy priorities and standards. However, sometimes exceptional developers who are committed to privacy lead team decisions and change their team culture. Tahaei, Frik, et al. (2021) have characterized those individuals as "privacy champions"—defined as those who strongly care about advocating for privacy on their teams—and found that they play important roles in cultivating a team culture that respects privacy in software development. They further found that privacy champions attempt to overcome prioritization conflicts, organizational ambiguity, and limited support through team-building practices such as informal discussions, promoting stakeholder communication, and developing documentation.

On the other hand, many developers' privacy decisions are subject to external influence. Developers tend to retain the default privacy settings on third-party services such as advertising networks (Mhaidli et al., 2019); meanwhile, platforms claim that developers are responsible for their own regulatory compliance—even though their configuration interfaces contain dark patterns that nudge developers toward "privacy-unfriendly defaults" (Tahaei & Vaniea, 2021). However, developers are often not aware of the data collected by third-party tools due to a lack

of resources, expertise, and time (Balebako et al., 2014). These findings are consequential for how and when developers think about privacy. For example, Li et al. (2018) found that although developers often claim to care about privacy, they carry a limited understanding of their own products' data collection practices, partly because of poor documentation. Instead, developers may be more likely to discuss privacy concerns in response to changing policies from platforms, app stories or regulators rather than features in their own products (Li et al., 2020). This may be because developers see a clear separation between their own code and third-party services, who they consider responsible for their own privacy practices (Mhaidli et al., 2019), or because developers feel that privacy work is more likely to impose costs on their own time rather than benefits (Li et al., 2020).

Understanding how developers approach privacy work is important because their decisions can carry material consequences for user privacy, including when implementing GDPR and CCPA compliance. Feng et al. (2021) have argued that compliant data protection options often do not help users make informed decisions, and thus argue that compliance work should be enhanced by providing users with "meaningful privacy choices" that exceed standards of usability. Indeed, several studies have found that the types of privacy choices mandated by these regulations are often placed in different locations on websites, some of which are difficult to find (Habib et al., 2019; Habib et al., 2020), and that design choices such as link text and choices between banners and overlays can affect user comprehension of and decisions about privacy choices (Degeling et al., 2019; O'Connor et al., 2021). In fact, a user study about icon design led to regulatory changes to CCPA (Habib et al., 2021). These criticisms apply to data subject access requests, as well. Although developers are not always responsible for fulfilling requests, they are often implicated in setting up workflows to verify and execute access requests, and Di Martino et al. (2021) have conducted audits that exposed how many large websites leaked personal information to unauthorized third parties after only minimal social engineering.

Overall, while previous studies have explored how developers make privacy decisions in general, and how their technical and design decisions shape user outcomes in particular, they have not defined the specific responsibilities assigned to developers to comply with data protection regulations or how they are situated within their organizational contexts. These dynamics are important because they shape users' experiences of privacy online—in essence, to what extent the regulations fulfill their intended goals—which is increasingly important as new data protection regulations have been passed or introduced around the world, often modeled after GDPR. This study fills that gap by evaluating how communicative and team-based processes identified in the literature interact with GDPR and CCPA—and what are the implications of developers' decisions on what privacy regulations actually look like in practice?

3. Methods

This study focuses on developers because the literature has already demonstrated that they play a particularly central role in between many stakeholders: they operationalize regulators' statutes, remain accountable to privacy lawyers, learn from and negotiate with platforms, and produce features for end-users (Greene & Shilton, 2018). However, this study also explores developers' experiences with data protection regulation compliance as collaborative processes with diverse organizational actors, including lawyers, executives, and other technical workers. Thus, it departs from previous research, which has either focused on developers alone (and excluded adjacent functional areas) or the outputs of their work (i.e., code,

features, and frameworks), by including workers from data, product, design, and operations teams.

The study was conducted through semi-structured interviews with 14 technical workers who have been responsible for GDPR and/or CCPA implementation in a variety of organizational environments. Technical workers are defined as individuals in engineering, product, design, user research, and compliance/operations departments. The key inclusion criterion was that participants should self-identify as having experience bearing some responsibility for an organization's compliance with GDPR or CCPA.

3.1 Recruitment

This research plan was designed to accommodate recruitment challenges documented by previous research. For example, Tahaei and Vaniea (2022) evaluated multiple channels, such as Prolific and MTurk, for recruiting probability samples for developer studies. They compared each channel's cost, quality, programming skills, and demographics, and found that recruiting from their university's listserv of computer science students returned the highest quality sample. However, recruiting participants with privacy and security experience and expertise adds an additional challenge that cannot be easily overcome with crowdsourcing platforms or by drawing from student listservs (Patnaik et al., 2022). Indeed, most studies about developers' privacy attitudes and practices are based on small samples of 10–20 interviews or 50–200 survey responses recruited from small sample pools.

Therefore, for this study participants were recruited using purposive sampling by drawing from the author's professional network. The author posted calls for participation on LinkedIn and Facebook, which yielded three interviews, and reached out to 26 individuals, which yielded 11 additional interviews. Individuals were targeted based on diverse organizational experiences, including Big Tech companies, startups, and non-profit organizations, and different global regions.

Purposive sampling was conducted in this way for two reasons. First, analytically, maximum variation purposive sampling helps discern more compelling common themes and differences across experiences. Second, purposive sampling was necessary to overcome the challenges previously described. Identifying individuals who have been responsible for privacy regulation compliance within a company is challenging from an external vantage point. Compliance instrumentation is often a single project—or even a series of tasks within a sprint—rather than a full job description for most developers. Therefore asking a potential participant if they have worked on privacy regulation compliance is open to interpretation and subject to recall from their project history from two (CCPA) or four (GDPR) years prior. Even if a participant is eligible for the study, they may be reluctant to discuss sensitive privacy regulation compliance issues on the record. Participants voiced this concern for two reasons: privacy regulations are considered sensitive topics given individuals' inability to speak for an entire company while potential complaints carry high fines, and because of the context of high-profile leaks from Facebook and ongoing controversy from companies like Twitter and Amazon. Therefore potential participants may harbor fear of retribution or may be bound by nondisclosure agreements in an industry marked by dramatic asymmetry between workers and owners.

3.2 Interviews

The fourteen interviews were scheduled for 60 minutes, but lasted between 30 and 90 minutes. They were conducted between October 2021 and August 2022 on Zoom, except for one

interview that was conducted in-person. Participants were offered an e-gift card worth USD 25 (or the equivalent in their local currency) regardless of their willingness or ability to answer every question. Several participants declined the gift card, most often because of restrictions imposed by their employer; in one case, a participant asked for their payment to be donated to a privacy-oriented advocacy organization.

Participant demographics are summarized in Table 1 and Table 2. The over-representation of developers and of North America-based participants reflects two factors: the author's professional network and experiences and participants' experiences being responsible for GDPR/CCPA compliance. Eleven out of 14 interviewees requested anonymity both individually and for their organizations. Therefore individual companies and countries are not identified, exact positions are not disclosed, and the gender-neutral "they" is used for all participants.

Table 1. Interview Participants by Functional Area

Engineering	8
Data	2
Operations/Compliance	2
Design	1
Product	1

Table 2. Interview Participants by Region

North America	9
Oceania	2
South Asia	2
Central and Eastern Europe	1

Eleven of the 14 interviews were recorded and transcribed, yielding approximately 75,000 words of transcription. The transcripts were analyzed using iterative open coding to identify common themes, exceptional experiences, and representative quotes. The other three interviews were not recorded or transcribed as requested by the participants. Instead, the author took non-identifiable notes during the interviews and compared them to the other interviews to synthesize broad findings.

The interviews were guided by four key questions, which are used to structure the results below. First, how do developers approach GDPR/CCPA compliance work? Second, how were developers' responsibilities situated within their organizational contexts? Third, what decisions did they make while doing compliance work? Fourth, what are the lasting effects and attitudes about data protection regulations today?

4. Findings

This section summarizes the findings from 14 interviews with developers and adjacent technical workers according to the four guiding questions described above.

4.1 How Did Developers Approach Compliance Work?

These questions were oriented toward discerning participants' attitudes toward data protection regulations and understanding of the scope of work.

4.1.1 Feeling Unprepared to Address Vague Regulations

The participants often expressed hesitation to describe the scope of GDPR and CCPA in detail. This was even true of several participants who were individually responsible for overseeing compliance work for a particular product or application. One participant, who was the lead engineer for a mobile app in a very large technology corporation, described GDPR compliance work in 2018 as such:

> We were in a mad dash to ultimately be able to say we're in compliance. That was, and is in all my experiences, the principal concern—being able to say we are compliant. I don't know that anybody I've encountered—and certainly not I—has a concrete grasp of exactly what that means.

Another participant, who was solely responsible for GDPR compliance at a small 40-person company, provided a more blunt response to a question about how prepared they felt to tackle GDPR: "Zero. I felt unprepared."

These two examples are representative of the experiences described by participants across different organizational contexts, which are also captured by the metaphors participants used to describe their experiences. One participant compared data protection compliance work to fulfilling accessibility requirements but with "less clear guidelines" in which "a lot of people and a lot of companies don't want to implement them to the point they are recommended, or at all, until it becomes a legal issue". Another participant described data protection compliance work as being "in the same bucket" as the US Health Insurance Portability and Accountability Act (HIPAA), but "less serious". The perception that data protection regulations are "less serious" than HIPAA was apparent in how participants reconciled their approaches to achieving compliance. Several participants used expressions such as "the spirit of GDPR/CCPA" to describe their aspirational level of compliance rather than trying to satisfy every component of the regulations.

4.1.2 Compliance Work as Risk Assessment Labor

The two exceptions among the participants—in that they felt well prepared for regulatory compliance work—were both consultants supporting compliance efforts at very large technology companies. One of these consultants "fell into" privacy work after being pulled into a project to respond to a data breach, which provided the necessary experience to become familiar with the technical and legal nuances of compliance work: "The best preparation happens on the job. I didn't plan to work in privacy. The breach happened, and then I got into it." The second consultant described how they drew upon their background in content moderation to make decisions about verifying and fulfilling data subject access requests:

> In the end, it's all about risk assessment. Is giving this person their data, or removing an account, based on the information they've given us more likely to result in the most correct outcome? Or is it more likely to result in some random person getting someone's data? It's the same as trying to trust someone to accept the rules of whatever platform they're on and not harass someone in the future. You're just basically just assessing what you have in front of you.

4.2 How is Compliance Work Situated Within Organizations?

4.2.1 High Autonomy

Participants on engineering teams described high levels of autonomy with minimal oversight over their work. In several instances, engineers were the first to bring GDPR or CCPA compliance requirements to their organizations. A mobile engineer described that their autonomy came from a combination of pressure from engineering leadership and lack of product strategy:

> We had a lot of autonomy in general, and in particular with GDPR and CCPA because they both came when we didn't have strong guidance on the mobile side. So this was a particular thing that engineering took upon itself to try and handle. We were responding to middle manager engineering pressure which is like, we have this SDK, this is our solution for all the products [across the company], this is what you're doing. And the sentiment was that Product didn't have that much to say about it—they'll want to inform where we put the menu in the app, and that's fine. But the underlying mechanics were fairly technical in nature… So it was very much Eng-owned, and it was a hodgepodge between direction that we were getting from senior Engineering leadership in the organization, and our own judgment and instincts about what the intention of the laws were.

This description is similar to the experiences articulated by other engineers, including those in smaller organizations that received pressure from executive leadership instead of middle manager engineers. In general, they described being saddled with, or perhaps entrusted with, interpreting compliance work because of their technical expertise—even if they did not feel equipped to do so.

For example, another participant described their experience instrumenting GDPR compliance as one of two contractors working for a large international non-profit organization without any full-time technical staff. They describe how they were saddled with the responsibility of ensuring that the organization was prepared for the beginning of the enforcement period:

> We were not really given clarity on what compliance was. We had a [marketing director]—she mostly just passed it off to [us] and had us kind of just figure it out… It involved Googling and seeing how to be compliant, and using online resources and not necessarily getting legal interpretations of how to do it.

4.2.2 Lack of Accountability

Since the participant was one of two contractors who worked on separate projects, their work was not code reviewed by another engineer, nor was their work fully tested for quality assurance. Instead, they were independently responsible for both executing and checking their own work—all while acting as a contractor. The participant ultimately ended their contract before the GDPR compliance work was complete, passing it back to the marketing director who hopefully found an alternate contractor to pick up the remaining work.

Lack of accountability was another theme from participants' experiences across diverse organizational contexts. A developer at a large corporation described how "it was ultimately up to us, the implementing team, to look at what was in the application, document all the [third-party services], and identify which ones were potentially at risk [for collecting personal data]". Notably, the process of researching third-party services was not verified by executive engineering leadership, compliance officers, or lawyers:

> One of the main things that I remember is a statement that basically 'information that you collect that applies directly to the operation and function of your application isn't really subject to the same restrictions'... We creatively applied that in a few places [where we couldn't verify the data collected by a third-party service] where it was like, okay, I *think* this is operational data. You could probably argue it in a different direction, but we don't have an easy workaround here, and we don't think it's a *flagrant* violation. So we're going to go with it. We're just going to say this is operationally relevant data that we're collecting, it's not being used for tracking or profiling. We're just going to ignore it.

In sum, as another participant described, "we could have easily flaunted it entirely and nobody would have known; nobody was in a position to question us." The exceptional level of autonomy was corroborated by a designer, a product manager, and a compliance consultant in different contexts. Each of them described providing suggestions or minimal requirements to engineers, but generally allowing engineers to lead in developing a compliance strategy. The compliance consultant described one instance where they were accosted by an engineer: "There was a person who screamed at me. They were like, why are you wasting my time? I said I'm not wasting your time, I'm just telling you what you need to do." Ultimately, the participant did not have the authority to block the engineer from launching their non-compliant feature.

4.2.3 Tenuous Relationships with Lawyers

Participants' high levels of autonomy and minimal accountability extended to their relationships with lawyers, albeit for different reasons. On one hand, a senior engineer at a large technology conglomerate described their team's reluctance to consult lawyers to reconcile questions about whether a particular third-party service was compliant:

> Nobody up or down the chain knew exactly what the answers were. But there was generally a hesitation to take things to lawyers just because it usually ended up being more work. We would often get the outcome that we might have suspected, but would prefer to avoid it in terms of the effort required or the implication.

On the other hand, another engineer at a startup had access to a lawyer and a full-time compliance officer to support GDPR and CCPA compliance, but they appeared to be more interested in upholding the "spirit of the law" in general simply to avoid a large penalty:

> It's like, just make a good faith effort and that's good enough. That's kind of our legal take on it… There's a lot of gray area on what is the right thing to do. But I think the lawyer probably wrote one paragraph about it and then stopped caring about it. Because there's very little chance we're going to get sued or anything, and that's what the lawyer cares about.

In several other cases, participants described how legal expertise was unavailable to them, either by design or because of lack of capacity. One participant described their experience

working with a small organization that was reluctant to consult any lawyers throughout the process:

> The struggle was that nobody wanted to talk to a lawyer to give us any understanding of the things we should do… It was something that, uh, felt like it was actively being avoided… So I don't know how close to compliance we were because I'm not a privacy expert and that was never what I was hired for.

4.3 How Did Developers Make Decisions About Compliance Work?

4.3.1 "Just a One-Time Thing"

The specific components of the regulations that participants focused on were implementing cookie consent notices, adding additional selections and disclosures to forms, updating privacy policies, purging data, reconfiguring third-party application integrations, and updating data infrastructure. However, one of the most common themes across all interviews was that participants described a lack of auditing or follow-up to ensure high quality compliance. On one hand, several participants described facing major bugs while instrumenting initial compliance. For example, one participant described how EU visitor traffic doubled immediately after deploying a redirect for GDPR because new cookies were being created for existing users. However, the issue was not caught immediately because a separate bug affecting a third-party analytics service was introduced at the same time, nearly canceling the inflated EU traffic.

Another participant also described a bug with cookie opt-outs during CCPA implementation that was caused by testing for one case—users opting out—but not for another—users *not* opting out:

> Our original code for [our third-party tracking service] wrapper was successfully turning off the cookies for people who didn't want them, but also had a bug where it was successfully turning them off for everyone else, too. We went like a week without collecting data before anyone noticed this. What happened was somebody wrote the code, somebody else ran it and opted out, and then looked at [the third-party tracking service] and confirmed that the correct data was there. That was our testing process. If that stopped working now, it would probably be months before we noticed it. We don't have any sort of regular audit.

When asked about the current status of GDPR and CCPA compliance, participants universally admitted that they had not reviewed or audited their implementation work, including confirming that third-party cookies were caught behind a wrapper, confirming that the appropriate disclosures and configurations were available on forms, and so forth. One participant suggested that their startup's privacy lawyer would likely not agree to a proactive compliance audit because it was not mandated and thus they could remain ignorant to potential bugs or violations while staying true to "the spirit of the law".

The mindset that regulatory compliance work was "just a one-time thing" applied to handling data subject access requests at small and medium-sized organizations. While some participants used self-service solutions such as privacy dashboards to enable users to request or delete their personal data, other participants set up manual processes that they planned to automate in the future in order to handle higher volumes of requests. However, none of the participants improved their data subject access request workflows after setting up their initial workflows. Participants generally attributed this to the low volume of requests:

> We wrote a "hacky" script that an admin runs with the intention that, if we get a lot of these requests, we'll make this a self-service tool so that we don't have to spend a lot of time doing this and so that it's easier for customers. But I think we've only had three people ever ask for their data. And I want to say maybe once a month somebody asks for their data to be deleted.

4.3.2 Limited Sense of Responsibility

In addition to expressing reservations about the status of their compliance instrumentation, several participants expressed that they were inevitably reliant upon several other services to maintain their compliance mechanisms. One participant described their GDPR and CCPA compliance work in a mobile app as flipping a series of switches without any visibility into what *actually* happens to user data stored on other companies' servers:

> I do not actually know—and I don't think anybody does—that [the service stops collecting data] when the user flips the switch. I'm the person building it—I'm responsible for making sure the switch is there, and that it calls the service somewhere when the user taps it. Beyond that, I have no idea what happens. I'm entrusting that what is supposed to happen, does. Somebody else is doing the same thing I am up the chain, probably, and maybe 100 more people up the chain after that are all passing it along, and hopefully the person in the place where that record exists—which is many places and many people—are treating it with due respect and honoring it the way they're supposed to. So everybody on my team and my peers and colleagues, we are doing what we believe needs to be done, by law, by intention, by principle, but without actually concrete assurance… We're sending them something and getting a response. That's the extent of our interaction. What happens on their systems is completely opaque to us.

4.4 What Do Developers Think About the Lasting Effects of GDPR and CCPA?

4.4.1 Strategic Value of Regulatory Ambiguity

Several participants described how they have been able to advance personal priorities under the umbrella of data protection regulation compliance. For example, a senior engineer at a nonprofit organization, who noted their background in the free software movement, noted that they found strategic value in drawing on GDPR to advance an anti-corporate ethos. They described how GDPR provided cover for decisions to, for example, avoid Big Tech companies.

> I would never be okay with uploading emails to Facebook to create targeted lists, because I imagine that Facebook is holding the data… We didn't really use Google Analytics for the same purpose. You put their tracker on your website and you're inviting the giants to your website and to all your users… These things influenced my decisions [about privacy]. For instance, I would prioritize not giving data to huge corporations. But I wouldn't—and I didn't—prioritize creating an automatic way for people to access their data. I would say, they can email us and ask to deliver their data, or change it, or delete it, and that's fine. But I wouldn't put so much priority on implementing something automatic to make it easier and more accessible. That's where there is imbalance. Also, when it comes to disclosures, GDPR would say, okay, you have to think about all these points, and there's this really long text. And I know a lot of organizations, including ours, were struggling to make this text a bit shorter, because it takes up a lot of space. But I wouldn't be super strict with this. I would think, okay, we are already trying to do a very good job [with privacy] here. I don't know if this text is so useful to these people… Privacy comes more from sentiment and the experience of being a member—rather than trying to follow the regulation to the letter.

In this example, the developer made technical decisions, such as avoiding Big Tech integrations, under the banner of advancing privacy—even though those services were not actually disallowed by GDPR. Instead, they described strategically interpreting the spirit of the regulation to advance a privacy framework driven by their personal politics. Later in the interview, the participant acknowledged that their values were not shared by the organization, and that once they left for another employer, the organization began employing Big Tech services again.

While this may seem like an extreme example, another participant described a similar appreciation for the regulations providing cover for advancing privacy as a human right:

> [GDPR and CCPA] are not the best thing for our business, but it is important to our values to try and deliver it to users to whom it matters and for whom it's important, because privacy as a value and a principle does matter to us as a group. It's not great for the advertising business so we would prefer that users did not elect to opt out. But, we do think it's the right thing to do to (A) be compliant because it's the law, and (B) because privacy is a fundamental right. So we are, in some ways, grateful that the mechanisms exist for us to be able to do that.

To this participant, advancing privacy is an important goal that is fundamentally incompatible with their business's corporate strategy, which is premised on an advertising model. Therefore, the mandates of GDPR and CCPA compliance offer the engineering team some reprieve to be able to advance data protection despite it potentially undercutting their employer's profitability.

4.4.2 "Regulations Do the Bare Minimum"

Participants responded to questions about the value of GDPR and CCPA with great skepticism. In general, they were not very hopeful about advancing a broader conceptualization of privacy through data protection regulations that were neither standardized nor enforced. However, several participants noted the value of compliance work instigating important discussions and processes within their teams. For example, one participant highlighted how the regulations allow developers on their team to signal their values to each other and promote discussion:

> I have kind of mixed feelings about [how important the regulations are]. I think it's important that they drive conversations at companies. I appreciate hearing someone say, "let's do more than we're legally required to do for our users." I would also appreciate hearing somebody say the opposite of that because then I would think—this isn't a company I want to work at. So I think that's probably the main benefit I see of it. I don't think the regulations do much to improve user data privacy. But I think they start conversations that that wouldn't happen if the regulations weren't in place.

Similarly, another participant appreciated that data protection regulations allow them to challenge their employer's business model:

> I think privacy is really important… In my work, I have tried to advance what I think is the spirit of the regulations in an environment even though it's sort of against the interests of the business that I'm working for. There's a conflict there…. But I've tried to do what I think the intention [of the regulations] was for our users, knowing that it's not enough. It's far from it. The regulations are ambiguous on the surface, and they're ambiguously implemented, although they're well intentioned. I don't know what the net benefit is, if any.

5. Discussion

5.1 Developers' Responsibility and Ability to Achieve Compliance

This study strengthens previous findings that developers are uniquely positioned among tech workers to shape what privacy means by being held responsible for implementing data protection regulations. While previous research has measured the influence of the "human factor" in software development in general, and privacy work in particular, this study adds a comparative perspective by interviewing tech workers with non-engineering functions, including design, data, and compliance. Participants' experiences suggest that product managers and engineering executives are well suited to shape data protection compliance work, but that developers are nonetheless often highly autonomous both by defining their work and by deciding when their decisions should receive external input, and from whom.

Several participants repeated that they were satisfied with trying to follow "the spirit of the law" rather than a highly determined notion of compliance. This was especially common when consulting lawyers, suggesting that "privacy" operates as a boundary object—a concept whose ambiguous definition enables cross-functional actors to collaborate without actually agreeing upon a single definition (Star & Griesemer, 1989). This ambiguity enables developers to define compliance work not only in terms of implementation decisions about specific technical features but also a more general sense of the scope of work. For example, the experiences cited in Section 4.4.1 about developers strategically interpreting the scope of GDPR and CCPA to advance goals outside the explicit scope of the regulations adds to Tahaei, Frik, et al.'s (2021) findings a new way that privacy champions influence their teams, and further suggests that privacy champions bear outsize influence beyond developer teams by shaping organizational strategy, especially in smaller organizations.

At the same time, this ambiguity also reflects developers' attitudes toward whether compliance seems possible at all. For example, several developers were frustrated by having to rely on internal and external systems and developer teams to have the capacity to accurately assess what they do with personal data—and then to also trustfully communicate those assessments. The magnitude of these dependencies suggests that developers feel that compliance is not attainable on an individual basis. Indeed, the anecdote from Facebook's leaked document at the beginning of this paper suggests that this inability to grasp and attain compliance on an individual level is, in fact, inherent in information technology systems.

Instead of measuring compliance by regularly reassessing what happens to specific types of data, then, developers often conceived of compliance as a one-time project that *other* actors are responsible for upholding through feedback. All the participants acknowledged that they have not returned to initial compliance instrumentation to ensure high quality, update temporary settings, or audit their work—even though several participants described implementing provisional configurations that they intended to return to in the future. Instead, they generally wait for bugs to be reported by users—which are unlikely to emerge because users are unlikely to notice or complain if their privacy preferences, such as cookie consent settings, do not work properly. In fact, several participants described bugs with their initial implementation of cookie consent notices and opt-ins. However, none of the participants reviewed that work on their own.

This suggests that end users can play an important role in increasing accountability for privacy work. For example, in the case of data subject access requests, several participants stated that they intended to enhance or automate their processes, but that they did not see the need to do

so because of the small number of access requests they have received. Moreover, participants agreed that very few users are likely to take advantage of features such as opting out of third-party cookies. Increasing pressure from end users, and making developers aware of user interest, may increase developers' motivations to return to and improve data protection features.

Finally, this study suggests that developers see their impact on privacy as highly limited in that much of their work is sending calls to third-party services and validating the response—without confirming what is actually happening to data upstream. This aligns with previous findings about the distributed nature of internet governance which manifests in a network of trust relationships (Mathew, 2014). Overcoming this opacity calls for further clarifying how this trust is built and operationalized, and potentially suggests that encouraging more collaborative communication may encourage developers to demonstrate and confirm the impact of their work.

5.2 Data Protection Compliance as Knowledge Production

These findings lead to an understanding of data privacy systems as fundamentally and necessarily marked by indeterminacy. Both the public and policymakers are unaware of the exact ways in which regulations are manifested in user-facing interfaces, so developers can be seen as experts who create and hold exclusive *knowledge* about what was actually enacted on a particular website or application. At the same time, developers are both strategically uninterested and structurally incapable of ascertaining a comprehensive, accurate, confident understanding of how and where data moves and is affected by privacy protections.

This conceptualization of *privacy knowledge* enables analytic distinction between the *context* in which privacy is constructed and the very *content* of privacy knowledge. In other words, it recasts privacy as not only a principle but a specific kind of expertise about the actual relationship between privacy expectations and privacy outcomes in a particular system. For example, certain developer teams are exclusively aware of the specific categories of data subject to user configuration in a given app, the advertisers and partners from whom that data is withheld, and the level of confidence, scrutiny, and due diligence applied to various components of data protection regulations. Once those decisions are enacted in code, the *actual* relationship between decisions and code is a category of knowledge exclusively but not fully available to developers.

Despite this indeterminacy, data compliance regulations often assume that data can be tracked in great detail. They expect that developers can confidently report on the status of both historical and future data: what kind of information they convey, where they are (and have been) stored, and which actors (humans, applications, and systems) have (and have had) access to them. Despite developers' admissions that such certainty is impossible, these systems are not seen as non-compliant by either developers or regulators. Instead, they see data privacy as a constant work-in-progress. Indeed, the leaked Facebook document states:

> It's simply inaccurate to conclude that it demonstrates non-compliance. New privacy regulations across the globe introduce different requirements and this document reflects the technical solutions we are building to scale the current measures we have in place to manage data and meet our obligations. (Franceschi-Bicchierai, 2022)

What can be ascertained by studying the content of privacy knowledge production through compliance work rather than focusing exclusively on the context? This paper argues that there are three primary implications of considering software developer teams as actors that hold and produce expert knowledge about privacy. First, thinking about organizations' tech teams as

sites of knowledge production opens up comparisons to scientific laboratories. Knorr Cetina (1992) argues that laboratories are important sites for sociological study because they provide a social context in which knowledge claims are made. This raises the stakes for building on previous research to examine how software developers approach privacy work (e.g., Greene & Shilton, 2018; Li et al., 2021; Shilton & Greene, 2019; Tahaei et al., 2019; Tahaei et al., 2020; Tahaei et al., 2021). In particular, it calls for further exploration of *articulation work*, which includes anticipating, planning, and organizing discontinuous elements to achieve a larger goal (Strauss, 1985), as a key component of defining privacy and enacting data protection regulations for the public. This is an alternative approach from studying outcomes by examining user interfaces or studying group processes for seeking information and making decisions.

Second, thinking about software developer teams as actors that hold and produce expert knowledge about privacy raises the urgency for an epistemic shift in conceptualizing what constitutes privacy. Specifically, Arora's (2018) call to "de-naturalize and estrange data from demographic generalizations and cultural assumptions" provides an important framework for reconceptualizing questions and approaches for understanding privacy. For example, she argues that the rise of the BRICS nations disrupts the core-periphery model that assumes a linear narrative of progress and standards emanating outward from the global North-West. Indeed, recent comparative studies of data protection regulations have found that strong conceptualizations of privacy in places like China, India, and South Korea predate and sometimes heavily deviate from the GDPR and CCPA. In addition, she argues that privacy scholarship should not rest upon an assumed rationality of either data subjects or data itself.

This argument about rejecting presumed rationality leads to a third implication of thinking about software developer teams as actors that hold and produce expert knowledge about privacy by calling for greater scrutiny of the specific social identities in such teams. Prior research about developers' approaches to privacy work often examines either small teams using qualitative methods (e.g., interviews with 8–12 developers in a single organization) or anonymous actors at scale using quantitative methods (e.g., content analysis of topics discussed in online forums). Both approaches divorce social identities from attitudes and work outputs, thus assuming a level of technocratic rationality. However, researchers have recently outlined a framework for comparative privacy research that has revealed significant differences in how people from different national or sociocultural contexts conceive of privacy differently (Masur et al., 2021). Given that GDPR and CCPA compliance labor is conducted by teams across national borders as well as, in particular, many Asian immigrants in the US and Europe, the relationship between software developers' social identities and their privacy work cannot be ignored.

6. Conclusion

This study has strengthened prior conclusions about the gap between privacy in law and in practice. On one hand, data protection regulations such as the GDPR and CCPA are seen as key strategies to protect individuals from the harms of digital data collection, including violating an expectation of privacy, collecting and selling personal information without consent, and threatening rights to free speech and association. On the other hand, prior research focusing on user-facing configuration options has revealed that features don't necessarily match expectations. While prior research has explored this gap by accounting for developers' information-seeking and communicative practices, this study has questioned whether there is truly a "gap" at all.

Instead, this paper has argued that data protection compliance work should be seen as a process of knowledge production in order to shift analytic focus from the exclusive study of social context to the content itself; that is, the knowability and facticity of compliance itself. This intervention is proposed to Waldman's (2019) diagnosis of *legal endogeneity*: symbols of compliance—such as training, audits, and documentation—standing in for real privacy protections. The stakes for understanding and reconciling this gap are increasingly high as more countries and jurisdictions around the world adopt data protection regulations, often modeled after the GDPR and CCPA, that assume that compliance is an inherently feasible concept in order to turn privacy into an individual responsibility in the form of "privacy self-management" (Hull, 2015). Therefore, data protection work merits particular attention because of its unique relationship to shaping public expectations of privacy vis-à-vis corporate and government surveillance and subjectification.

In summary, this paper argues that software developers should be understood not only as the "human factor" in sociotechnical systems, but instead as actors who hold and produce expert knowledge in the privacy domain. Understanding this role is important because privacy knowledge is by nature inaccessible to both policymakers and the public, while also encoding privacy norms for users. This paper has highlighted three implications: it opens up comparisons between developer teams and scientific laboratories as sites of knowledge production; raises the urgency of decolonizing privacy studies; and calls for scrutiny of developers' social identities as opposed to presuming rationality in interpreting privacy expectations. These implications provide a framework for exploring data privacy compliance as not only a socially constructed process but also a form of knowledge production uniquely yet not fully available to a population of developers whose approaches to compliance work have high stakes for the very concept of privacy itself.

7. Addenda

Funding: This study was supported by the USC Center on Science, Technology, and Public Life and the AEJMC Law and Policy division's Michael Hoefges Graduate Student Research Fund.

Acknowledgements: Many thanks to Mike Ananny, Jeeyun (Sophia) Baik, Christina Dunbar-Hester, Larry Gross, Andy Lakoff, Jennifer Petersen, Richmond Wong, and participants of the 2022 UC Berkeley CLTC Symposium on GDPR/CCPA for their feedback on earlier iterations of this project.

About the author: Rohan Grover is a PhD student at the Annenberg School for Communication and Journalism at the University of Southern California. He is also an affiliate of the UNC Center for Information, Technology, and Public Life (CITAP). His research focuses on privacy, internet governance, platform studies, and political communication technology, and his work has been published in *Telecommunication Policy*; *Information, Communication & Society*; and the *Asian American Policy Review*.

8. References

Arora, P. (2019). Decolonizing Privacy Studies. *Television & New Media*, *20*(4), 366–378. https://doi.org/10.1177/1527476418806092

Balebako, R., Marsh, A., Lin, J., Hong, J., & Faith Cranor, L. (2014). The Privacy and Security Behaviors of Smartphone App Developers. *Proceedings of the 2014 Workshop on Usable Security*. https://doi.org/10.14722/usec.2014.2300

Cavoukian, A. (2009). *Privacy by Design - The 7 Foundational Principles*. Information and Privacy Commissioner of Ontario.

Degeling, M., Utz, C., Lentzsch, C., Hosseini, H., Schaub, F., & Holz, T. (2019). We Value Your Privacy ... Now Take Some Cookies: Measuring the GDPR's Impact on Web Privacy. *Proceedings 2019 Network and Distributed System Security Symposium*. Network and Distributed System Security Symposium, San Diego, CA. https://doi.org/10.14722/ndss.2019.23378

Di Martino, M., Robyns, P., Weyts, W., Quax, P., Lamotte, W., & Andries, K. (2019.). Personal Information Leakage by Abusing the GDPR ``Right of Access''. *Proceedings of the Fifteenth Symposium on Usable Privacy and Security*, 371–386.

Feng, Y., Yao, Y., & Sadeh, N. (2021). A Design Space for Privacy Choices: Towards Meaningful Privacy Control in the Internet of Things. *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, 1–16. https://doi.org/10.1145/3411764.3445148

Franceschi-Bicchierai, L. (2022, April 26). *Facebook doesn't know what it does with your data, or where it goes: Leaked document*. Vice. https://www.vice.com/en/article/akvmke/facebook-doesnt-know-what-it-does-with-your-data-or-where-it-goes

Greene, D., & Shilton, K. (2018). Platform privacies: Governance, collaboration, and the different meanings of "privacy" in iOS and Android development. *New Media & Society*, *20*(4), 1640–1657. https://doi.org/10.1177/1461444817702397

Habib, H., Pearman, S., Wang, J., Zou, Y., Acquisti, A., Cranor, L. F., Sadeh, N., & Schaub, F. (2020). "It's a scavenger hunt": Usability of Websites' Opt-Out and Data Deletion Choices. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (pp. 1–12). Association for Computing Machinery. http://doi.org/10.1145/3313831.3376511

Habib, H., Zou, Y., Jannu, A., Sridhar, N., Swoopes, C., Acquisti, A., Cranor, L. F., Sadeh, N., & Schaub, F. (2019). An empirical analysis of data deletion and opt-out choices on 150 websites. *Proceedings of the Fifteenth Symposium on Usable Privacy and Security (SOUPS)*, 387–406.

Habib, H., Zou, Y., Yao, Y., Acquisti, A., Cranor, L., Reidenberg, J., Sadeh, N., & Schaub, F. (2021). Toggles, Dollar Signs, and Triangles: How to (In)Effectively Convey Privacy

Choices with Icons and Link Texts. *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, 1–25. https://doi.org/10.1145/3411764.3445387

Hadar, I., Hasson, T., Ayalon, O., Toch, E., Birnhack, M., Sherman, S., & Balissa, A. (2018). Privacy by designers: Software developers' privacy mindset. *Empirical Software Engineering*, *23*(1), 259–289. https://doi.org/10.1007/s10664-017-9517-1

Hull, G. (2015). Successful failure: What Foucault can teach us about privacy self-management in a world of Facebook and big data. *Ethics and Information Technology*, *17*, 89–101. https://doi.org/10.1007/s10676-015-9363-z

Knorr Cetina, K. (1995). Laboratory studies: The cultural approach to the study of science. In S. Jasanoff, G. E., Markle, J. C. Petersen, & T. Pinch (Eds.), *Handbook of Science and Technology Studies* (1st ed, pp. 140-167). Sage.

Li, T., Agarwal, Y., & Hong, J. I. (2018). Coconut: An IDE Plugin for Developing Privacy-Friendly Apps. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, *2*(4), 1–35. https://doi.org/10.1145/3287056

Li, T., Louie, E., Dabbish, L., & Hong, J. I. (2020). How developers talk about personal data and what it means for user privacy: A case study of a developer forum on Reddit. *Proceedings of the ACM on Human-Computer Interaction*, *4*(CSCW3). https://doi.org/10.1145/3432919

Masur, P. K., Epstein, D. Quinn, K., & Wilhelm, C. (2021, May). "Comparative Privacy Research Framework." 71st Annual International Communication Association Conference. https://osf.io/preprints/socarxiv/fjqhs/

Mathew, A. J. (2014). Where in the world is the internet? Locating political power in internet infrastructure (Publication No. 3685949) [Doctoral dissertation, University of California, Berkeley]. ProQuest Dissertations Publishing.

Mhaidli, A. H., Zou, Y., & Schaub, F. (2019). "We Can't Live Without Them!'' App Developers' Adoption of Ad Networks and Their Considerations of Consumer Risks. *Proceedings of the Fifteenth Symposium on Usable Privacy and Security*, 225–244.

O'Connor, S., Nurwono, R., Siebel, A., & Birrell, E. (2021). (Un)clear and (In)conspicuous: The Right to Opt-out of Sale under CCPA. *Proceedings of the 20th Workshop on Workshop on Privacy in the Electronic Society*, 59–72. https://doi.org/10.1145/3463676.3485598

Patnaik, N., Hallett, J., Tahaei, M., & Rashid, A. (2022). If you build it, will they come? Developer recruitment for security studies. *Proceedings of 1st International Workshop on Recruiting Participants for Empirical Software Engineering - The 44th International Conference on Software Engineering (ROPES-ICSE'22)*.

Shilton, K., & Greene, D. (2019). Linking platforms, practices, and developer ethics: Levers for privacy discourse in mobile application development. *Journal of Business Ethics*, *155*(1), 131–146. https://doi.org/10.1007/s10551-017-3504-8

Star, S. L., & Griesemer, J. (1989). Institutional ecology, 'translations', and boundary objects: Amateurs and professionals on Berkeley's Museum of Vertebrate Zoology. *Social Studies of Science*, *19*(3), 387-420. https://doi.org/10.1177/030631289019003001

Strauss, A. (1985). Work and the division of labor. *The Sociological Quarterly*, *26*(1), 1–19. https://doi.org/10.1111/j.1533-8525.1985.tb00212.x

Tahaei, M., Frik, A., & Vaniea, K. (2021). Privacy champions in software teams: Understanding their motivations, strategies, and challenges. *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. https://doi.org/10.1145/3411764.3445768

Tahaei, M., Jenkins, A., Vaniea, K., & Wolters, M. (2021). "I don't know too much about it": On the security mindsets of computer science students. In T. Groß & T. Tryfonas (Eds.), *Socio-Technical Aspects in Security and Trust* (pp. 27–46). Springer International Publishing. https://doi.org/10.1007/978-3-030-55958-8_2

Tahaei, M., Li, T., & Vaniea, K. (2022). Understanding privacy-related advice on Stack Overflow. *Proceedings on Privacy Enhancing Technologies*, *2022*(2), 114–131. https://doi.org/10.2478/popets-2022-0038

Tahaei, M., & Vaniea, K. (2019). A survey on developer-centred security. *2019 IEEE European Symposium on Security and Privacy Workshops*, 129–138. http://doi.org/10.1109/EuroSPW.2019.0002

Tahaei, M., & Vaniea, K. (2021). "Developers Are Responsible": What Ad Networks Tell Developers About Privacy. *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems*, 1–11. https://doi.org/10.1145/3411763.3451805

Tahaei, M. & Vaniea, K. (2022). Recruiting Participants With Programming Skills: A Comparison of Four Crowdsourcing Platforms and a CS Student Mailing List. *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. https://doi.org/10.1145/3491102.3501957

Tahaei, M., Vaniea, K., & Saphra, N. (2020). Understanding privacy-related questions on Stack Overflow. *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. https://doi.org/10.1145/3313831.3376768

Waldman, A. E. (2019). Privacy Law's False Promise. *Washington University Law Review*, *97*(1), 773–834.