Privacy-Preserving Analytics on the Ground

Ryan Steed[1] and Alessandro Acquisti

Heinz College of Information Systems and Public Policy

Carnegie Mellon University

Working draft—please ask before quoting or circulating.

Abstract: Techniques for privacy-preserving analytics (PPA) offer organizations a way to maintain and expand access to valuable insights while preserving the privacy of data subjects. Adoption of PPA is growing in both the private and public sectors, but their impacts are not yet clear: small differences in technical design can have great impacts on data privacy, and little work examines the decisions that determine whether adoption lives up to the expectations of consumers and regulators. With a qualitative study of executives, lawyers, engineers, and scientists implementing PPA across 21 large tech firms and startups, non-profits, and government agencies, we investigate the organizational processes driving the adoption and deployment of PPA systems. We develop a process model of the key drivers of adoption, the ways practitioners interpret those drivers into specific designs, and the ways organizations justify designs. We also highlight ways that organizations 1) center managerial concerns during these processes, 2) decouple the representations they make about privacy from the specifics of their implementation, and 3) use PPA adoption to influence privacy expectations in turn. But, we also find that morally motivated practitioners in these organizations leverage their expertise to spread adoption and maintain their own standards. We explore the consequences of these findings for the future of privacy regulation and research.

Keywords: privacy-preserving technology, technology adoption, organizational decoupling, legal endogeneity, algorithmic systems

[1]Email: ryansteed@cmu.edu.

Privacy-Preserving Analytics on the Ground

We continue to believe personalized ads and privacy can co-exist... that's why we're investing in research and development of privacy-enhancing technologies.

— Erin Egan, Meta Chief Privacy Officer (Egan, 2020)

1 INTRODUCTION

Facing consumers' calls for privacy and policymakers' threats of regulatory action, tech leaders are endorsing "privacy-preserving" analytics techniques as the future of digital privacy (Egan, 2020). Organizations in industry and government—including Meta, Google, Apple, Microsoft, Wikipedia, Mozilla, the U.S. Census Bureau, and the Internal Revenue Service (Desfontaines, 2021)—are pioneering deployments of differential privacy (Dwork et al., 2006), federated learning (McMahan et al., 2017), multiparty computation (Goldreich, 2009), and other approaches to data privacy.

Yet, despite a flourishing technical literature in computer science and statistics, little economic or social science research has examined organizations' strategic use of these techniques. What drives organizations to adopt privacy-preserving technologies? And how do these drivers inform design choices made during deployment?

On its face, consumers and regulators have reason to hope adoption will improve digital privacy: these mathematical, statistical, and computational practices—which we refer to as *privacy-preserving analytics* (PPA)—are used to produce useful insights from people's personal data while still preserving some technical definition of data privacy. But, these techniques are complex in theory and implementation, and in many cases their impacts on consumer privacy and social welfare are untested and uncertain. The U.S. Census Bureau's differentially private 2020 Decennial Census tabulations, for instance, have been emphatically protested by stakeholders concerned that the change will sacrifice research and governance for little privacy benefit (boyd & Sarathy, 2022). Advertisers (Nguyen, 2022) and privacy advocates (Cyphers, 2021) alike have serious doubts about Google's plan for the future of targeted advertising, which is centered on a suite of new PPA technologies (Goel, 2022). In short, PPA adoption could present new benefits to researchers and

consumers—or, as some advocates fear, it could simply preserve extractive, surveillance-based economies (Cohen, 2018; Zuboff, 2019).

In absence of conclusive evidence, organizational research offers several theoretical explanations for the spread of PPA adoption and what it might mean for digital privacy. One explanation is that changing consumer expectations and regulatory pressure force organizations to innovate more privacy-preserving approaches to data analytics, similar to innovation in environmental protection (Ashford et al., 1985). But other organizational research suggests that formal policies adopted to satisfy consumers and regulators may be merely symbolic—"decoupled" from substantive changes to daily practice (Meyer & Rowan, 1977; Edelman, 2016). On this trajectory, PPA adoption could amount to nothing more than "privacy theater": gestures towards privacy accorded deference incommensurate with their actual benefits to people online (Soghoian, 2011).

In this qualitative study, we use grounded theory methodology (Charmaz, 2014) to better explain the motivations and decision-making processes behind PPA adoption and design (§4). We interviewed 28 practitioners involved in PPA adoption efforts across 21 tech companies, privacy startups, internet non-profits, and government agencies. Analyzing transcripts and documents, we developed an emergent theoretical model of PPA adoption. We find that organizations transitioned to PPA primarily to preserve existing modes of operation against new regulations and consumer expectations, but also to gain advantages against competitors (for businesses) and to open up new opportunities for researchers. As organizations developed systems to satisfy regulators and consumers, PPA design required extensive technical and legal interpretation from experts. In the absence of technically prescriptive regulation, these practitioners were left to choose and justify specific parameter settings and other technical designs. Through infrastructure sharing and active standardization work, these choices set the bar for future industry implementations and for regulatory guidance.

Though we identified many potential paths to decoupling, we also found that morally-motivated institutional entrepreneurs (DiMaggio, 1988) constituted a countervailing force within the organization. The privacy practitioners we interviewed—executives, managers, and key contributors in

technical, legal, and policy roles—were responsible for deciding whether to adopt and how to deploy PPA systems. As experts, they had significant leverage over the interpretation of privacy rules (from law or internal policy) on the one hand and the technical properties of PPA systems on the other, mediating between managerial, legal, societal, and technical concerns. An influential subset of these practitioners used their influence to act as "privacy champions" within their organizations (Tahaei et al., 2021), pushing for PPA adoption out of their own personal ethics.

This study makes two theoretical contributions to research on information systems (IS), organizational behavior, and privacy. First, by comparing our theoretical model to research from sociology, law, and technology studies, we identify several ways that the adoption of technology—in particular, algorithmic technology—may be decoupled from institutional expectations in unique ways, a concept we call *policy-technology* decoupling. As technological artifacts, PPA systems admit to technological interpretation and practitioners' technological "readings"—in addition to legal readings (Fuller et al., 2000)—leaving more room for managerial influence, but also empowering technologists to exercise moral leadership (Solinger et al., 2020). Often instantiated as *algorithmic* technologies in particular, PPA practices may receive less scrutiny in practice and in court (K. A. Bamberger, 2010; Metcalf et al., 2023). Technology also has influence on the institutional environment in turn: norms embedded in technological infrastructure can scale and spread quickly through centralized, shared infrastructure like open source libraries or cloud services (Lessig, 1999). Second, our findings suggest that PPA adoption may not be an unalloyed benefit to society, if consumer and regulatory privacy expectations are overly mediated by managerial concerns. Our process model, connected with work on privacy regulation, offers a guide for policymakers to more critically evaluate and intervene in PPA adoption.

2 THEORETICAL BACKGROUND

In this section, we situate our theoretical model of PPA adoption within existing theories from IS, organization behavior, and other fields. First, we describe our setting: the class of technologies adopted to preserve privacy in data analytics (§2.1). Then, we outline potential institutional pressures

leading organizations to adopt socially desirable technologies like PPA (§2.2)—these pressures constitute the institutional environment for our model of adoption within the organization. Finally, inspired by the processes that emerged in our analysis, we summarize other research suggesting that those drivers may not be entirely exogenous—their impacts on daily practice are mediated by the organization, and adoption may reshape institutional norms and expectations (§2.3).

2.1 Privacy-Preserving Analytics

While all PPA technologies are used in practice to preserve data subjects' privacy, the methods used—and the precise definition of privacy preserved—vary widely. We consider techniques used to preserve privacy in either the inputs to data analysis—e.g., secure multiparty computation (Goldreich, 2009)—or to preserve privacy in the outputs. A differentially private mechanism, for example, injects additional noise into statistics to formally ensure that they are not sensitive to the inclusion of any one individual's information (Dwork et al., 2006). We also consider technologies used to preserve privacy in practice even when they do not always satisfy a formal guarantee of privacy. Federated learning, for example, is a technique for training machine learning models on multiple devices without transferring raw data off-device (McMahan et al., 2017). Table 1 lists all the PPA practices used by our participants.

Privacy technology is not new—PPA techniques are a subset of privacy enhancing technologies (PETs),[2] a variety of tools used by consumers, regulators, and organizations to negotiate information privacy issues for at least the last three decades (Goldberg, 2007). But organizational interest in PPA technologies specifically has increased markedly in the last decade. Deployed systems in industry and government use these techniques to share data with researchers, suggest emojis, and more (Desfontaines, 2021). New startups—such as Tumult Labs—are offering PPA services and software. Still, the impacts of PPA—for digital privacy on the one hand and for data science, social science research, policymaking, and other data-dependent processes on the other (Abowd & Schmutte, 2019; Hotz et al., 2022)—are largely untested outside primarily theoretical research in computer science

[2]We do not include PETs for private communication or authentication (see, e.g., Domingo-Ferrer & Blanco-Justicia, 2020)—we are specifically concerned with technologies used in data analytics.

and statistics (Acquisti & Steed, 2023). A few interview studies explore practitioners' challenges with differential privacy adoption specifically, but these studies mostly hone in on specific usability issues (Dwork et al., 2019; Garrido et al., 2022; Sarathy et al., 2023). Others investigate related practices for business data ethics (Hirsch et al., 2020) and privacy in artificial intelligence (AI) products (H.-P. Lee et al., 2024). Our study contributes to the ongoing project of documenting and describing the social impact of PPA technologies, starting with the organizational processes explaining their adoption and deployment.

2.2 Reasons to Adopt

Our work adds to decades of research on the interaction between institutional expectations and technological innovation, particularly related to health and the environment. The practice of adopting socially beneficial technologies is a form of corporate social performance (CSP): efforts to accomplish social benefits alongside economic gains (K. Davis, 1973)—social benefits which, some argue, now include data privacy (Soghoian, 2011; Pollach, 2011). We focus on the specific case where CSP is accomplished at least in part with technology adoption—for example, the adoption of "green" IS practices to improve environmental sustainability (Leidner et al., 2022).

In a traditional conceptions of the firm—in which organizations seek to rationally economize transaction costs, from employment relations to contract law—CSP practices are a rational adaptation to changes in regulation and consumer demand, particularly when there is a sense of crisis (Oliver, 1991; Silva & Hirschheim, 2007). Stricter environmental standards, for example, forced firms to innovate new technologies, such as the catalytic converter, to avoid financial penalties (Ashford et al., 1985; Orlitzky & Benjamin, 2001) or, in instrumental stakeholder theory, to improve relationships with instrumental stakeholders (Aguilera et al., 2007; Jones, 1995).

In an institutional view, organizational strategy is is also a product of "rational myths"—widely accepted prescriptions for successful strategy (Meyer & Rowan, 1977; DiMaggio & Powell, 1983). The relationship between market conditions and organizational strategy is mediated by the *institutional environment* in which an organization operates (Meyer & Rowan, 1977)—normative

pressures from activists, industry norms, stakeholders, and independent watchdogs (Campbell, 2007; Boldosova, 2019), all loosely coordinated (rather than strictly imposed) by regulations (Kagan et al., 2003). In this setting, CSP contributes to organizational legitimacy: the ways that organizations justify their existence (Weber, 1978; Meyer & Rowan, 1977; M. C. Suchman, 1995). Legitimacy may earn private and public organizations alike increased deference from external stakeholders (Cavazos & Rutherford, 2012)—their social license to operate (Gunningham et al., 2004).

Implications for PPA adoption: In the last decade, Chief Privacy Officers (CPOs)—responding to changes in consumer expectations and impending regulation—recognized an opportunity go "beyond compliance" with ambitious privacy agendas (Gunningham et al., 2004; K. A. Bamberger & Mulligan, 2015). Organizations had similar motives to adopt corporate data ethics (Hirsch et al., 2020) and privacy practices related to AI (H.-P. Lee et al., 2024). Our model of PPA adoption starts with these institutional pressures—namely, regulatory and consumer expectations—and our participants' motivations to adopt echoed this literature. But unlike the catalytic converter or green IT, or other previously studied social innovations, PPA is still on the cusp of widespread adoption. Our study thus uniquely focuses on the earliest stages of adoption, documenting the ways that technological designs come to be accepted as appropriate solutions to institutional pressures. In particular, we examine the ways external pressures are channeled into specific organizational narratives and the ways technologists translate those expectations into specific designs.

2.3 Decoupling & Legal Endogeneity

Organizations certainly do react to exogenous institutional pressures, but further institutional research complicates that simple model. In the process of responding to exogenous factors, organizations may 1) gain the legitimacy benefits of compliance by implementing formal policies without substantively changing practice (organizational decoupling) or 2) attempt to influence the social expectations that constitute compliance (legal endogeneity). Thus, organizations *mediate*—or influence—the impact of institutional pressures on daily practice (Edelman, 1992).

Organizational decoupling: First, research on organizational *decoupling* shows that institutional

pressures push organizations towards certain formal structures and legitimacy-seeking behavior, but those formal intentions may be "decoupled" from daily actions (Meyer & Rowan, 1977; Weick, 1976). For example, some firms obtained green technology certificates from the Korean government without actually implementing those technologies in daily operations (S. Park & Cha, 2019). Decoupling results when institutional pressures on the organization are mediated internally—the success of enterprise IS adoption, for example, is mediated by top management participation (Liang et al., 2007) and organization structure (Silva & Hirschheim, 2007). More recent research distinguishes policy-practice decoupling, in which policies are adopted without the accompanying effort required to alter work routines, from means-ends decoupling, in which policies change routines but not outcomes (Bromley & Powell, 2012).

Legal endogeneity: Second, organizations model their early "educated guesses" at compliance to the rest of the market. As in the case of equal opportunity, employment law, and insider trading, these guesses may be legitimated by courts and policymakers (Edelman et al., 1999; Edelman, 1992; Bozanic et al., 2012; Sutton & Dobbin, 1996). After defendants began instituting grievance procedures to forestall unionization and insulate against discrimination suits, for example, courts and legal journals increasingly considered those procedures relevant to liability despite little evidence that they actually reduced complaints (Edelman et al., 1999). Courts have shown similar deference to algorithmic decision systems (Selbst et al., 2023).

This phenomenon is called *legal endogeneity*: after organizations decide what forms of compliance are reasonable (*managerial mediation*), those practices are institutionalized as a rational market response to regulation (Edelman, 2016). Organizations may also engage directly in lobbying, corporate-sponsored research, and other forms of direct regulatory capture to promote their versions of compliance (Hillman et al., 2004; Kamieniecki, 2006). Tech companies like Airbnb, Uber and DraftKings in particular are exemplars of "regulatory entrepreneurship"—the pursuit of business models that are predicated on changing the law (Pollman & Barry, 2016).

Implications for PPA adoption: Could the marketing materials and privacy policies used to

describe PPA adoption diverge from the realities of code and operation? There is evidence of organizational decoupling in research on privacy practices in general: Waldman (2017), for example, distinguishes CPOs' privacy myth-making efforts from their actual performance by technologists on the ground, who had little material incentive to enact new privacy agendas. In general, policy-practice decoupling is more likely when adoption is nascent, but also depends on other factors like executives' personal commitment to ethics and the extent to which compliance is externally motivated (Weaver et al., 1999; Bromley & Powell, 2012).

While evidence suggests consumer and regulatory pressures have changed privacy policies in the past (K. A. Bamberger & Mulligan, 2015), the organizational motivations for PPA adoption specifically have not been thoroughly explored. And nearly all studies of organizational decoupling focus on primarily non-technological practices (e.g., Table 2, Bromley & Powell, 2012). PPA adoption differs from these practices in key ways—for example, PPA systems are technological artifacts, often involving complex algorithms and applied to algorithmic systems. Our grounded theory, developed in the remainder of this paper, explicates the organizational mechanisms governing PPA adoption (§4) and complicates existing theory, highlighting additional avenues for organizations and morally-motivated technologists to influence the course of adoption (§5).

3 METHODS

To investigate the drivers of PPA adoption and how they affect PPA deployment, we conducted a qualitative study of PPA adoption through semi-structured interviews with PPA practitioners—including engineers, lawyers, managers, researchers, policy experts, and executives—at technology firms, privacy-focused startups, non-profits, and government agencies. Qualitative methods have a long and impactful history of helping researchers theorize abductively about emerging phenomena in IS (Monteiro et al., 2022; Wiesche et al., 2017). Drawing from grounded theory methodology (Charmaz, 2014) and the "Gioia methodology" for rigorous qualitative research (Gioia et al., 2013; Strike & Rerup, 2016), we conducted a first-order analysis—with line-by-line annotations—to represent the experiences and perspectives of PPA practitioners. From this analysis emerged key

concepts describing the processes involved with adopting and deploying PPA. In a second-order analysis, we organized these concepts into a process model of adoption, presented in §4, explaining the processes driving organizations to adopt PPA and the processes governing their design. We then compared that theoretical model to prior research §5.

3.1 Sample

Our data are comprised primarily of semi-structured interviews with the individuals responsible for helping their organizations decide whether and how to implement PPA systems. Our questions focused on these practitioners' experience with researching, choosing, deploying, and justifying PPA techniques in their organizations. We contacted practitioners working on PPA products or services, with an emphasis on recruiting participants underrepresented backgrounds in computing. We first identified candidates from a small, initial set of contacts mostly within the second author's network as well as prominent PPA community members known to both authors through their public work in the field. We then asked those candidates to recommend one or two others we should speak to (snowball sampling), requesting especially participants from underrepresented backgrounds in computing. We also contacted individuals at particular organizations we knew to have deployed or considered deploying PPA to round out the diversity of organizations in our sample, though not all organizations in our sample have actually deployed a PPA system. (Most have made it at least as far as prototyping.) After we developed an initial model of development and deployment based on a first round of interviews, we conducted a second round interviews—with an investigator at a regulatory agency, as well as an additional privacy lawyer and a public policy representative in tech—to round out our understanding of how private firms communicate with regulators about PPA adoption, a form of theoretical sampling (Charmaz, 2014).

Though our study relies on analytical, rather than statistical, generalization (A. S. Lee & Baskerville, 2003), we sought to thoroughly explore different adoption settings in our data. Of the 21 organizations represented in our sample, 8 are technology companies ($N = 10$), 6 of which are members of the Fortune 500 ($N = 8$). Five privacy-focused startups are represented ($N = 6$)—organizations

Employer	PPA practices mentioned	Roles	Participants
Startup	DP, SDL, pseudonymization, encryption, minimization, cohort analytics, k-anonymity, deletion, PII detection, other cryptography	Director/Executive, Software/Privacy Engineer	P8, P13, P15,P17, P18, P24
Other for-profit	DP, SDL, k-anonymity, l-diversity, FL, HE, SMC, synthetic data, PPML, private set intersection, encryption, access control, retention limits, cohort analytics, other cryptography	Director/Executive, Manager, Software/Privacy Engineer, Data Scientist, Researcher	P2, P3, P4, P7, P11, P12, P14, P19, P20, P23, P25, P26, P28
Non-profit	k-anonymity, DP, minimization, deletion, retention limits, SDL, other cryptography	Director/Executive, Engineer, Researcher	P1, P9, P10, P22, P27
Government	DP, SDL, noise infusion, SMC	Director/Executive, Software Engineer	P6, P16, P21

DP: differential privacy. FL: federated learning. HE: homomorphic encryption (Gentry, 2009). SMC: secure multi-party computation. PPML: privacy-preserving machine learning (e.g., Abadi et al., 2016). SDL: other statistical disclosure limitation (Matthews & Harel, 2011), specifically aggregation, cell suppression, or data swapping. Cohort analytics: other privacy-preserving methods for clustering individuals and analyzing the groups.

Table 1: Current employer, PPA techniques used, and participant roles at the time of interview. Participants were given differential privacy (DP) and federated learning (FL) as examples, but allowed to name any practices they used for PPA.

whose products are differentiated by their focus on privacy, or organizations who offer PPA consultation or other services. Our sample also includes 4 non-profits ($N = 5$) and representatives from 3 U.S. government agencies ($N = 3$). Two of the government agencies were responsible for federal data collection and public statistics; one was a regulatory agency. For large organizations with large-scale or wide-ranging PPA activities, we recruited at least two or more participants, to add alternative perspectives on the same processes.

Our sample is predominantly American (90%), white (70%), non-Hispanic (85%), heterosexual (65%) and cisgender male (55%), as a proportion of participants who chose to self-identify for each category (22 for race, ethnicity, and gender; 17 for sexuality)—similar in race- and gender- diversity to typical high-tech professions (U.S. Equal Employment Opportunity Commission, 2016).

3.2 Data Collection

Interviews: The bulk of our study—interviews with 25 practitioners—was conducted between September 2021 and January 2022. We conducted the second round of interviews with 3 practitioners in August 2023. Each participant ($N = 28$) first completed a short, 5-minute demographic survey, then participated in a 50–100 minute interview over Zoom, under the condition that their identities and the names of their organizations will be kept confidential. Participants were compensated with their choice of a $30 gift card or donation to a charity. Each interview was recorded (with participant consent and IRB approval) and transcribed verbatim by the lead author.

Our interview focused on broad, open-ended questions about organizational processes involved with: 1) the decision to adopt PPA, including motivations, major trade-offs, and research; 2) design and deployment, especially communication, common challenges faced, and future trends. While most of our questions focused on specific projects our participants have been involved in—to determine why their organizations considered adoption and how adoption progressed—we also asked a few very broad questions about perceived trends in the field. The interview guide (Appendix B) includes a list of participant-oriented, open-ended, non-sequential questions with multiple formulations, along with suggested follow-up questions. Interviews were *semi-structured*—that is, co-constructed by the interviewer and the participant, who guides the tone, pace, and path of discussion (Kallio et al., 2016; Spradley, 1979). Our interview guide was designed to be flexible but still grounded in our area of interest. We tested our interview guide using two think-aloud interviews (Willis & Artino, 2013) with colleagues and three practice interviews with volunteer graduate students who had experience with building PPA products in industry. To ensure rigor and clarity, we adopted an insider-outsider approach to interview design and analysis (Strike & Rerup, 2016). The first author lead all the interviews, and the second author also participated in a random selection of interviews to help constantly tune the interview guide throughout the data collection period.

Documents: We supplemented our interview data with other documents: internal documents and public white papers provided by participants, public information, news articles, and other archival

documents. Official documents produced by the organization, such as internal design documents or published press releases, often reflect shared meaning and denote formal structures (Charmaz, 2014). We did not code these documents line-by-line, but they did inform our analysis.

3.3 Analysis

After transcribing the interviews, our analysis cycled continuously between 1) first-order participant-centered, inductive creation of analytic codes and categories, 2) second-order abductive aggregation and theoretical analysis, and 3) written and visual presentation of our emergent theoretical model, grounded in first-order quotations (*in vivo* codes) and emergent concepts. Analysis did not only occur after interviews—the interviewer was empowered to follow up on new topics of theoretical interest or on theoretical concepts that were not yet fully described, or saturated (Charmaz, 2014).

In first-order analysis (also called open coding; Charmaz, 2014; Saldaña, 2013), the first author annotated lines of data with short, precise, very granular descriptions—over 2,500 unique codes—of what the data describes, staying grounded in the participants' point-of-view (often by using the participants' own words) and focusing on actions and processes to avoid imposing preconceived framing on their responses (Charmaz, 2014). For example, initial coding of practitioner motivations included "'keeping data around longer'" (a participant's words) and "following an internal leader" (our words). In our codes, we aimed to avoid uncritically adopting the participants' views or stock disciplinary categories, also paying attention to things left unsaid or implied by the data (Charmaz, 2014). Additionally, the second author re-coded a sample of four of these interviews, providing critical feedback which we used to calibrate the rest of our coding.

In second-order analysis (also called focused coding; Charmaz, 2014), we took a critical view of the initial codes, sorting and synthesizing open codes (axial coding) to draw out theoretical hypotheses and narratives (theoretical coding). Where the first-order analysis stays close to processes and events in the data, in this stage we actively reduced the many first-order codes, grouping codes by their similarities and theoretical relevance to establish tentative concepts. While constructing these emergent concepts, we maintained a focus on processes—the actions participants and their

organizations took. Figure 1 displays a key selection of these concepts. For example, the concept "preserving data-driven operations" helped us explain the first-order observation that practitioners used PPA to keep data around longer.

We gradually formed and grouped those concepts, bottom-up, into multi-layer hierarchies and eventually aggregate themes describing the overarching processes involved with adoption and deployment. We labelled new concepts with descriptors—still preferring the participants' own words to avoid introducing prior concepts and terminology—and began to map the relationships between them with written notes and iterative process diagramming. As in the first-order analysis, the authors discussed and exchanged notes about the concepts and their relationships. Because our sample includes multiple organizations, as well as multiple participants from the same organization, we also compared and generalized across organizations and industries, similar to a case study design (Eisenhardt, 1989). We refer to these concepts and the relationships between them as an emergent process model because they *emerged* from our first-order annotations, which describe our data (Gioia et al., 2013).

After we developed an emergent process model based on the data, we also began to consider how our results compared to prior research. This comparison surfaced three key lenses to consider in our analysis: managerial mediation (which may result in organizational decoupling) and legal endogeneity. We examined our model through these lenses, highlighting mechanisms by which PPA adoption mediates the impacts of institutional pressures or influences the institutional environment in turn. We also highlight a third interest: areas where PPA experts had high influence. Figures 2–5 depict our final process diagrams and highlight key processes associated with these theoretical lenses. §5 draws out new dimensions for decoupling and endogeneity revealed by this analysis.

4 FINDINGS

In the following analysis, we explain how organizations narrated and constructed PPA adoption as an appropriate response to changes to the institutional environment (§4.1). Our emergent process model of PPA adoption is segmented into three main themes. First, we describe the internal organizational

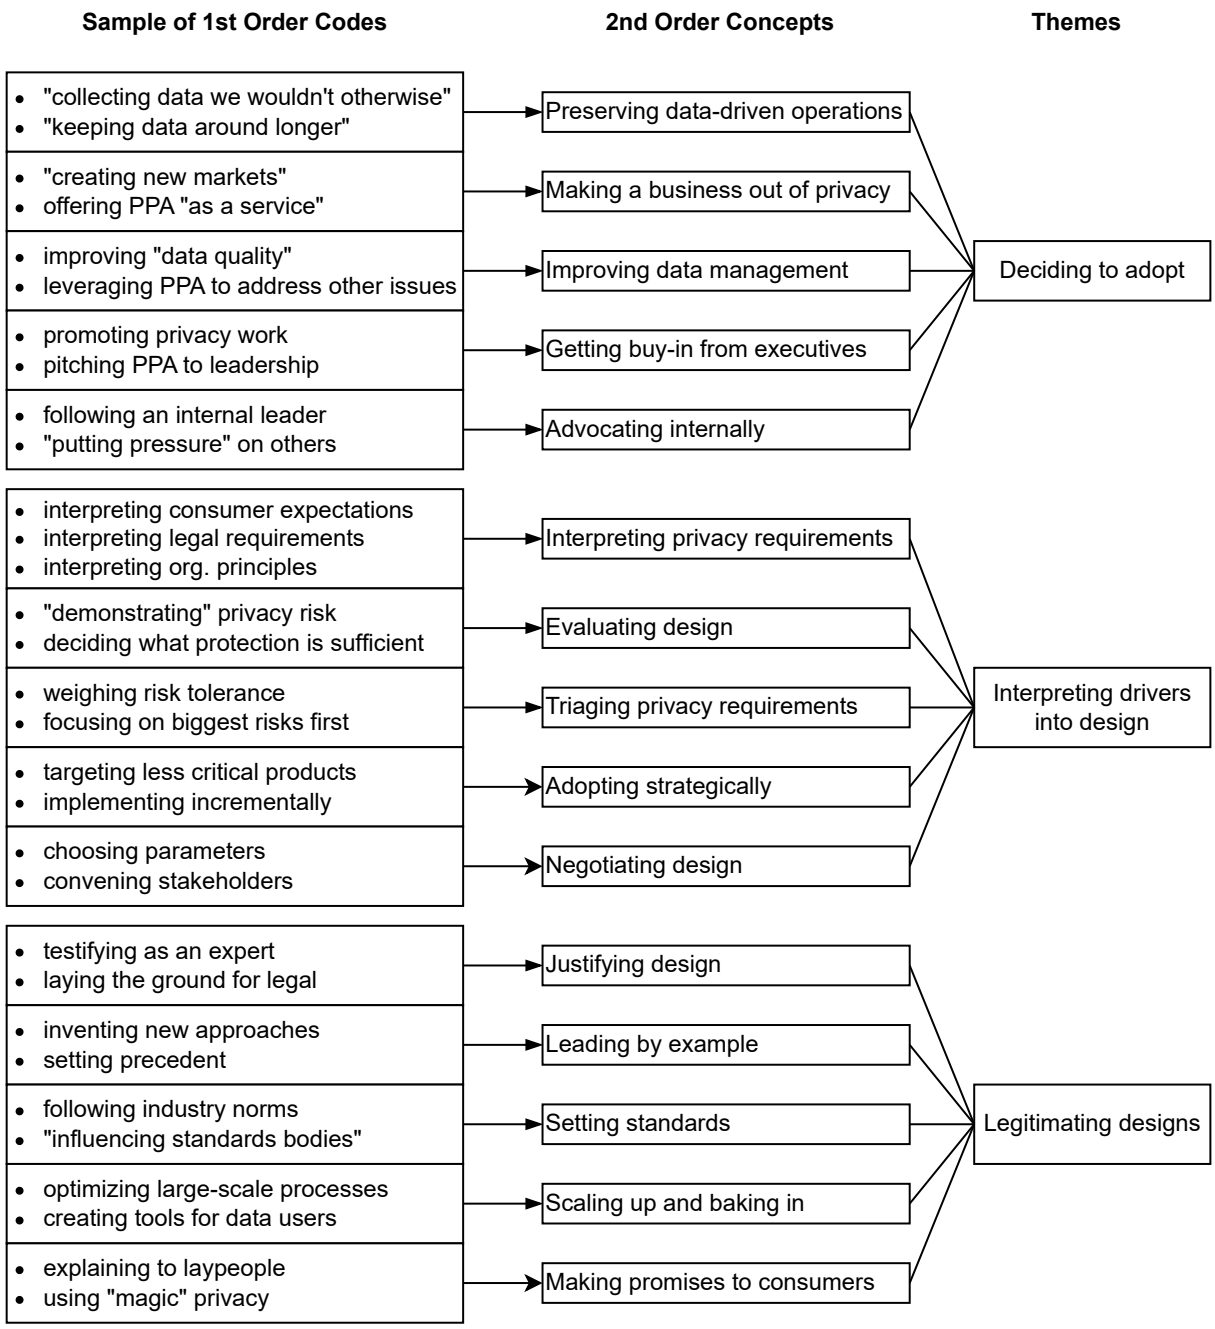


Figure 1: Themes arising from initial codes. Figure inspired by Gioia et al. (2013).

narratives practitioners use to motivate PPA adoption (§4.2). Second, we describe the processes by which those adoption drivers are interpreted into specific technological design choices (§4.3). In the course of our analysis, we also discovered that the decision to adopt and subsequent design choices are not only driven by external privacy expectations—in some ways, organizations influence those

expectations in turn. So third, we also describe the ways that organizations represent and legitimate PPA adoption to stakeholders, influencing the external drivers of adoption in turn (§4.4).

4.1 Institutional context

Across the array of organizations we studied, PPA adoption was overwhelmingly driven by the expectations of regulators, consumers, and their advocates. Practitioners across organization types said that the threat of regulation (including lawsuits and enforcement fees) contributed to their organizations' final decision to adopt—sixteen (especially those at tech for-profits and privacy startups) mentioned specific existing legal requirements or agreements with regulators. Seven participants at tech companies and non-profits described other negative consequences of bad privacy practices, including negative media coverage and public backlash after scandals that could harm their organization's brand and make it more difficult to recruit talented employees.

Public pressure was most pronounced in the wake of a data scandal—the Cambridge Analytica scandal (Confessore, 2018), for example—but even when social norms were not "loud and furious" (P4), practitioners anticipated institutional pressure. One employee of a tech company said:

> *I think a lot of folks in the industry are concerned about regulators stepping in and saying okay here's how the web should work. And all of a sudden, the whole business model of the web would fall apart. (P9)*

Google Chrome, for example, intends to replace third party cookies, a ubiquitous technology for tracking web activity, with a suite of new techniques in their Privacy Sandbox (Goel, 2022). Meta, another large online advertiser, has likewise invested in research on "novel privacy-preserving technologies" for ad measurement because of their belief that "it is possible to have privacy and a thriving, free, ad-supported Internet... personalized ads and privacy can coexist" (Egan, 2020). Shifts in regulation and consumer expectations were the most commonly cited changes contributing adoption, but some participants some organizations also pursued PPA after realizing that their current privacy protections were inadequate against modern linkage and reconstruction attacks (e.g.

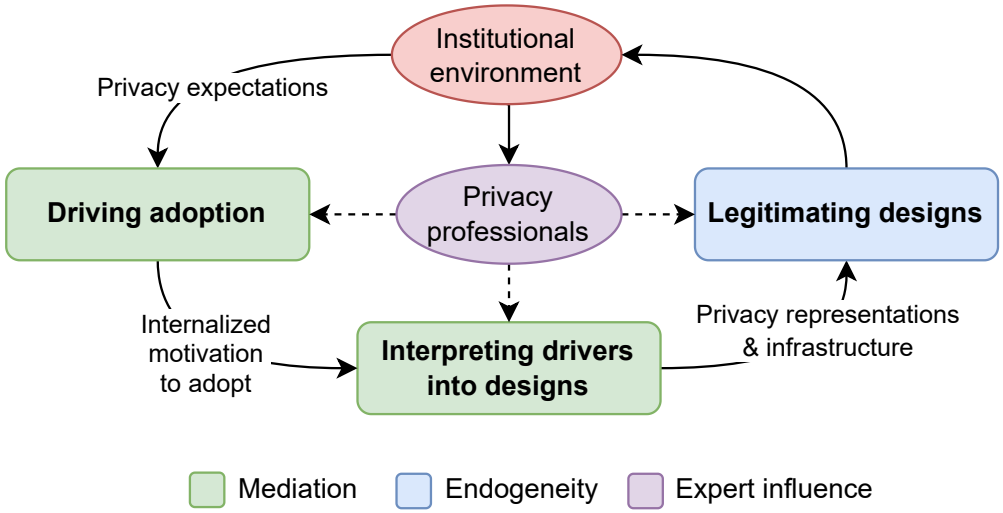


Figure 2: Summary of emergent process theory. Key processes associated with managerial mediation (and possibly decoupling), institutional endogeneity, or expert influence are highlighted.

the U.S. Census Bureau Abowd & Hawes, 2023).

4.2 Driving adoption

Why do organizations adopt PPA? Our second order analysis surfaced three main narratives practitioners used to motivate PPA adoption in their organizations (§4.2.1): first, preserving data-driven operations (accommodating privacy expectations); second, making a business out of privacy (leveraging privacy to open new markets); and third, improving data management (meeting expectations as a side benefit of good data practices). These narratives are the accounts practitioners gave of their organizations' motivations to adopt, and in some cases, the stories practitioners told to convince decision-makers. Figure 3 depicts our emergent model of this decision-making process.

4.2.1 Motivating adoption

The narratives practitioners used to motivate PPA adoption are an important initial mediator of institutional pressures—they describe the ways adoption becomes relevant to the organization.

Preserving data-driven operations: Most commonly, we find that organizations sought to leverage PPA to preserve a model of operation—a business model, for the firms in our sample—imperiled by privacy scandals and new or impending data protection regulations. PPA is framed as a compromise

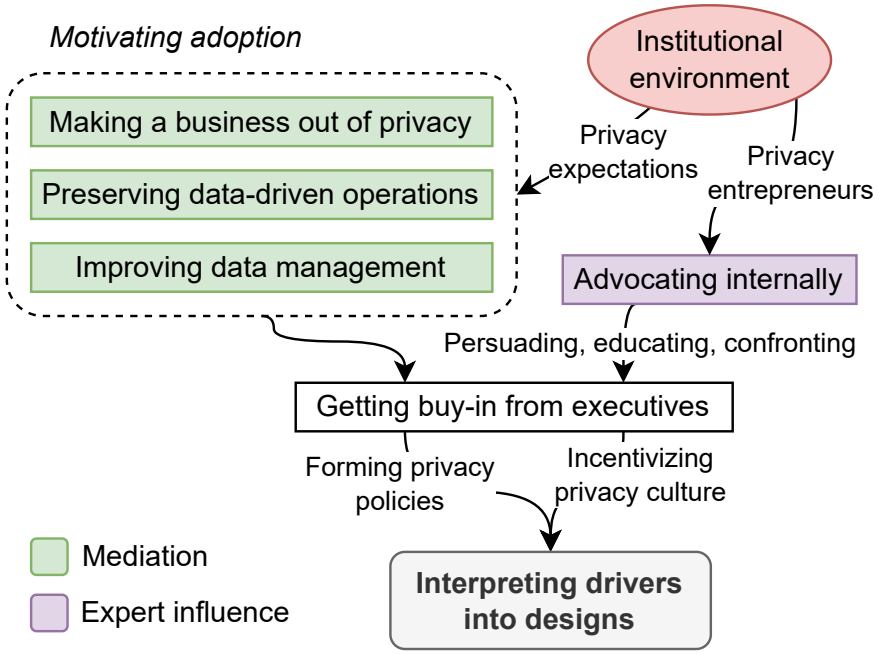


Figure 3: Driving adoption. Key processes associated with managerial mediation (and possibly decoupling), institutional endogeneity, or expert influence are highlighted.

between business interests and privacy concerns. As Meta's Chief Privacy Officer said in a blog post, "We continue to believe personalized ads and privacy can co-exist... that's why we're investing in research and development of privacy-enhancing technologies." Not all of the tech industry may believe that PPA is a viable option–an employee at a regulatory agency observed a "race to the bottom" in industry, in which "everyone thinks that if they jump to something that's more privacy-preserving and lose [the capability to precisely track individuals], then they'll be worse off" (P5). But when the alternative was losing analytics altogether, organizations adopted PPA to collect the data they otherwise could not due to a regulation or other expectation, or to keep data around longer than would otherwise be acceptable to consumers and regulators. One privacy director said, "People say that [personal] data is the new oil. I say data is the new toxic waste... How can we actually get what we need, without having that personal identifying information?" (P9)

Seven participants said their organizations (all for-profit firms but one) adopted PPA to "keep data around longer," usually by anonymizing data to increase the length of time it can legally be stored. A research scientist for a privacy-as-a-service startup noted that clients were interested in "the ability

to monetize data more properly and outside the purview... while still complying with regulations" (P17). Regulations were most commonly cited as the main constraint against keeping old data or collecting new data, but some participants cited brand safety, too—as a manager at a tech non-profit put it, "protecting the brand is just as much a liability thing as protecting you from creating a compliance disaster where you have to pay millions of dollars in fines" (P10).

For twenty-one of our participants (including all the non-profit and government employees), PPA also helped to legally leverage *new* data sources for existing operations. One privacy engineer at a large tech company said: "We are trying to make sure that we're doing things on the correct side of the law while getting ourselves some sort of benefit...for example, being able to collect data we wouldn't otherwise collect..." (P3). Some participants in private industry viewed this strategy for adoption as a "cloak for just collecting lots and lots of data" (P22) or "legal cover for hoovering up as much information as possible" (P15). Organizations in both private and public organizations tended to "collect the data first and then figure out what it's useful for later" (P25).

Making a business out of privacy: At the same time, organizations—especially startups—also welcomed PPA as an opportunity for *new* business or a competitive edge. Four of our participants (including one at a non-profit) framed privacy-preserving analytics as an opportunity to access new, privacy-conscious markets—and seven (all for-profit tech companies but one) were motivated by the prospect of more competitive marketing. A director said:

> *Before... we would appeal to the principles of the company in order for something to happen. Now... you can start saying things like, 'It probably will appeal to this market,' or, 'This segment actually has expressed an interest in this kind of thing.' (P10)*

Another manager said: "privacy is a business proposition... that customers are happy and willing to pay for" (P3). For example, web browsers like DuckDuckGo and Brave have grown small market shares around privacy-preserving branding. Other startups like Tumult Labs or Leap Year Technologies, recently acquired by the data cloud giant Snowflake, aim to offer PPA directly as a

service to businesses, non-profits, and government agencies.

For non-profit organizations in particular, PPA also offered an opportunity to open up new sources of data to researchers and other third parties, particularly for statistical agencies and other public-interest organizations. For example, the first use of differential privacy at the U.S. Census Bureau was the OnTheMap tool (Machanavajjhala et al., 2008), which made available new data on commuting patterns in the U.S. The Wikimedia Foundation, in collaboration with Tumult Labs, released additional Wikipedia usage statistics with differential privacy (Adeleye et al., 2023).

Improving data management: Still, PPA development was often costly and lengthy. Many participants in the private sector felt the need to ship products quickly to maintain competitive advantage—but PPA adoption threatened to slow down their businesses through additional privacy design review or infrastructure development. To combat the view that adoption was "costing resources with no clear benefits" (P17), some practitioners argued that PPA would serve internal business purposes other than privacy compliance. One employee of a startup offering compliance services said, "I've seen cases where if [a PPA proposal] is solely a privacy use case and not integrated with the overall data analytics security use cases, then it actually fails" (P13). Several practitioners argued internally that the improvements to data management and data quality required to adopt PPA would also reduce "bad science," and ultimately make for better products. Discussions about PPA adoption often offered practitioners a chance to bring up older, long-standing issues like data minimization: "it gives you the ability to talk about that like it's a fresh thing" (P22).

4.2.2 Deciding to adopt

How exactly did these narratives drive the decision to adopt? PPA practitioners built these narratives as they attempted to persuade executives, product teams, lawyers, and privacy practitioners that PPA adoption was a relevant solution for their organization. As a result, an organization's adoption and deployment of PPA is not merely strategic—it is shaped and modified by a distributed network of internal and external stakeholders and their concerns. In particular, while the drivers of adoption named in the previous section may be primarily economic or political, privacy-preserving analytics

are also inherently moral: the decision to adopt is often tied up with employees' moral judgments. Here, we describe two key internal stakeholders: the decision-makers, usually executives, who decide whether to prioritize PPA adoption; and the privacy practitioners who advocate for adoption internally. (We discuss other internal stakeholders further in §4.3.)

Getting buy-in from executives: Across organizations, the primary audience for our participants' adoption narratives were the key decision-makers within the organization controlling the resources necessary to develop and deploy a new PPA system. An employee of a privacy compliance startup observed that whether a client adopted their services depended on "how strong their chief privacy officer is or deputy officer, how influential the legal and risk teams are and the executive board as well, the CEO, and the board members in prioritizing this effort" (P13). And an employee of a government agency described how having support from leadership helped them close off the possibility of returning to previous privacy practices: "Your own leadership has said, 'we are doing this'—going back to [prior practices] is not an option, so let's make it work" (P6).

Twelve participants described pitching PPA projects to higher-ups as part of their roles, appealing to the external privacy pressures from the previous section. In meetings with executives, one privacy engineer at a large tech company made it a habit to "go in armed with a bucket of consumer research and case studies," including internal studies aimed at estimating "how much bad privacy can potentially cost you, based on historical data" (P3). Which narrative and other arguments practitioners employed depended on the audience. One director said, "Trying to understand what kind of executive stakeholder you're dealing with is like a dirty secret" (P10).

Executive buy-in helped push forward adoption on a case-by-case basis and also drove the formation of new internal policies. For example, one large tech company in our sample integrated PPA adoption into its normal privacy review process. While participants at government agencies relied mostly on hierarchy to spread adoption, participants at other organizations discussed the importance of a more general "privacy culture," especially for companies that relied on product teams to "self-forward" (P8) relevant cases or take on adoption independently. Ten participants—mostly at

large tech companies—reported efforts to foster this culture, including pair mentorship or extra training during onboarding. These efforts were not always successful—a manager at a large tech company emphasized the lack of incentives for privacy work: "Privacy is defensive, and the entire industry is terrible respecting or rewarding consistently defensive work... It takes expert eyes to be able to vet the value of the work, so it feels riskier for people to take on" (P23). Only one participant mentioned specific plans to reverse these incentives—such as having executives highlight privacy as an area of importance or providing a "laundry list" of promotion-eligible privacy work.

Advocating internally: Privacy practitioners' efforts towards PPA adoption often stemmed not only from their job duties, but also from their personal ethics. Six of our participants—all at for-profit tech companies and start-ups—said they considered adopting PPA because it was the "right thing" to do for users. One privacy engineer at a tech company said:

> *I think at any company from little to big you're going to find that there's some set of people who are genuinely deeply ethical people—and I have met many of those at [my organization] and they're fantastic to work with—and there are people who recognize that privacy is a business proposition... (P3)*

Whatever their motivation, practitioners leveraged formal structures like privacy reviews to persuade engineers or resist opposing pressure from executives: "everything goes through privacy, so we can just tell the leadership 'Hey, we're not letting this launch unless you do this'" (P11). For a few participants, all at for-profit firms, this stance engendered an adversarial relationship with others in the organization. From a former employee of a large tech company:

> *The inherent relationship between the people on the ground doing the privacy reviews and the leadership, even within the privacy org, is very confrontational... the [product] team is trying to do the least amount of privacy that we will approve... You are fighting against different teams and organizations and you might as well act like it. (P19)*

Still, privacy engineers tried to be constructive: "If people have a positive experience with privacy,

they are more likely to come to us when they feel there's a problem" (P26). Internal advocacy also depended on the leadership of influential practitioners. Two participants at large tech companies observed practitioners "pushing hard" for particular PPA techniques—one of whom a former employee said "probably is uniquely responsible for driving adoption in the company" (P19).

But these methods only went so far to spread adoption internally. One participant said that when an influential leader left, internal adoption of PPA dwindled. At least two of the organizations in our sample were taking steps to deconstruct these central teams, distributing privacy professionals to product and infrastructure teams across the organization. Some felt there was a limit to adoption: "Once there's real money on the line, you get leadership involved. And someone doesn't care about protecting privacy because they're going to get promoted if you make however many billion dollars" (P19). In that engineer's view, it took a "push from outside"—in their case, a court ruling—to get their organization to invest in privacy work.

4.3 Interpreting drivers into designs

How do these external and internal drivers of adoption inform choices made during deployment? After all, none of our participants described a regulation or public campaign that called for specific PPA technologies—instead, organizations used the narratives described below to rationalize PPA as an appropriate response to general privacy expectations. And adopting PPA—and negotiating its design—is not as simple as choosing a vendor or product off the shelf. Most of the participants we talked to shipped their own PPA system or had to make significant effort to integrate external technical infrastructure or PPA products. During this process, organizations made a number of specific design decisions—and choices about how to define and measure privacy—to make their analytics techniques "privacy-preserving." PPA systems are thus constructed by organizations as a managerial solution to institutional expectations.

While engineers in product teams may think about privacy rarely and never interact with a privacy professional, our participants—PPA practitioners—contributed to the design of PPA systems as part of their primary roles. First, practitioners interpreted the (usually legal) privacy requirements that

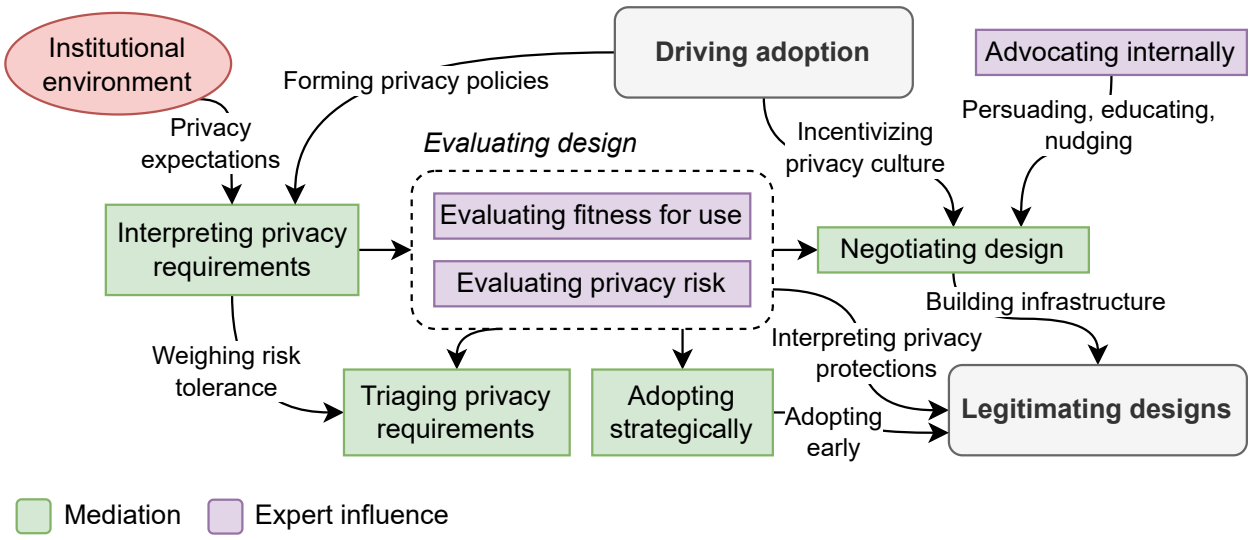


Figure 4: Interpreting drivers into designs. Key processes associated with managerial mediation (and possibly decoupling), institutional endogeneity, or expert influence are highlighted.

systems must satisfy and evaluated trade-offs in proposed designs. Then, practitioners iteratively negotiated with product teams, lawyers, and executives to triage privacy requirements, choose uses cases to tackle, and settle on an appropriate solution. Figure 4 depicts the relationships between these key interpretation processes that emerged from our second-order analysis.

Interpreting privacy requirements: PPA adoption was often motivated by external privacy expectations, but those expectations were broader than the technical specificity of PPA system design. In private firms and public organizations alike, a key task for PPA practitioners was to "interpret" or "translate" consumer expectations, regulatory requirements, and organizational standards into technical specifications for a new PPA system.

Twenty-three participants across organization types described interpreting or translating legal requirements as part of their work. From a former director at a tech company:

> *The way [the organization] did GDPR was we had this privacy legal department like they'd spent a huge amount of time with the law... We had them dump the entire law into my head and I wrote the engineering requirements! (P7)*

Though specific technical standards are not written in current law, practitioners could sometimes rely on regulatory guidance—from the European Courts, for example (Data Protection Working Party, 2014, e.g.). Likewise, one employee of a large tech company described receiving "hints" from regulators about which practices would be considered unacceptable. Other participants interpreted more specific company policies determined by executives or lawyers: "These regulations would get translated into internal policies, and so what people inside the company care about are the internal policies... I very rarely had to care about what the actual regulations were" (P19).

Privacy requirements did not always come from articulated laws or policies. For example, one policy researcher at a large tech company was responsible for understanding the "key opinion leaders" in policy and regulation: "What are their expectations, and what do we need to do to meet them?" (P12) Less commonly, participants mentioned requirements related to user privacy expectations. While academic researchers have explored user expectations for PPA (Cummings et al., 2021), most of our participants did not have direct contact with data subjects.

The process of interpreting "external mandates," as one participant called them, was not straightforward. Fourteen participants perceived conflicting, deficient or "unreasonable" expectations in external regulations. For practitioners who view their work as ahead of regulation, this was especially frustrating. One executive at a tech company lamented, "Regulations come in and kind of break what I'm doing... you [regulators] just made our system work worse and I'm very grumpy about it" (P7). Eleven complained about inconsistencies in privacy laws across domains, and at least one created a "pecking order" (P23) of privacy rules to follow. Some even pointed out specific technical errors in regulatory guidance, speculating that "people who knew how this stuff really worked or could work weren't necessarily at the table" (P26) when the regulation was written.

Triaging privacy requirements: No single PPA system or design met all privacy requirements, especially for the largest organizations we studied. Instead, practitioners weighed the privacy risks against their organization's tolerance for risk. As one executive put it, "you want to ask yourself what risk are you willing to take, how much uncertainly are you willing to live with" (P14)? Another

manager at a large tech company recounted,

> *When I first started working on this, I accepted the culture of 'Oh gosh, the sky is falling, it's all important, we've got to get it all!'... It turns out that we have permission to fail smaller risks... [executives] are fine with us taking misses... because they can't imagine a future in which we don't take another fine... (P23)*

A privacy engineer noted that risk tolerances varied even within their company:

> *[In design meetings] there's usually a representative from the business and their job is to sit there and advocate for the continued health of the business... I think the lawyers and I are both more risk averse and the business is the counterweight to that. (P3)*

Which protections are prioritized depends on both the costs of privacy harms to the organization—usually fines or loss of consumer trust—as well as practitioners' estimations of privacy risk.

Evaluating privacy risk: Of course, simply evaluating the reduction of privacy risk provided by a particular design involves more expert interpretation. Legal requirements that data be "anonymous," for example, have been satisficed by successively stronger technical standards in just the past decade, including the use of k-anonymity (Malin et al., 2011) to satisfy HIPAA standards and more recently the use of differential privacy in the 2020 Decennial Census (Abowd et al., 2022).

But twelve of our participants (especially those from government agencies and non-profits) mentioned difficulties translating theoretical privacy guarantees to practice. An engineer recounted, "The sorts of guarantees that a privacy-enhancing technology offers almost never line up with anything a lawyer would recognize... and so we wind up in a room with a whiteboard sort of scribbling frantically at each other trying to do a lingo match" (P3). Five participants even observed misinterpretations that would compromise privacy guarantees in practice. Twelve reported difficulty judging the sensitivity of personal information, and ten debated with stakeholders over whether new protections were even necessary. Four fell back on "experimental" or "practical" guarantees, rather

than formal methods, to convince stakeholders of privacy risk.

Without a clear conception of privacy risk, tuning the strength of privacy protections was more art than science. As one privacy engineer at a large tech company put it,

> *There's nothing out there to guide you, so you're just kind of winging it. We've gotten as far as lining [differential privacy] up with things like GDPR's singling out clause, but it still doesn't give us any notion of how to do things like tuning parameters. (P3)*

In practice, various stakeholders—including legal counsel interested in shielding the organization from liability, or external partners with a vested interest in the data—often interpreted privacy risk differently. For example, Social Science One, an association of scholars partnering with Facebook to release a large research dataset with differential privacy, argued that less stringent standards were enough to satisfy the General Data Protection Regulation and Facebook's FTC consent decree, but Facebook disagreed (King & Persily, 2020).

Negotiating design: These interpretations—of privacy requirements and associated risks—guided design discussions between privacy practitioners and the teams responsible for implementing the PPA system. Most of our participants described a process of negotiation with representatives of internal (and sometimes external) stakeholders—primarily in roles related to legal liability or organizational performance (financial performance or product quality, in the case of a firm).

These negotiations usually centered managerial and technical concerns—with exceptions, our participants framed the design process as maximizing fitness for use, subject to a minimum "privacy bar." When asked what trade-offs they considered in PPA design, nearly all of our participants discussed compromising between some conception of privacy and other desirable properties (what some of our participants call "fitness for use"), including economic cost (both costs of development and opportunity costs), speed and efficiency, user friendliness, customer satisfaction, interoperability, granularity, freedom of future applications, equity between data users, and the ability to audit for discrimination. Notably, only one of our participants admitted that utility trade-offs were

unavoidable (P19)—practitioners, especially at organizations with large volumes of data and expertise, often found ways to work around particularly uncomfortable sacrifices either with process improvements or by changing product requirements.

Adopting strategically: In addition to deciding on design parameters, higher-level managers and executives also decided on where and when to apply PPA techniques. Given the limited time and resources available for PPA adoption, most of the organizations we studied did not apply PPA uniformly across products and features—in fact, many adopted PPA for only one or two. The rollout strategy depended on how privacy requirements were triaged—for example, larger organizations had systems for prioritizing data deemed more "sensitive" or risky. Though most executives and higher-level managers agreed that it was better to start adoption early—to promote communication and reduce disruption—they disagreed on whether it was better to target mission-critical products first (to set a precedent) or to start with "easy" use cases (to build momentum). Five privacy engineers, all in private industry, said they experienced mostly the latter—adoption for only peripheral use cases. One noted that organizations "don't use the distributed machine learning approaches in things that are really mission critical... Where it really matters, they just collect a bunch of data [centrally] and make some promises around it" (P15).

4.4 Legitimating designs

The previous sections describe organizations' motivations to adopt—including external privacy expectations—and the ways practitioners' interpretation of those motivations affected design and adoption strategy. However, our analysis also revealed ways that PPA adoption influences the institutional environment in turn (Figure 5). To satisfy the privacy expectations that drove adoption, organizations needed to convince regulators, consumer advocates, and other audiences that their design choices were acceptable. To do so, organizations—relying on privacy experts—first had to decide how to interpret and represent the privacy properties of their PPA designs (§4.4.1). Organizations prepared these interpretations to legally justify PPA adoption to regulators (e.g., to preserve existing operations) and to make privacy promises to consumers (e.g., to compete in

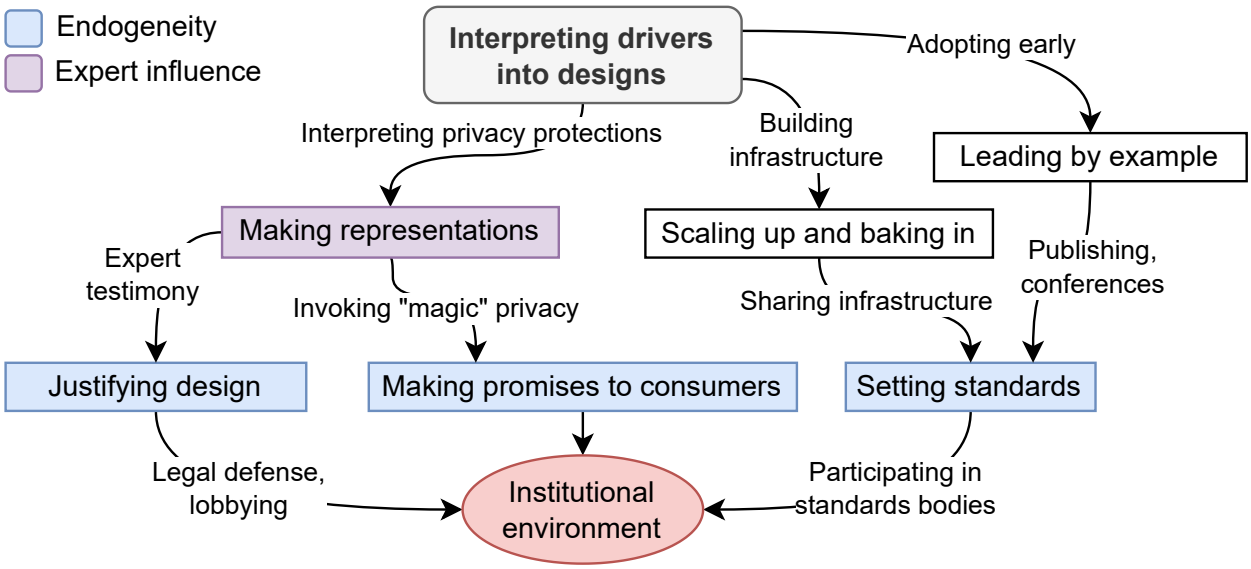


Figure 5: Legitimating design. Key processes associated with managerial mediation (and possibly decoupling), institutional endogeneity, or expert influence are highlighted.

new markets). Organizations also scaled up their PPA practices, setting standards internally and promoting those standards externally to the rest of the industry (§4.4.2). By making representations and setting standards, organizations both legitimated design choices and created a precedent for future responses to institutional pressure.

4.4.1 Making representations

All the organizations in our sample who deployed a PPA system translated the properties of that system into some sort of (often broad) external representation.

Justifying design: Practitioners—particularly those who adopted to preserve existing operations—often framed the suitability of a PPA proposal in terms of eventual legal justification or "defensibility." As a policy researcher put it, "Can you back that [public statement] up with your systems in an investigation?" (P12) One engineer said,

> *At some point the lawyer and I just kind of have to sit down and be like... 'This seems both ethically justifiable and probably reasonable in the eyes of the law based on some esoteric U.K. law from the 1600s...' or 'We think this is defensible and we think it's not*

a crappy thing to do.' (P3)

Sometimes justification had less to do with specific law and more to do with perceptions of the regulatory climate. As one privacy engineer recalled telling internal lawyers,

> '*We don't give legal opinions, we're technical people. But, this new thing is the gold standard. If a regulator says anonymization is a thing that's possible, and academia says this is the best thing you can hope for—the gold standard—then, by definition, if you use this probably you're going to be fine.' (P18)*

Practitioners were prepared to rely heavily on their own expertise. A researcher at another large tech company defended their new technique internally:

> *Ultimately we had a very, very senior statistician who basically just got up and told the lawyers, 'I believe that risk of leakage is very low in this model' and that's it. They didn't go into the math, they didn't look at any of the other stuff. (P20)*

For lawyers and policy researchers, justification was also a political matter. Several practitioners described steps their organizations (two large tech companies and a privacy startup) took to influence regulation and develop relationships with policymakers—as that director put it, "glad handing" (P2)—outside the normal course of fact-sharing. Describing their strategy for justifying their privacy practices to regulators, one privacy director said: "Regulators are political creatures just by definition... They have their own policy preferences. They care about outcomes. They're not automatons who apply the law" (P2). Another manager put it more bluntly, describing meetings with regulators who "just looked like they were there to throw a wrench at the company" (P23). These practitioners used PPA adoption to help make their practices more palatable to policymakers.

Making promises to consumers: In addition to relying on PPA to satisfy privacy laws, organizations also relied on PPA to fulfill privacy promises to consumers. Sixteen participants related PPA design choices to public representations their organizations made about privacy practices (e.g.,

public privacy policies), particularly when adoption aimed to preserve existing operations. Most participants, though, did not have direct contact with data subjects—they explained their work to legal and policy teams internally. Those that did described the difficulty of accurately representing PPA protections to laypeople. One participant in our study said, "You're asking [customers] to trust an algorithm they will never understand. It's not that easy to prove to someone that these things are actually going to work" (P25). Eight participants said they preferred to choose simpler PPA techniques to avoid misinterpretation, particularly for privacy-branded products.

Others were less concerned about the need to comprehensively educate consumers.

> *Privacy-preserving technologies [are] almost like whiz-bang technology... it doesn't have to be something that everybody needs a detailed awareness of, because it just makes life easier in the background. (P12)*

Another privacy director at a tech company pointed out a practical upside to operating PPA "in the background": companies would no longer have to ask for consent and risk "creep[ing] people out" (P7). At least three participants—in both government and industry—worried that marketing new PPA protections would "muddy the waters" (P21) by surfacing existing data practices.

Six of our participants, mostly in private industry, feared uses of PPA that amounted to "magic privacy" or privacy "pixie dust"—black box techniques that, when invoked in marketing or legal copy, assure consumers and regulators of strong privacy protections without any further inspection. An executive at a tech company distinguished "things that I can make a strong representation about and the things I can't" (P7), depending on the technology, parameter settings, use case, and other factors. Later, they put it more bluntly: "most of this marketing is bullshit... they're writing checks that their tech cannot cash" (P7). The policy researcher who advocated for less detailed awareness also advocated for clearly disclosing the limitations of PPA: "Technologists always have to be really careful... because you can make this sound so much more impressive. You can make it seem like snake oil" (P12). Six other participants—who mostly worked to preserve existing products in at

for-profit firms—brought up the possibility that their PPA efforts were just "good theater" (P28) or "privacy whitewashing" (P12). An employee at a regulatory agency distinguished "adoption for show" from adoption that changes "the broader field of analytics," referring to industry proposals that—since they do not affect first-party tracking—could just be a "side story, instead of moving privacy along holistically" (P5). Some participants mentioned Google's Privacy Sandbox as an example of this kind of proposal.

Another way organizations sought to justify and legitimate PPA design was by increasing transparency and inviting stakeholders to participate. The U.S. Census Bureau, for example, submitted their differentially private 2020 disclosure avoidance system to multiple rounds of public comment, adjusting parameters and post-processing along the way, and commissioned reviews from organizations like JASON and Mitre (boyd & Sarathy, 2022). Nine participants, especially in government, reported documenting their PPA methods publicly. A handful, all but one in industry, even developed their PPA technologies as open source platforms. Seven reported making changes in response to feedback from the public, especially at non-profit organizations. Even when code was public, though, our participants and other researchers found it difficult to judge how exactly that code is being used from the outside, and key privacy parameters and product details were not always disclosed (Dwork et al., 2019; Gong, 2022).

4.4.2 Setting standards

Another way organizations influenced the institutional environment was through technical standards. While organizations may benefit from influencing external, industry standards, organizations also developed standards as a practical necessity internally. In our sample, sixteen participants mentioned setting internal "rules" based on their organization's anticipation of regulation or industry standards, especially for cases where adoption preserved an existing business model. Standards have several uses for PPA practitioners. One is to push adoption internally, "irrespective of if [product teams] want to or not" (P11). One privacy engineer at a large tech company said, "The [product] team would come in and say okay well how about this epsilon. So a lot of the time, we can really easily

just point to, 'No we have a standard, it's at this epsilon, use it'" (P19).

Theoretically, standards also help practitioners ensure PPA design is tightly coupled to internal and external expectations. But practitioners disagreed on what designs and parameter settings are truly "privacy-preserving" for any given use case. Some practitioners in our sample believed that DP is a "gold standard" approach for managing privacy risk that provides "meaningful" privacy if implemented properly, but disagreed over whether a PPA standard is still meaningful after certain relaxations or under certain parameter choices. One privacy engineer said,

> *A lot of what I've heard is people who want to use differential privacy with more relaxed standards to not quite count as anonymous, which sounds nice. But... that means that in the worst case, you could be releasing a lot of information, if you have a very high epsilon. And that's not meaningful [privacy], at least in my opinion. (P19)*

Others doubted further whether differential privacy is even an appropriate technical conception of privacy (for a representative view, see Hotz et al., 2022; Seeman & Susser, 2022). Even among those who believe differential privacy is an appropriate standard for a majority of cases, the appropriate parameter settings for various application was context-dependent and subjective.

As a result, standards often seemed arbitrary. As one manager put it, "a lot of this we've just had to come up with out of thin air" (P23). Another put it even more bluntly: "The definition of what we consider to be anonymous is completely arbitrary. There's some legal things informing that but, for the most part we just made up ours" (P19). For setting the k parameter in k-anonymity, for example, a privacy engineer at a privacy-focused startup said, "We have some magic numbers in the company, like 20. $k = 20$ is like a minimum that we don't go under. But what we usually ask for as like we started the discussion from like $k = 1000$ or $k = 100$ at least and then go down from there" (P8). A lawyer for a tech company described their "magic number" for k-anonymity similarly: "we kind of feel like [a k value of] 20 is usually not super necessary, but if it gets down to under 5 it's kind of a little dicey" (P28). Not all organizations had "magic numbers." Other, larger organizations had

teams of privacy experts who helped product teams set parameters on a case-by-case basis.

Twenty-one of our participants—nearly all, and especially those not at privacy-focused startups—mentioned the need for better standards for PPA systems. At least four, all at for-profit tech companies, felt that regulators did not provide enough guidance. From a lawyer at a tech company: "I guess I've been taken aback a lot in the private sector [at] how out in the cold companies feel—like they want to do something, but they just don't know what is required" (P28).

Leading by example: Standards may be developed to serve internal needs, but organizations also aimed to affect standards in industry and policy. Some internal standard-setting merely attempted to forestall future regulation. One privacy engineer at a large tech company said:

> *An organization like [mine] really would like to avoid getting to court for every little thing. You need some kind of consistent internal standard... so that you can avoid getting into situations where you are in a public sense told, 'You have to do this.' (P19)*

Six practitioners at tech companies, startups, and non-profits spent time anticipating (sometimes "speculating" about) and preempting future regulation:"let's try to fix this two years before they make it mandatory" (P3). In particular, a few participants felt that waiting for new regulations left them short on time to implement—scrambling to catch up to new regulations impairs organizational efficiency. As one manager at a large tech company said,"We're either going to eat a big penalty for the process inefficiency or we're occasionally going to eat penalties by sending engineers after something that didn't turn out to make it into law" (P23).

Other standard setting actively aimed to actively *influence* future regulation. Practitioners—particularly at privacy-focused startups—explicitly to be a model for future regulations. Executives at two privacy boutiques said, "We're eager to demonstrate that legislation is catching up to [us], instead of [us] catching up to legislation" (P15); "We want to be seen as thought leaders in this area as it continues to evolve" (P13). As one privacy engineer at a large tech company said:

> *Imagine we end up in a state where when we say... this is the bar, this is what we consider to be acceptable. The first big tech company to be able to do that is actually in a leadership position... (P18)*

For example, some practitioners in legal and policy teams at large private companies advocated to policymakers for regulation that would provide "safe harbor" (P2) from certain regulatory requirements for organizations who implemented certain PPA techniques, and a policy researcher at an international tech company described efforts to "put [PPA] on the radar" of countries in the Middle East and Asia as they wrote their first privacy laws (P12). Another way organizations performed this kind of "standardization work"—as one participant described it—was to work with standards bodies such as the International Organization for Standardization (ISO), the National Institute of Standards and Technology (NIST) or the World Wide Web Consortium (W3C). Five participants, all at for-profit tech companies and startups, mentioned joining standards organizations with the goal of guiding future PPA deployments.

Standard-setting and early adoption also influenced other organizations. Seven participants—disproportionately from non-profits and smaller organizations—modeled their PPA systems after others' prominent deployments. For example, multiple participants mentioned watching the U.S. Census Bureau's deployment of differential privacy to see how it will change statistical practices in government or change regulatory expectations for private organizations. An engineer at one internet non-profit wished that other companies would publish guidance on how they protect different types of user data so that professionals could develop more shared norms (P23).

Scaling up and baking in: Standardization is also a byproduct of organizations' desire to "scale up" PPA systems within their organizations. Thirteen participants, mostly at smaller organizations, mentioned the practical difficulties of scaling PPA systems and privacy review processes to handle dispersed troves of complex personal data and wider varieties of use cases, often due to a lack of engineering talent, labor and computing resources.

Many of our participants—particularly those at tech companies and non-profits—turned to tooling to help amplify their limited expertise and limited resources. Sixteen participants, especially those using PPA to create new products, mentioned building or updating tools for developers—mostly software libraries and internal platforms—as part of their PPA development work; 22 described building or updating software infrastructure in the course of adoption. One manager said, "We bake [PPA] into infrastructure. We make it so it can't be screwed up" (P23). Four participants likewise aimed to construct a "well-lit path" for developers—as one executive at a tech company explained, "You want to make it easy to do the right thing and hard to do the wrong thing" (P7).

By virtue of their market positions, some companies also had the ability to lay a well-lit path for developers across industry by sharing software infrastructure. Twelve participants, disproportionately from smaller non-profit and government organizations, mentioned integrating external PPA software libraries or other infrastructure—built by large tech companies like IBM or by open source communities such as OpenDP–to deploy their own PPA systems, though most deployments still required extensive customization. Moreover, cloud service companies such as Databricks have included PPA tooling and features in their data analytics products (Zaharia et al., 2023).

Combined with external "standardization work", these efforts may allow large tech companies to simultaneously avoid regulation and deny certain forms of data use to their competitors. Five of our industry practitioners noted that PPA infrastructure can help firms box out competitors. Several mentioned Apple's new App Tracking Transparency feature, an ostensibly privacy-motivated feature that allows users to opt out of one form of tracking used by ad brokers including Meta and Google (Morrison, 2022), and Google's Privacy Sandbox. Both projects drew anti-competition criticism from advertisers in France and the United Kingdom (Lomas, 2021). As one PPA researcher at a large tech company observed, "When [corporations] make a bid on a privacy-preserving technology, it's not just that they want to do it quietly. They want to do it spectacularly with regulations that ensconce what they've done at the expense of their competitors" (P4).

5 THEORETICAL INTEGRATION & IMPLICATIONS

The emergent process model described in the previous sections elucidates the institutional view that 1) organizations have latitude to negotiate the design of PPA systems to accommodate managerial concerns and 2) organizations aim to justify these choices to regulators and consumers and normalize them in industry. Our process model aligns with theorized stages of managerial mediation and legal endogeneity (Edelman, 2016, p. 27-41): 1) organizations encounter ambiguous law; 2) compliance professionals frame the legal environment (here, by interpreting privacy requirements); 3) symbolic forms of compliance (here, the privacy "pixie dust" some practitioners observed—spread across industry); 4) organizations create proxy "company laws" (here, privacy rules and other internal policies); 5) organizations mobilize symbolic structures (PPA) to avoid liability (e.g., legally storing data for longer); 6) courts, legislatures, and administrative agencies endorse organizations' symbolic compliance practices. It is not yet clear the extent to which PPA systems receive legal deference, but practitioners expected it—one engineer said, "A lot of [adoption] is not so much 'this is what the law says' as 'differential privacy seems to make the regulators happy'" (P19). Some organizations hoped to codify deference to PPA—tech companies such as Google are already advocating for regulatory exceptions for the PPA techniques they use (Google, 2022).

Our emergent process model also reveals mediators unique to PPA. PPA systems may 1) increase the ease with which policies are decoupled from practice through an additional dimension we call *policy-technology* decoupling, 2) provide unique avenues for organizational practice to influence the institutional environment, and 3) shift organizational practice in turn (§5.1). We also highlight two mediators of policy-technology decoupling revealed in our process model: the countervailing influence of morally-motivated practitioners and the narratives driving adoption (§5.2).

5.1 Policy-technology decoupling and technological endogeneity

Technological infrastructure—the code bases and cloud services that industry "runs on"—is inextricably linked with practice (L. Suchman et al., 1999). Just as routines executed by people mediate the impacts of institutional pressures, so can routines executed in code (Lessig, 1999).

Policy-technology decoupling: Technology may increase the likelihood and permanence of decoupling and managerial mediation. Prior work documents how the average employee forms their own interpretation of their organization's legal structure and mobilizes that "legal reading" in their every day work (Fuller et al., 2000). Our analysis shows that PPA experts also employed their own *technological* "readings" in their attempts to represent privacy protections in social or legal terms—in other words, they frame not only the legal environment, but also the technological environment. As Wu (2003, p. 682) writes, "The programmer is not unlike the tax lawyer, exploiting differences between stated goals of the law, and its legal or practical limits."

In PPA systems, these legal and technological readings are especially consequential. Particularly for recently developed techniques, modern privacy laws rarely admit to straightforward, specific translations to technical implementation (Nissim & Wood, 2018). Differential privacy, for example involves parameters and variations that are notoriously difficult to interpret (Nanayakkara et al., 2023). There is little consensus on key parameter choices that determine the "meaningful degree of privacy" provided by a differentially private mechanism (Dwork et al., 2019), and there are heated epistemic disagreements over its value (boyd & Sarathy, 2022). Because of this uncertainty, important technological practices affecting digital privacy may be easily obfuscated, subject to less scrutiny, and therefore more easily decoupled from formal policy. We call this gap between formal policy and technological design *policy-technology decoupling*—a special form of decoupling arising from the gray area between a formal policy (here, an organizations' external representations) and its technological manifestation. Research on social performance often frames the adoption of technologies like the catalytic converter as categorically beneficial (Ashford et al., 1985); policy-technology decoupling models the very plausible possibility that technology adoption is contextual and variable, and could therefore be decoupled from policy.

Our findings do not suggest that PPA designs are ubiquitously decoupled from institutional privacy expectations. In our setting, the distinction between substantive and merely symbolic adoption is context-dependent and non-binary—PPA technologies encompass several definitions of privacy that

may be preserved or neglected to varying degrees. But our model of PPA adoption does includes several mechanisms by which this may easily happen—most notably the processes of interpretation and evaluation required to negotiate PPA design. Information technologies generally are layers of digital devices, networks, and services that scale through translation and loose coupling (Faik et al., 2020; Berente & Yoo, 2012). The relative invisibility of algorithms, as well as the legally incognizable harms they cause, can make information technologies difficult to interrogate in court (Metcalf et al., 2023; Selbst et al., 2023). For example, a privacy impact assessment requirement for federal agencies arguably fell short in part due to "the inaccessible idiom of technology" (K. A. Bamberger & Mulligan, 2008). For this reason, policy-technology decoupling may be more durable than typical policy-practice decoupling, which some argue does not withstand public scrutiny for long (Bromley & Powell, 2012; Haack et al., 2012).

Technological endogeneity: Rational myths about acceptable privacy practice are not the only means by which initial "guesses" at compliance endogenously alter industry practices and institutional expectations (Edelman et al., 1999). As technological artifacts, PPA systems represent material standards and guidelines for the ethical treatment of personal data (Verbeek, 2006; Lampland & Star, 2009). They encourage technological, in addition to organizational, isomorphism (DiMaggio & Powell, 1983). For example, developers' tendency to rely on only a few entrenched software libraries means that early adopters have even more influence over practices in the rest of industry (Gururaja et al., 2023), as our analysis corroborates. And algorithmic technologies in particular receive additional legal deference as courts often treat their design as logistical inevitability rather than question the design choices that lead to their creation (Selbst et al., 2023).

Technological—and particularly algorithmic—artifacts, then, have a special potential for lasting influence on the institutional environment (Lessig, 1999). As PPA infrastructure is shared within and across organizations through open source libraries or cloud computing services, the early instantiations of PPA—usually created by the best-resourced organizations—may close off alternative formulations of privacy. For example, nearly all of our participants framed privacy risks as invasions

or exploitation by state actors or other "attackers," shifting the focus away from own organization (Seeman & Susser, 2022). And they focused almost entirely on privacy for individuals rather than privacy with respect to information flow or social relations more generally (Viljoen et al., 2021; Barocas & Levy, 2020)—some argue the technical property of anonymity, the focus of many PPA techniques, has been used to bypass privacy rather than protect it (Barocas & Nissenbaum, 2014). Insights from Microsoft's differentially private Workplace Analytics tool (Bird, 2020), for example, may still be used to increase managerial control and restrict workers' autonomy (Levy, 2022).

Technology shifts organizational behavior in turn: Once adoption is underway, technological requirements also shift organizational values and routines in turn, as demonstrated in research from information systems (Leonardi, 2011; Davidson & Chismar, 2007; Scott & Orlikowski, 2022) as well as technology studies (Winner, 1980; L. Suchman et al., 1999) and organizational adaptation (Rossetti & DeZoort, 1989; S. Y. Park et al., 2015). In our setting, interviews with practitioners (Dwork et al., 2019; Garrido et al., 2022; Sarathy et al., 2023), a 2022 workshop (Cummings et al., 2023), and reflections on the 2020 Decennial Census (Hawes, 2020; boyd & Sarathy, 2022) all elaborate on the ways that the use of differential privacy changes organizational processes surrounding data analysis. These studies provide early evidence not only that perceived usefulness and ease of use are crucial to convincing engineers to adopt PPA technologies (F. D. Davis, 1989), but also that organizational practice—and organizational values—adapts to PPA technology as adoption takes place (Verbeek, 2006). The use of local processing techniques, for example, implicitly frames individuals' devices as resources to be mined in the process of private computation, rarely providing them an opportunity to opt out (Veale, 2023).

As PPA designs are similarly oriented towards surveillance and the extraction of value from data (Cohen, 2019; Zuboff, 2019), scholars may worry that the ripple effects of PPA adoption may be at fundamentally odds with values like democracy and social equality (Waldman, 2021; Viljoen, 2021a; Viljoen, 2021b). Intense focus on scale and growth—evident in our analysis—involves reduction and inequality (Hanna & Park, 2020) and may conflict with other values such as justice by

overlooking context, relations, and individuality (Hasinoff & Schneider, 2022). As one participant put it, PPA "can't make [data collection] ethical just because it makes it private" (P15).

5.2 Mediators of policy-technology decoupling

Prior work suggests several determinants of policy-practice decoupling, including the stage of adoption and the extent to which adoption is motivated by social performance rather than operational demands (Bromley & Powell, 2012). Here, we identify two additional variables that influence the extent and direction of policy-technology decoupling.

5.2.1 Privacy champions

Several components of our process model are explicitly mediated by managerial interests—for example, the negotiation of trade-offs between privacy parameters and operational concerns. But our study also identifies a force in favor of tighter technology-practice coupling: the moral leadership displayed by PPA practitioners themselves.

PPA technologies require expertise to interpret: Prior work suggests policy-practice decoupling is more likely when internal constituents fail to reinforce external pressures (Bromley & Powell, 2012; Edelman, 2016). Our process model suggests the influence of internal experts was especially salient—even dominant—in PPA adoption. First, organizations and courts rely on performances of expertise to judge algorithmic systems (Metcalf et al., 2023), as our analysis corroborates. Second, the process of interpreting drivers into designs was heavily dependent on PPA experts—often only one or two individuals. Eleven of our participants—especially at non-profits and start-ups—mentioned lack of expertise and experience as a barrier to PPA adoption and eight consulted with external experts. One employee at a large non-profit in tech said that their PPA deployment may not have "gotten off the ground" (P19) without the expertise of an external consultant. Similar to "street-level bureacrats" (Lipsky, 1980)—public service workers whose individual decisions collectively add up to government policy—privacy practitioners on the ground effectively define the practical limits of PPA technologies that likewise bridge the gap between policy and practice (Alkhatib & Bernstein, 2019).

PPA professionals leverage their expertise to tightly couple policy and practice: In our study, some practitioners sought to exercise this latitude to endogenize and promote PPA adoption for their own reasons in an act of institutional entrepreneurship (DiMaggio, 1988). Recent work chronicles the role of privacy "champions": individual practitioners who advance privacy through informal education and daily processes when official policies are missing or insufficient (Tahaei et al., 2021). Privacy champions are tempered radicals: an individual committed to their organization but are also committed to the cause of data privacy, possibly in opposition to the reigning culture in their organization (Meyerson & Scully, 1995). Many of our participants or their colleagues were privacy champions. They promoted PPA adoption through informal mechanisms—particularly technological mechanisms, such as the "well-lit paths" of privacy by default—and belonged to "special interest" groups, usually centralized privacy teams.

PPA adoption may be accomplished through moral leadership: These privacy champions led efforts to fully integrate and meaningfully implement PPA technology, building coalitions and safeguarding adoption routines (Solinger et al., 2020); but, they also struggled to prioritize, incentivize, and organize around ethics in corporate environments, as do software engineers more generally (Widder et al., 2023; Ali et al., 2023; H.-P. Lee et al., 2024). The approaches these participants took to internally advocate for adoption mirror the process of emergent moral leadership (Solinger et al., 2020), including a precipitating moral scandal, internal coalition building, negotiation with other stakeholders, and the establishment and maintenance of formal structures.

5.2.2 Motivational narratives

Our analysis suggests that policy-technology decoupling is also mediated by the three narratives motivating organizations to adopt PPA: preserving data-driven operations (the most common), making a business out of privacy, and (less frequently) improving data management. Cross-organization analysis of our (small) sample suggests that 1) decoupling may be more difficult for organizations making a business out of privacy, and 2) legal and technological endogeneity is stronger for larger organizations preserving existing data-driven operations.

Decoupling is more difficult for organizations making a business out of privacy: Our participants reported decoupling between the stated aims of PPA adoption and the technical specifics most commonly in organizations that adopted to preserve existing data-driven operations and in private industry—disproportionately large tech companies. For example, the strategy of adopting only for peripheral, less "mission critical" use cases (§4.2) was reported disproportionately by practitioners in private industry, as were concerns about deceptive invocations of "magic privacy" techniques (§4.4). PPA adoption may then be a case of "difficult decoupling"—when employees' identities are tied up with their organizations' external marketing, decoupling is less likely to succeed (Turco, 2012). Indeed, our analysis suggests that privacy-motivated practitioners at privacy-branded startups were more positive about PPA standards at their companies. Meanwhile, at large tech companies adopting mostly to comply with existing regulations, employee-led adoption efforts are more likely to falter as key privacy leaders left the company and the company broke up central privacy groups.

Legal and technological endogeneity is stronger in large organizations: Likewise, in our study, institutionally endogenous processes such as making representations or influencing industry standards were disproportionately discussed by practitioners at large, for-profit companies. For example, large tech companies used representations and standardization work to convince regulators that applying PPA to their existing practices was sufficient to comply with privacy laws. Organizational theory suggests these industry leaders have an advantage in influencing the institutional environment—both technology adoption (Rogers, 2003) and legitimacy (DiMaggio & Powell, 1983; Campbell, 2007) spread through networks of influence.

5.3 Boundary conditions

In our study, we analyze a very specific context: the adoption of privacy-preserving analytics. Besides providing thick descriptions of important phenomena like PPA adoption (A. S. Lee & Baskerville, 2003), unconventional contexts can be useful for developing newly insightful theory (P. A. Bamberger & Pratt, 2010; Monteiro et al., 2022). We chose to study PPA adoption not only because of its potential impact on digital privacy, but also because it represents an understudied

intersection of technological, managerial, and societal concerns common in a broader class of information systems technologies transforming the technology industry. We believe our model applies particularly to technological practices that 1) are socially motivated (or morally weighted), 2) involve routines enacted by machines, in addition to people, and 3) have properties or impacts that require technological expertise to understand. Algorithmic systems in particular fit these criteria—pushes towards "responsible" artificial intelligence (AI) and ethical data science, for example, are motivated by similar institutional pressures to PPA adoption (Hirsch et al., 2020; H.-P. Lee et al., 2024), especially as AI regulation matures. Thus, many of the mechanisms we identify in this unconventional context—for example, policy-technology decoupling—could also help explain the interaction between regulation and technology adoption in fields such as artificial intelligence, environmental protection, or digital social innovation (Qureshi et al., 2021).

6 PRACTICAL IMPLICATIONS

Given our findings, policymakers and privacy-minded managers face a conundrum. Within existing privacy laws, there is room for technological interpretation, as our analysis explicates. But, the capacity to form and propagate these interpretations is concentrated in mostly large and mostly private organizations. These early adopters hold particular sway over the techniques that may become synonymous with privacy compliance and social performance. One practitioner, for example, optimistically described a future of "automatically invokable" mathematical techniques that would assure users their privacy is preserved. Of course, it may be that the PPA systems deployed today *are* a significant, substantive step, and many of our participants were quite optimistic about their organizations' efforts. But left to self-regulate, prominent early adopters may exercise their influence in ways that are contrary to the public interest.

One remedy to managerial mediation is closer scrutiny (Edelman, 2016)—mechanisms that keep designs tightly coupled to institutional expectations. Researchers and managers should investigate not only the theoretical properties of PPA techniques but also their empirical manifestations and impacts. Prior work offers guidelines for accountability in algorithmic systems (Kaminski, 2019;

Selbst, 2021), as do areas of law where technologies are routinely deconstructed—products liability, for example (Selbst et al., 2023). And independent academic researchers are already scrutinizing the most visible PPA systems deployed by organizations such as Apple (Tang et al., 2017) and Google (Berke & Calacci, 2022). Some of our participants—echoed by differential privacy scholars (Cummings et al., 2023; Gong, 2022)—called for additional efforts to make PPA techniques more transparent and usable, such as a registry of PPA parameter choices (Dwork et al., 2019).

Two contemporary cases illustrate the benefits and limits of public scrutiny. In their updated 2020 disclosure avoidance system, the U.S. Census Bureau adapted existing differentially private mechanisms into a novel algorithm that preserves utility for certain critical statistics includes exceptions for certain statutory requirements such as redistricting (Abowd et al., 2022). This design resulted from years of consultation and efforts at transparency, which facilitated heated debate about whether the Bureau's design properly balanced confidentiality with other policy goals (Hotz et al., 2022). On the other hand, one of Google's publicly documented proposals, Federated Learning of Cohorts (FLoC)—a plan to automatically cluster users into interest groups based on their behavior—drew opposition from competing browsers Brave and Vivaldi (Bohn, 2021) and from the Electronic Frontier Foundation, who called FLoC "the opposite of privacy-preserving technology" (Cyphers, 2019). In response, Google eventually replaced FLoC with Topics API, which groups users into fewer and purportedly less sensitive clusters (Goel, 2022). Here, scrutiny helped to overturn an unpalatable design on technical grounds. But the Topics API still gives advertisers the power to target users based on their behavior, leaving the core business model relatively untouched and hampering third party competition (Lomas, 2021).

Scrutiny and transparency are not an unalloyed goods—they must be accompanied by formal mechanisms for accountability (Han, 2015; Birchall, 2014). Judges, policymakers, investigators, and managers tasked with enforcing regulations should look beyond the invocation of a technique like differential privacy to examine its full technical manifestation—including parameters, definitions of sensitivity, pre- and post-processing steps. Privacy and algorithm impact assessments optimistically

could, in addition to providing a formal structure for internal advocates, raise public awareness and provide valuable information to outside observers (Selbst, 2021).

In addition to scrutiny, policymakers could offer additional resources, including guidance or development funding that could help other organizations contribute to PPA practice. But even as policymakers seek to advance the development PPA technologies (e.g., Office of Science and Technology Policy, 2022; Fast-Track Action Committee on Advancing Privacy-Preserving Data Sharing and Analytics et al., 2023), they should also take care not to endorse entire classes of PPA techniques symbolically or outside of context. Regulations should also avoid the assumption that a technological, formal privacy solution is the appropriate solution (Green & Viljoen, 2020) or that "magic privacy" systems even function as advertised (Raji et al., 2022).

7 CONCLUSION, LIMITATIONS, AND FUTURE RESEARCH

Our study analyzes the experiences of the expert practitioners who help organizations decide to adopt, interpret PPA designs, and justify designs to external stakeholders. Through this unique perspective, we develop a theoretical model that captures interactions between institutional privacy expectations and the design of PPA technologies—interactions that have great implications for digital privacy. We intersect work on information systems, organizational decoupling, and legal endogeneity, highlighting salient new technological dimensions to decoupling (policy-technology decoupling) and legal endogeneity.

Our work has limitations. First, because PPA techniques originate mostly in computer science, mathematics, and statistics this study focuses on a particular set of formal, technical, and algorithmic approaches to privacy. Likewise, our analysis is influenced by our own backgrounds in computer science, economics, and public policy (see Appendix A for further reflection on the perspectives that shaped our analysis). Second, while our sample did include organizations who considered PPA techniques but failed to deploy them, our sample does not include organizations which have never seriously considered adoption. Third, our participants' responses may be prone to social desirability bias or crafted to serve other professional motives, though we do not repeat their

views uncritically—we constantly compare data within and between participants and organizations. Finally, our sample—like the pool of PPA practitioners at large—is predominantly American, white, and cisgender male. As a result, our study examines the culturally dominant perspectives in PPA work, but by no means do those perspectives encompass all possible conceptions of PPA—for example, our participants' views of PPA rarely included documented inequalities in access to privacy (Skinner-Thompson, 2018) or the impacts of surveillance and datafication for Black and low-income communities (Allen, 2022; Gilman & Green, 2018–2019; Madden et al., 2017–2018). Further research could also examine how PPA might develop to account for other perspectives—for example, the particularly salient privacy concerns of Black and low-income communities .

Other directions of future research include the nascent drive towards more usable, transparent tooling could help reduce the barriers to interpreting PPA systems. User studies of differential privacy practitioners, for example, highlight several gaps in differential privacy tooling that further research could help resolve (Garrido et al., 2022; Sarathy et al., 2023). Research from disciplines outside computer science and statistics could further elucidate the social trade-offs facing practitioners and the downstream consequences of adoption. And researchers and policymakers could imagine alternative PPA technologies that elevate public interests—for example, as a means to facilitate algorithm auditing (Xu & Zhang, 2021).

Large tech companies are promoting PPA techniques as the future of online advertising and the digital economy (Egan, 2020). We hope our study will inspire further rigorous empirical investigation and theorization about the ways regulators and consumer advocates may shape the development of these and other socially motivated technologies.

ACKNOWLEDGEMENTS

We thank our participants for their time and insight. Thanks also to those who provided feedback on earlier versions of this research: Taya Cohen, Anna Mayo, Roy Rinberg, seminar participants at Carnegie Mellon University, and participants at the the 2023 USENIX Conference on Privacy Engineering Practice and Respect and the 2023 Privacy Law Scholars Conference—including

Nikita Aggarwal, Jason Cronk, Elizabeth Edenberg, Thomas Haley, Cameron Kerry, Anne Klinefelter, Siona Listokin, Scott Skinner-Thompson, Jeremy Seeman, Alexis Shore, Harry Surden, and Alexandra Wood. This study was approved by the Institutional Review Board at Carnegie Mellon University in 2021, #00000327.

Funding: Both authors are supported by the Heinz College of Information Systems and Public Policy at Mellon University. Ryan Steed also receives support from a Meta PhD Research Fellowship. This work is not subject to pre-publication review or veto, and only the authors control its content.

REFERENCES

Abadi, M., Chu, A., Goodfellow, I., McMahan, H. B., Mironov, I., Talwar, K., & Zhang, L. (2016). Deep Learning with Differential Privacy. *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, 308–318.

Abowd, J. M., Ashmead, R., Cumings-Menon, R., Garfinkel, S., Heineck, M., Heiss, C., Johns, R., Kifer, D., Leclerc, P., Machanavajjhala, A., Moran, B., Sexton, W., Spence, M., & Zhuravlev, P. (2022). The 2020 Census Disclosure Avoidance System TopDown Algorithm. *Harvard Data Science Review*.

Abowd, J. M., & Hawes, M. B. (2023). Confidentiality Protection in the 2020 US Census of Population and Housing. *Annual Review of Statistics and Its Application*, *10*(1), 119–144.

Abowd, J. M., & Schmutte, I. M. (2019). An Economic Analysis of Privacy Protection and Statistical Accuracy as Social Choices. *American Economic Review*, *109*(1), 171–202.

Acquisti, A., & Steed, R. (2023). Learning to Live with Privacy-Preserving Analytics. *Communications of the ACM*, *66*(7), 24–27.

Adeleye, T., Berghel, S., Desfontaines, D., Hay, M., Johnson, I., Lemoisson, C., Machanavajjhala, A., Magerlein, T., Modena, G., Pujol, D., Simmons-Marengo, D., & Triedman, H. (2023). *Publishing Wikipedia usage data with strong privacy guarantees*. arXiv: 2308.16298 [cs].

Aguilera, R. V., Rupp, D. E., Williams, C. A., & Ganapathi, J. (2007). Putting the S Back in Corporate Social Responsibility: A Multilevel Theory of Social Change in Organizations. *The Academy of Management Review*, *32*(3), 836–863.

Ali, S. J., Christin, A., Smart, A., & Katila, R. (2023). Walking the Walk of AI Ethics: Organizational Challenges and the Individualization of Risk among Ethics Entrepreneurs. *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, 217–226.

Alkhatib, A., & Bernstein, M. (2019). Street-Level Algorithms: A Theory at the Gaps Between Policy and Decisions. *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, 1–13.

Allen, A. (2022). Dismantling the "Black Opticon": Privacy, Race Equity, and Online Data-Protection Reform. *Yale Law Journal Forum*.

Ashford, N. A., Ayers, C., & Stone, R. F. (1985). Using Regulation to Change the Market for Innovation. *Harvard Environmental Law Review*, *9*(2), 419–466
Accepted: 2002-08-05T20:06:51Z.

Bamberger, K. A. (2010). Technologies of Compliance: Risk and Regulation in a Digital Age. *Texas Law Review*, *88*(4), 669–739.

Bamberger, K. A., & Mulligan, D. K. (2008). Privacy Decisionmaking in Administrative Agencies. *The University of Chicago Law Review*, *75*(1), 75–107.

Bamberger, K. A., & Mulligan, D. K. (2015). *Privacy on the ground: Driving corporate behavior in the United States and Europe*. The MIT Press.

Bamberger, P. A., & Pratt, M. G. (2010). Moving forward by looking back: Reclaiming unconventional research contexts and samples in organizational scholarship. *Academy of Management Journal*, *53*(4), 665–671.

Barocas, S., & Levy, K. (2020). Privacy Dependencies. *Washington Law Review*, *95*(2), 555–616.

Barocas, S., & Nissenbaum, H. (2014). Big Data's End Run around Anonymity and Consent. In H. Nissenbaum, J. Lane, S. Bender, & V. Stodden (Eds.), *Privacy, Big Data, and the Public Good: Frameworks for Engagement* (pp. 44–75). Cambridge University Press.

Berente, N., & Yoo, Y. (2012). Institutional Contradictions and Loose Coupling: Postimplementation of NASA's Enterprise Information System. *Information Systems Research*, *23*(2), 376–396.

Berke, A., & Calacci, D. (2022). Privacy Limitations of Interest-based Advertising on The Web: A Post-mortem Empirical Analysis of Google's FLoC. *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*, 337–349.

Birchall, C. (2014). Radical Transparency? *Cultural Studies ↔ Critical Methodologies*, *14*(1), 77–88.

Bird, S. (2020). *Putting differential privacy into practice to use data responsibly*. AI for Business. https://blogs.microsoft.com/ai-for-business/differential-privacy/

Bohn, D. (2021). *Nobody is flying to join Google's FLoC*. The Verge. https://www.theverge.com/2021/4/16/22387492/google-floc-ad-tech-privacy-browsers-brave-vivaldi-edge-mozilla-chrome-safari

Boldosova, V. (2019). Deliberate storytelling in big data analytics adoption. *Information Systems Journal*, *29*(6), 1126–1152.

boyd, d., & Sarathy, J. (2022). Differential Perspectives: Epistemic Disconnects Surrounding the U.S. Census Bureau's Use of Differential Privacy. *Harvard Data Science Review*.

Bozanic, Z., Dirsmith, M. W., & Huddart, S. (2012). The social constitution of regulation: The endogenization of insider trading laws. *Accounting, Organizations and Society*, *37*(7), 461–481.

Bromley, P., & Powell, W. W. (2012). From Smoke and Mirrors to Walking the Talk: Decoupling in the Contemporary World. *The Academy of Management Annals*, *6*(1), 483–530.

Campbell, J. L. (2007). Why Would Corporations Behave in Socially Responsible Ways? An Institutional Theory of Corporate Social Responsibility. *The Academy of Management Review*, *32*(3), 946–967.

Cavazos, D. E., & Rutherford, M. A. (2012). Bringing Regulatory Agencies Into Organizational Studies: Broadening the Lens Used to Examine the State. *Journal of Management Inquiry*, *21*(1), 4–12.

Charmaz, K. (2014). *Constructing grounded theory* (2nd edition). Sage. OCLC: ocn878133162.

Cohen, J. E. (2018). The Biopolitical Public Domain: The Legal Construction of the Surveillance Economy. *Philosophy & Technology*, *31*(2), 213–233.

Cohen, J. E. (2019). *Between Truth and Power: The Legal Constructions of Informational Capitalism*. Oxford University Press.

Confessore, N. (2018). Cambridge Analytica and Facebook: The Scandal and the Fallout So Far [newspaper]. *The New York Times: U.S.*

Cummings, R., Desfontaines, D., Evans, D., Geambasu, R., Jagielski, M., Huang, Y., Kairouz, P., Kamath, G., Oh, S., Ohrimenko, O., Papernot, N., Rogers, R., Shen, M., Song, S., Su, W., Terzis, A., Thakurta, A., Vassilvitskii, S., Wang, Y.-X., … Zhang, W. (2023). *Challenges towards the Next Frontier in Privacy*. arXiv: 2304.06929 [cs].

Cummings, R., Kaptchuk, G., & Redmiles, E. M. (2021). "I need a better description": An Investigation Into User Expectations For Differential Privacy. *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security*, 3037–3052.

Cyphers, B. (2019). *Don't Play in Google's Privacy Sandbox*. Electronic Frontier Foundation. https://www.eff.org/deeplinks/2019/08/dont-play-googles-privacy-sandbox-1

Cyphers, B. (2021). *Google's FLoC Is a Terrible Idea*. Electronic Frontier Foundation. https://www.eff.org/deeplinks/2021/03/googles-floc-terrible-idea

Data Protection Working Party. (2014). *Opinion 05/2014 on Anonymisation Techniques* (Opinion No. WP216).

Davidson, E. J., & Chismar, W. G. (2007). The Interaction of Institutionally Triggered and Technology-Triggered Social Structure Change: An Investigation of Computerized Physician Order Entry. *MIS Quarterly*, *31*(4), 739–758.

Davis, F. D. (1989). Perceived Usefulness, Perceived Ease of Use, and User Acceptance of Information Technology. *MIS Quarterly*, *13*(3), 319–340.

Davis, K. (1973). The Case for and Against Business Assumption of Social Responsibilities. *Academy of Management Journal*, *16*(2), 312–322.

Desfontaines, D. (2021). *A list of real-world uses of differential privacy*. Ted is writing things. https://desfontain.es/privacy/real-world-differential-privacy.html

DiMaggio, P. J. (1988). Interest and agency in institutional theory. In L. G. Zucker (Ed.), *Research on Institutional Patterns*. Ballinger Publishing Co.

DiMaggio, P. J., & Powell, W. W. (1983). The Iron Cage Revisited: Institutional Isomorphism and Collective Rationality in Organizational Fields. *American Sociological Review*, *48*(2), 147–160.

Domingo-Ferrer, J., & Blanco-Justicia, A. (2020). Privacy-Preserving Technologies. In M. Christen, B. Gordijn, & M. Loi (Eds.), *The Ethics of Cybersecurity* (pp. 279–298, Vol. 21). Springer International Publishing.

Dwork, C., Kohli, N., & Mulligan, D. (2019). Differential Privacy in Practice: Expose your Epsilons! *Journal of Privacy and Confidentiality*, *9*(2).

Dwork, C., McSherry, F., Nissim, K., & Smith, A. (2006). Calibrating Noise to Sensitivity in Private Data Analysis. In S. Halevi & T. Rabin (Eds.), *Theory of Cryptography* (pp. 265–284). Springer.

Edelman, L. B. (1992). Legal Ambiguity and Symbolic Structures: Organizational Mediation of Civil Rights Law. *American Journal of Sociology*, *97*(6), 1531–1576.

Edelman, L. B. (2016). *Working Law: Courts, Corporations, and Symbolic Civil Rights*. University of Chicago Press.

Edelman, L. B., Uggen, C., & Erlanger, H. S. (1999). The Endogeneity of Legal Regulation: Grievance Procedures as Rational Myth. *American Journal of Sociology*, *105*(2), 406–454.

Egan, E. (2020). *A Path Forward for Privacy and Online Advertising*. Meta. https://about.fb.com/news/2020/10/a-path-forward-for-privacy-and-online-advertising/

Eisenhardt, K. M. (1989). Building Theories from Case Study Research. *The Academy of Management Review*, *14*(4), 532–550.

Faik, I., Barrett, M., & Oborn, E. (2020). How Information Technology Matters in Societal Change: An Affordance-Based Institutional Perspective. *MIS Quarterly*, *44*(3), 1359–1390.

Fast-Track Action Committee on Advancing Privacy-Preserving Data Sharing and Analytics, Networking and Information Technology Research and Development Subcommittee, & National Science and Technology Council. (2023). *National Strategy to Advance Privacy-Preserving Data Sharing and Analytics*. Executive Office of the President.

Fuller, S. R., Edelman, L. B., & Matusik, S. F. (2000). Legal Readings: Employee Interpretation and Mobilization of Law. *The Academy of Management Review*, *25*(1), 200–216.

Garrido, G. M., Liu, X., Matthes, F., & Song, D. (2022). *Lessons Learned: Surveying the Practicality of Differential Privacy in the Industry*. arXiv: 2211.03898 [cs].

Gentry, C. (2009). Fully homomorphic encryption using ideal lattices. *Proceedings of the Forty-First Annual ACM Symposium on Theory of Computing*, 169–178.

Gilman, M., & Green, R. (2018–2019). The Surveillance Gap: The Harms of Extreme Privacy and Data Marginalization. *New York University Review of Law & Social Change*, *42*(2), 253–308.

Gioia, D. A., Corley, K. G., & Hamilton, A. L. (2013). Seeking Qualitative Rigor in Inductive Research: Notes on the Gioia Methodology. *Organizational Research Methods*, *16*(1), 15–31.

Goel, V. (2022). *Get to know the new Topics API for Privacy Sandbox*. Google. https://blog.google/products/chrome/get-know-new-topics-api-privacy-sandbox/

Goldberg, I. (2007). Privacy-Enhancing Technologies for the Internet III: Ten Years Later. In A. Acquisti, S. Gritzalis, C. Lambrinoudakis, & S. di Vimercati (Eds.), *Digital Privacy* (0th ed., pp. 25–40). Auerbach Publications.

Goldreich, O. (2009). *Foundations of Cryptography: Volume 2, Basic Applications* (1st ed.). Cambridge University Press.

Gong, R. (2022). Transparent Privacy Is Principled Privacy. *Harvard Data Science Review*.

Google. (2022). *Request for Information on Advancing Privacy-Enhancing Technologies* (87 Fed. Reg. 35250 No. 2022-12432).

Green, B., & Viljoen, S. (2020). Algorithmic realism: Expanding the boundaries of algorithmic thought. *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 19–31.

Gunningham, N., Kagan, R. A., & Thornton, D. (2004). Social License and Environmental Protection: Why Businesses Go beyond Compliance. *Law & Social Inquiry*, *29*(2), 307–341.

Gururaja, S., Bertsch, A., Na, C., Widder, D. G., & Strubell, E. (2023). *To Build Our Future, We Must Know Our Past: Contextualizing Paradigm Shifts in Natural Language Processing*. arXiv: 2310.07715 [cs].

Haack, P., Schoeneborn, D., & Wickert, C. (2012). Talking the Talk, Moral Entrapment, Creeping Commitment? Exploring Narrative Dynamics in Corporate Responsibility Standardization. *Organization Studies*, *33*(5-6), 815–845.

Han, B.-C. (2015). *The Transparency Society*. Stanford University Press.

Hanna, A., & Park, T. M. (2020). *Against Scale: Provocations and Resistances to Scale Thinking*. arXiv: 2010.08850 [cs].

Haraway, D. (1988). Situated Knowledges: The Science Question in Feminism and the Privilege of Partial Perspective. *Feminist Studies*, *14*(3), 575–599.

Hasinoff, A. A., & Schneider, N. (2022). From Scalability to Subsidiarity in Addressing Online Harm. *Social Media + Society*, *8*(3), 20563051221126041.

Hawes, M. B. (2020). Implementing Differential Privacy: Seven Lessons From the 2020 United States Census. *Harvard Data Science Review*, *2*(2).

Hillman, A. J., Keim, G. D., & Schuler, D. (2004). Corporate Political Activity: A Review and Research Agenda. *Journal of Management*, *30*(6), 837–857.

Hirsch, D. D., Bartley, T., Chandrasekaran, A., Norris, D., Parthasarathy, S., & Turner, P. N. (2020). *Business Data Ethics: Emerging Trends in the Governance of Advanced Analytics and AI*.

Hotz, V. J., Bollinger, C. R., Komarova, T., Manski, C. F., Moffitt, R. A., Nekipelov, D., Sojourner, A., & Spencer, B. D. (2022). Balancing data privacy and usability in the federal statistical system. *Proceedings of the National Academy of Sciences*, *119*(31), e2104906119.

Jones, T. M. (1995). Instrumental Stakeholder Theory: A Synthesis of Ethics and Economics. *The Academy of Management Review*, *20*(2), 404–437.

Kagan, R. A., Thornton, D., & Gunningham, N. (2003). Explaining Corporate Environmental Performance: How Does Regulation Matter? *Law & Society Review*, *37*(1), 51–90.

Kallio, H., Pietilä, A.-M., Johnson, M., & Kangasniemi, M. (2016). Systematic methodological review: Developing a framework for a qualitative semi-structured interview guide. *Journal of Advanced Nursing*, *72*(12), 2954–2965.

Kamieniecki, S. (2006). *Corporate America and Environmental Policy: How Often Does Business Get Its Way?* Stanford University Press.

Kaminski, M. (2019). Binary Governance: Lessons from the GDPR's Approach to Algorithmic Accountability. *Southern California Law Review*.

King, G., & Persily, N. (2020). *Unprecedented Facebook URLs Dataset now Available for Academic Research through Social Science One*. Social Science One. https://socialscience.one/blog/unprecedented-facebook-urls-dataset-now-available-research-through-social-science-one

Lampland, M., & Star, S. L. (2009). *Standards and Their Stories: How Quantifying, Classifying, and Formalizing Practices Shape Everyday Life*. Cornell University Press.

Lee, A. S., & Baskerville, R. L. (2003). Generalizing Generalizability in Information Systems Research. *Information Systems Research*, *14*(3), 221–243.

Lee, H.-P., Gao, L., Yang, S., Forlizzi, J., & Das, S. (2024). "I Don't Know If We're Doing Good. I Don't Know If We're Doing Bad": Investigating How Practitioners Scope, Motivate, and Conduct Privacy Work When Developing AI Products. *USENIX Security Symposium (Forthcoming)*.

Leidner, D. E., Sutanto, J., & Goutas, L. (2022). Multifarious Roles and Conflicts on an Interorganizational Green Is. *MIS Quarterly*, *46*(1), 591–608.

Leonardi, P. M. (2011). When Flexible Routines Meet Flexible Technologies: Affordance, Constraint, and the Imbrication of Human and Material Agencies. *MIS Quarterly*, *35*(1), 147–167.

Lessig, L. (1999). *Code and Other Laws of Cyberspace*. Basic Books, Inc.

Levy, K. (2022). *Data Driven: Truckers, Technology, and the New Workplace Surveillance*. Princeton University Press.

Liang, H., Saraf, N., Hu, Q., & Xue, Y. (2007). Assimilation of Enterprise Systems: The Effect of Institutional Pressures and the Mediating Role of Top Management. *MIS Quarterly*, *31*(1), 59–87.

Lipsky, M. (1980). *Street Level Bureaucracy: Dilemmas of the Individual in Public Services*. Russell Sage Foundation.

Lomas, N. (2021). *France's competition authority declines to block Apple's opt-in consent for iOS app tracking*. TechCrunch. https://techcrunch.com/2021/03/17/frances-competition-authority-declines-to-block-apples-opt-in-consent-for-ios-app-tracking/

Machanavajjhala, A., Kifer, D., Abowd, A., John M, Gehrke, J., & Vilhuber, L. (2008). Privacy: Theory meets Practice on the Map. *2008 IEEE 24th International Conference on Data Engineering*, 277–286.

Madden, M., Gilman, M., Levy, K., & Marwick, A. (2017–2018). Privacy, Poverty, and Big Data: A Matrix of Vulnerabilities for Poor Americans. *Washington University Law Review*, *95*(1), 53–126.

Malin, B., Benitez, K., & Masys, D. (2011). Never too old for anonymity: A statistical standard for demographic data sharing via the HIPAA Privacy Rule. *Journal of the American Medical Informatics Association*, *18*(1), 3–10.

Matthews, G. J., & Harel, O. (2011). Data confidentiality: A review of methods for statistical disclosure limitation and methods for assessing privacy. *Statistics Surveys*, *5*, 1–29.

McMahan, B., Moore, E., Ramage, D., Hampson, S., & Arcas, B. A. y. (2017). Communication-Efficient Learning of Deep Networks from Decentralized Data. *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, 1273–1282.

Metcalf, J., Singh, R., Moss, E., Tafesse, E., & Watkins, E. A. (2023). Taking Algorithms to Courts: A Relational Approach to Algorithmic Accountability. *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, 1450–1462.

Meyer, J. W., & Rowan, B. (1977). Institutionalized Organizations: Formal Structure as Myth and Ceremony. *American Journal of Sociology*, *83*(2), 340–363.

Meyerson, D. E., & Scully, M. A. (1995). Crossroads Tempered Radicalism and the Politics of Ambivalence and Change. *Organization Science*, *6*(5), 585–600.

Monteiro, E., Constantinides, P., Scott, S., Shaikh, M., & Burton-Jones, A. (2022). Editor's Comments: Qualitative Research Methods in Information Systems: A Call for Phenomenon-Focused Problematization. *Management Information Systems Quarterly*, *46*(4), iii–xix.

Morrison, S. (2022). *The winners and losers of Apple's anti-tracking feature*. Vox. https://www.vox.com/recode/23045136/apple-app-tracking-transparency-privacy-ads

Nader, L. (1972). *Up the Anthropologist: Perspectives Gained From Studying Up*. ERIC Number: ED065375.

Nanayakkara, P., Smart, M. A., Cummings, R., Kaptchuk, G., & Redmiles, E. (2023). *What Are the Chances? Explaining the Epsilon Parameter in Differential Privacy*. arXiv: 2303.00738 [cs].

Nguyen, G. (2022). *Google's Topics API: Advertisers share concerns about topic diversity and other potential challenges*. Search Engine Land. https://searchengineland.com/googles-topics-api-advertisers-share-concerns-about-topic-diversity-and-other-potential-challenges-379247

Nissim, K., & Wood, A. (2018). Is privacy privacy? *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, *376*(2128), 20170358.

Office of Science and Technology Policy. (2022). Request for Information on Advancing Privacy-Enhancing Technologies. *Federal Register*.

Oliver, C. (1991). Strategic Responses to Institutional Processes. *The Academy of Management Review*, *16*(1), 145–179.

Orlitzky, M., & Benjamin, J. D. (2001). Corporate Social Performance and Firm Risk: A Meta-Analytic Review. *Business & Society*, *40*(4), 369–396.

Park, S., & Cha, H. (2019). Institutional decoupling and the limited implementation of certified environmental technologies. *Journal of Environmental Management*, *247*, 253–262.

Park, S. Y., Chen, Y., & Rudkin, S. (2015). Technological and Organizational Adaptation of EMR Implementation in an Emergency Department. *ACM Transactions on Computer-Human Interaction*, *22*(1), 1:1–1:24.

Pollach, I. (2011). Online privacy as a corporate social responsibility: An empirical study. *Business Ethics: A European Review*, *20*(1), 88–102.

Pollman, E., & Barry, J. M. (2016). Regulatory Entrepreneurship. *Southern California Law Review*, *90*(3), 383–448.

Qureshi, I., Pan, S. L., & Zheng, Y. (2021). Digital social innovation: An overview and research framework. *Information Systems Journal*, *31*(5), 647–671.

Raji, I. D., Kumar, I. E., Horowitz, A., & Selbst, A. (2022). The Fallacy of AI Functionality. *2022 ACM Conference on Fairness, Accountability, and Transparency*, 959–972.

Rogers, E. M. (2003). *Diffusion of Innovations, 5th Edition* (5th edition). Free Press.

Rossetti, D. K., & DeZoort, F. A. (1989). Organizational Adaptation To Technology. *S.A.M. Advanced Management Journal*, *54*(4), 29.

Saldaña, J. (2013). *The coding manual for qualitative researchers* (2nd ed). SAGE. OCLC: ocn796279115.

Sarathy, J., Song, S., Haque, A., Schlatter, T., & Vadhan, S. (2023). *Don't Look at the Data! How Differential Privacy Reconfigures the Practices of Data Science*. arXiv: 2302.11775 [cs].

Scott, S., & Orlikowski, W. (2022). The Digital Undertow: How the Corollary Effects of Digital Transformation Affect Industry Standards. *Information Systems Research*, *33*(1), 311–336.

Seeman, J., & Susser, D. (2022). *Between Privacy and Utility: On Differential Privacy in Theory and Practice*.

Selbst, A. D. (2021). An Institutional View Of Algorithmic Impact Assessments. *Harvard Journal of Law & Technology*, *35*(117).

Selbst, A. D., Venkatasubramanian, S., & Kumar, I. E. (2023). Deconstructing Design Decisions: Why Courts Must Interrogate Machine Learning and Other Technologies. *Ohio State Law Journal*, *85*, forthcoming.

Silva, L., & Hirschheim, R. (2007). Fighting against Windmills: Strategic Information Systems and Organizational Deep Structures. *MIS Quarterly*, *31*(2), 327–354.

Skinner-Thompson, S. (2018). Privacy's Double Standards. *Washington Law Review*, *93*(4), 2051–2106.

Soghoian, C. (2011). An End to Privacy Theater: Exposing and Discouraging Corporate Disclosure of User Data to the Government. *Minnesota Journal of Law, Science & Technology*, *12*(1), 191–238.

Solinger, O. N., Jansen, P. G., & Cornelissen, J. P. (2020). The Emergence of Moral Leadership. *Academy of Management Review*, *45*(3), 504–527.

Spradley, J. P. (1979). *The Ethnographic Interview*. Holt, Rinehart and Winston.

Strike, V. M., & Rerup, C. (2016). Mediated Sensemaking. *Academy of Management Journal*, *59*(3), 880–905.

Suchman, L., Blomberg, J., Orr, J. E., & Trigg, R. (1999). Reconstructing Technologies as Social Practice. *American Behavioral Scientist*, *43*(3), 392–408.

Suchman, M. C. (1995). Managing Legitimacy: Strategic and Institutional Approaches. *The Academy of Management Review*, *20*(3), 571–610.

Sutton, J. R., & Dobbin, F. (1996). The Two Faces of Governance: Responses to Legal Uncertainty in U.S. Firms, 1955 to 1985. *American Sociological Review*, *61*(5), 794–811.

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*, 1–15.

Tang, J., Korolova, A., Bai, X., Wang, X., & Wang, X. (2017). *Privacy Loss in Apple's Implementation of Differential Privacy on MacOS 10.12*. arXiv: 1709.02753 [cs].

Turco, C. (2012). Difficult Decoupling: Employee Resistance to the Commercialization of Personal Settings. *American Journal of Sociology*, *118*(2), 380–419.

U.S. Equal Employment Opportunity Commission. (2016). *Diversity in High Tech* (Special Report).

Veale, M. (2023). *Rights for Those Who Unwillingly, Unknowingly and Unidentifiably Compute!*

Verbeek, P.-P. (2006). Materializing Morality: Design Ethics and Technological Mediation. *Science, Technology, & Human Values*, *31*(3), 361–380.

Viljoen, S. (2021a). The Promise and Limits of Lawfulness: Inequality, Law, and the Techlash. *Journal of Social Computing*, *2*(3), 284–296.

Viljoen, S. (2021b). A Relational Theory of Data Governance. *Yale Law Journal*, *131*(2), 573–655.

Viljoen, S., Goldenfein, J., & McGuigan, L. (2021). Design choices: Mechanism design and platform capitalism. *Big Data & Society*, *8*(2), 20539517211034312.

Waldman, A. E. (2017). *Designing Without Privacy*. https://papers.ssrn.com/abstract=2944185

Waldman, A. E. (2021). *Privacy, Practice, and Performance*.

Weaver, G. R., Treviño, L. K., & Cochran, P. L. (1999). Integrated and Decoupled Corporate Social Performance: Management Commitments, External Pressures, and Corporate Ethics Practices. *The Academy of Management Journal*, *42*(5), 539–552.

Weber, M. (1978). *Economy and Society: An Outline of Interpretive Sociology*. University of California Press.

Weick, K. E. (1976). Educational Organizations as Loosely Coupled Systems. *Administrative Science Quarterly*, *21*(1), 1–19.

Widder, D. G., Zhen, D., Dabbish, L., & Herbsleb, J. (2023). It's about power: What ethical concerns do software engineers have, and what do they (feel they can) do about them? *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, 467–479.

Wiesche, M., Jurisch, M., Yetton, P., & Krcmar, H. (2017). Grounded Theory Methodology in Information Systems Research. *Management Information Systems Quarterly*, *41*(3), 685–701.

Willis, G. B., & Artino, A. R. (2013). What Do Our Respondents Think We're Asking? Using Cognitive Interviewing to Improve Medical Education Surveys. *Journal of Graduate Medical Education*, *5*(3), 353–356.

Winner, L. (1980). Do Artifacts Have Politics? *Daedalus*, *109*(1), 121–136.

Wu, T. (2003). When Code Isn't Law. *Va. L. Rev.*, *89*, 679.

Xu, H., & Zhang, N. (2021). Implications of Data Anonymization on the Statistical Evidence of Disparity. *Management Science*.

Zaharia, M., Li, A., Efeoglu, C., Simeone, C., Thakur, S., & Tenedorio, D. (2023). *Introducing Lakehouse Federation Capabilities in Unity Catalog*. Databricks. https://www.databricks.com/blog/introducing-lakehouse-federation-capabilities-unity-catalog

Zuboff, S. (2019). *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power* (1st edition). PublicAffairs.

A REFLEXIVITY STATEMENT

In the tradition of ethnography and qualitative methods more generally (Charmaz, 2014), we tried to remain aware of our own cultural and epistemic perspectives in the course of this study. Reflexivity is not the same as accounting and correcting for bias—rather, we acknowledge that research is not a "view from nowhere" (Haraway, 1988). The first author (RS) is a Ph.D. student with training in computer science and economics with experience researching the social implications of algorithms, including differentially private mechanisms. The second author (AA) is a tenured professor with years of research experience studying the economics of privacy and privacy-enhancing technology. Both authors are white men based in the United States working at an elite Western university. Both authors have worked for tech companies in the past, and RS receives stipend support from Meta. Neither author has experience building and deploying PPA—rather, PPA systems have been the subject of our largely empirical research. We do not have strong political or intellectual convictions about the value or future of PPA in particular—this study grew out of our curiosity to understand how these technologies fit into privacy practice. However, both authors have an interest and stake in the protection of online privacy.

In our analysis, we considered how our backgrounds might lead us towards certain framings of our results—particularly those that frame PPA in terms of empirically quantifiable trade-offs—or close us off to certain possibilities—for example, that the adoption of PPA is not inevitable. We also considered how it may also lead our interviews towards certain topics (e.g. economic trade-offs or privacy scholarship) more than others. As described in §3, we communicated frequently and adopted a "devil's advocate" approach to coding to challenge our suppositions about what the data communicate or suggest. We made a special effort to recruit and consider perspectives from practitioners with marginalized identities. Recognizing that we cannot encompass all perspectives—we also leverage our backgrounds to "study up" (Nader, 1972), highlighting in the discussion where our analysis—and the subject of our research—reflect the views and interests of culturally hegemonic groups and institutions.

B INTERVIEW GUIDE

The final version of our semi-structured interview guide can be viewed here.