ALGORITHMIC DECOUPLING AND THE ADOPTION OF 'PRIVACY-PRESERVING' ANALYTICS

Ryan Steed* and Alessandro Acquisti†

Carnegie Mellon University

February 17, 2025

Abstract

Algorithmic techniques for privacy-preserving analytics (PPA) offer organizations a way to maintain and expand access to valuable data while preserving individuals' privacy. Adoption of PPA techniques is growing in industry and government. However, their impacts on privacy protection are not yet clear. Small differences in design can have significant downstream consequences for data privacy, but little research has yet examined the motivations and decision-making processes behind PPA adoption, and how those decisions ultimately influence whether adoption lives up to consumers' and regulators' expectations. We investigate the organizational processes driving adoption and deployment of PPA systems in a qualitative study based on interviews with executives, lawyers, engineers, and scientists implementing PPA across large technology firms, startups, non-profits, and government agencies. We develop grounded theory to describe the processes by which organizations interpret social expectations into algorithmic designs and justify their design choices. We find several mechanisms by which organizations decouple representations about privacy from the specifics of their implementation with algorithms. We introduce a new dimension of organizational decoupling, *algorithmic decoupling*, to describe the ways algorithmic systems contribute to gaps between organizational policy, practice, and outcomes. We explore the consequences of algorithmic decoupling for the future of regulation and research related to privacy and other issues.

*ryansteed@cmu.edu
†acquisti@andrew.cmu.edu

1 Introduction

Facing consumers' calls for privacy and policymakers' threats of regulatory action, technology leaders are endorsing "privacy-preserving" techniques for data analytics 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), and other algorithmic approaches to reconcile analytics and privacy.

On its face, consumers and regulators have reason to hope adoption will improve digital privacy: these mathematical, statistical, and algorithmic techniques—which we will broadly 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 their actual impacts on consumer privacy and social welfare are often untested. Privacy advocates (Cyphers, 2019) and advertisers (Lomas, 2021) alike, for instance, have raised doubts about Google's privacy-preserving plans for the future of targeted advertising (Goel, 2022). PPA adoption could present new benefits to researchers and consumers—or, as some fear, it could simply preserve extractive, surveillance-based economies (Cohen, 2018; Zuboff, 2019; McGuigan et al., 2023).

Despite a flourishing technical literature in computer science and statistics, little IS, economics, or social science research has examined organizational processes driving the adoption and deployment of PPA systems. What drives organizations to adopt privacy-preserving technologies? How do these drivers inform design choices made during deployment? And how does PPA adoption change the relationship between policy and practice in the organization?

In absence of empirical studies of PPA adoption in organizations, prior research does, however, offer possible theoretical answers to these questions. Some streams of organizational research suggest that organizations innovate socially beneficial technologies in response to changing consumer expectations and to regulatory pressure, such as for environmental protection (Ashford et al., 1985). But institutional research also suggests that policies adopted to satisfy consumers and regulators may at times be merely symbolic—"decoupled" from substantive changes to daily practice (Edelman, 2016; Bromley & Powell, 2012). On this trajectory, PPA adoption could amount to little more than "privacy theater": gestures accorded deference incommensurate with their actual benefits to online privacy (Soghoian, 2011).

We investigate the motivations and decision-making processes behind PPA adoption and design across 21 large technology firms and startups, non-profits, and government agencies. We interviewed

28 executives, managers, and key contributors in technical, legal, and policy roles responsible for deciding whether to adopt and how to deploy PPA systems. The firms and agencies they work for include many of the organizations leading PPA adoption in the United States. Analyzing transcripts and documents, we used grounded theory methodology (Charmaz, 2014) to develop a model of the organizational processes that led to and shaped adoption.

Our primary finding is that, during implementation, PPA design choices were constrained and sometimes dominated by operational concerns, disconnecting algorithmic system design from both internal policies and external representations—a special form of organizational decoupling we term *algorithmic decoupling*. Both public and private sector organizations adopted PPA systems primarily to preserve existing modes of operation against new regulations or consumer expectations, though some organizations leveraged PPAs to use data in new ways. The processes involved with deciding when to use PPA techniques and interpreting privacy expectations into algorithmic designs often prioritized these managerial interests. Through infrastructure sharing and active standard-setting, these practices set the bar for future industry implementations and even for regulatory guidance. However, we also find that morally motivated privacy "champions" (Tahaei et al., 2021) constituted a countervailing force against algorithmic decoupling within the organization. They had significant leverage over the interpretation of privacy laws and internal policies on the one hand and the technical properties of PPA systems on the other, mediating between managerial, legal, and technical concerns. A subset of these practitioners—particularly in the private sector—used their influence to push for PPA adoption and defend privacy standards, often because of their own ethical and professional commitments.

Connecting our findings with research from sociology, law, and technology studies, this study makes two key contributions to research on information systems (IS), organizational behavior, and privacy. First, we conceptualize algorithmic decoupling: ways in which the aspects of organizational practice embedded in algorithmic systems may be uniquely decoupled from formal policy and external expectations and may uniquely alter those expectations in turn. Prior work shows how practitioners' interpretations of the legal environment mediate the effect of regulation (Fuller et al., 2000). Algorithmic decoupling draws on another dimension of ambiguity: practitioners' *technological* interpretations of algorithmic systems employed to fulfill formal policies. Algorithmic decoupling helps to explain why even significant investments in PPA adoption often fall short of public privacy expectations (Martin et al., 2023). While our current study focuses on algorithms for private data processing, the proliferation of algorithmic systems in other parts of the market and society suggests that the propagation of similarly motivated innovations may not be an unalloyed benefit to society, if they are overly mediated by managerial concerns. Second, our findings suggest that although algorithmic ambiguity leaves room for decoupling, it also uniquely empowers expert

PPA practitioners. Where prior work focused on the role of executives' ethical commitments in preventing decoupling (Weaver et al., 1999) or the role of lawyers in framing the legal environment (Edelman, 2016), algorithmic decoupling is influenced most by technologists' commitments to privacy. Our findings offer a guide for scholars, managers, and policymakers to more critically evaluate and intervene in the use of algorithmic systems to fulfill social responsibilities.

2 Theoretical Background

2.1 Technology Adoption and Social Performance

Though organizational IS research on PPA adoption is scant, our work builds upon decades of research on the social benefits of digital innovation in IS (Yoo et al., 2010), particularly related to sustainability (Malhotra et al., 2013; Hanelt et al., 2017). At the firm-level, much IS research is devoted to identifying the antecedents (resources, competitive environment, management style, etc.) of information technology adoption and its effects on financial and operational performance (Fichman, 2004). Less research explores how organizations adopt IS technologies to improve *social* performance, efforts to fulfill social responsibilities alongside economic gains (Davis, 1973; Carroll, 1979; Orlitzky & Benjamin, 2001)—responsibilities which, some argue, now include data privacy (Pollach, 2011). In the last decade, technology firms have developed and adopted a number of algorithmic innovations to address privacy, fairness, sustainability, and other ethical concerns with data processing and artificial intelligence (AI) practices (Bamberger & Mulligan, 2015; Hirsch et al., 2020; Metcalf et al., 2019; Morozov, 2013). Recent IS research, for example, examines the drivers and performance benefits of "green" IS practices (including smart grids and building automation) adopted to improve environmental sustainability (Seidel et al., 2013; Hanelt et al., 2017; Malhotra et al., 2013; Leidner et al., 2022; Hu et al., 2016; Loeser et al., 2017; Ketter et al., 2023).

Earlier economic theories explain these practices simply as strategic, cost-saving adaptations to changes in regulation (Oliver, 1991). Stricter environmental standards, for example, forced firms to innovate new technologies such as the catalytic converter to avoid financial penalties (Ashford et al., 1985). A wide-ranging literature from institutional theory (Scott, 2007), on the other hand, explains organizational behaviors as a product of "rational myths"—widely-accepted ideas about how organizations should act, conditioned on historical, cultural, and social context (Meyer & Rowan, 1977; DiMaggio & Powell, 1983). Adoption is mediated not only by strict economic rationality but also by the expectations of activists, competitors, investors, employees, consumers, and other stakeholders (Campbell, 2007; Boldosova, 2019; Aguilera et al., 2007; Jones, 1995). Under this theory, adopting socially beneficial technologies helps organizations maintain their social license to operate (Gunningham et al., 2004).

2.2 Algorithmic Decoupling as a Dimension of Organizational Decoupling

A key observation of contemporary institutional theory is that organizations facing these external pressures may partially or completely *decouple* the performance of daily practices ("performative" aspect) from their presentation in formal structures and policies ('ostensive" aspect) (Oliver, 1991; Bromley & Powell, 2012; Boxenbaum & Jonsson, 2017; Feldman & Pentland, 2003). The concept of decoupling arises from early observations that organizations can maintain contradictory institutional logics by insulating inconsistent, yet responsive, practices from one another—a phenomenon referred to in institutional theory as "loose coupling" (Weick, 1976; Meyer & Rowan, 1977; Orton & Weick, 1990; Hallett & Hawbaker, 2021). This perspective has proved useful for analyzing IS adoption (Strong & Volkoff, 2010; Chen et al., 2011). Berente and Yoo (2012), for example, use loose coupling to explain improvisational user responses to enterprise IS adoption as a resolution to the friction between abstract software and local contexts.

For organizations facing strong but ambiguous or contradictory external and internal expectations (Powell & DiMaggio, 2023; Scott, 2007), decoupling—sometimes referred to as "organized hypocrisy" (Brunsson, 2003; Lim & Tsutsui, 2011)—serves as a buffer between internal practices and external pressures and as a key component of organizational legitimacy (Meyer & Rowan, 1977; Weber, 1978; Suchman, 1995). Several studies examine decoupling in the context of social performance efforts (Lim & Tsutsui, 2011; Schoeneborn et al., 2020; Dobbin & Kalev, 2022; Marquis & Qian, 2014; Li & Wu, 2020). For example, some firms obtained green technology certificates from the Korean government without actually implementing those technologies in daily operations (Park & Cha, 2019). Decoupling the formal adoption and espousal of these efforts from their practice allows organizations to reduce costs while avoiding legal sanctions and reputational harms (Bromley & Powell, 2012). However, nearly all studies of organizational decoupling focus on primarily non-technological practices (see, e.g., Bromley & Powell, 2012, Table 2). And while IS research explores the possibility of loose coupling as a response to institutional contradictions in enterprise system implementations (Berente & Yoo, 2012; Keller et al., 2019; Baptista et al., 2021; Chen et al., 2011), the interaction between organizational *decoupling* and IS has not been fully explored.

IS research on technology-mediated organizational change explores how organizational routines are *materially embedded* in enterprise IS (Volkoff et al., 2007). When information systems are adopted, routines, roles, and other organizational structures are constrained and modified by material aspects of the system as built (Silva & Hirschheim, 2007; Berente & Yoo, 2012). IS artifacts, then, constitute a form of embedded, often invisible, organizational regulation (de Vaujany et al., 2018; Hennigsson & Eaton, 2024)—a topic of nascent IS research agendas (Butler et al., 2023; de Vaujany et al., 2018). de Vaujany et al. (2018) call for further research on the "materialization" of rules

in IT artifacts, in addition to temporal decoupling between design time and use time; this study explores the processes involved with rules materialization and the consequences for regulation & compliance. Technological embeddedness introduces the possibility of decoupling the presentation of organizational practice (its ostensive aspect) not only from its performance (performative aspect), but also from the aspect of practice embedded in information systems (material aspect).

Algorithmic decoupling helps describe this additional dimension: the gap between policy and the technical and material properties of the algorithmic system deployed to fulfill it. Research on technology and social performance often frames the adoption of technologies like the catalytic converter as uniformly implemented and categorically beneficial (Ashford et al., 1985); algorithmic decoupling accounts for the reality that technology adoption is contextual and adapted, its design mediated by the organization. In the context of PPA, privacy advocates' criticisms of several prominent, public proposals provide an early indication that organizations' claims have not been fulfilled by their technological designs (Cyphers, 2019; McGuigan et al., 2023; Martin et al., 2023); algorithmic decoupling helps to explain these shortcomings. In our study, we explore how the use of algorithmic systems complicates existing theory about the mediators of and remedies to organizational decoupling.

2.3 From Algorithmic Decoupling to Perverse Innovation

Decoupling makes clear that organizational responses to external pressures are not determined. Organizations and their constituents mediate the impact of the institutional environment on the adoption of new practices (Edelman, 2016; Oliver, 1991). But institutional research on *heterogeneous diffusion* also explores the influence of adoption on the institutional environment in turn (Powell & DiMaggio, 2023).

In particular, organizations model their early "educated guesses" at compliance to their peers and competitors. As in the case of equal opportunity, employment law, and insider trading, these initial guesses are often legitimated by courts and policymakers and become standard practice (Edelman, 2016; Bozanic et al., 2012). 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 (Sutton & Dobbin, 1996; Edelman et al., 1999; Dobbin & Kalev, 2022). Edelman (2016) calls this phenomenon *legal endogeneity*: after organizations decide what forms of compliance are reasonable, those practices become institutionalized as rational responses to regulation. Private organizations may also engage directly in lobbying, corporate-sponsored research, and other forms of regulatory capture to promote their versions of compliance (Hillman et al., 2004; Kamieniecki, 2006). Technology firms in particular, such as Airbnb and Uber, are

exemplars of "regulatory entrepreneurship": the pursuit of business models that are predicated on changing the law (Pollman & Barry, 2016).

When organizational practices are mediated by algorithms and IS, regulatory entrepreneurship may be accomplished with technological innovation. Burk (2016) uses the term "perverse innovation" to describe technological innovation directed at exploiting loopholes in formal rules. Seed producers in the E.U., for example, avoided restrictions on genetically-modified crops by replacing recombinant DNA technologies with mutagenic chemicals, an alternative approach with possibly greater health and safety risks; and the PT Cruiser was designed with the footprint of a "small truck" to allow Chrysler to avoid stricter EPA fuel efficiency requirements for "passenger cars" (Burk, 2016).

Algorithmic decoupling is perverse innovation when and if the implemented algorithmic system is not only disconnected from but contrary to expected social benefits (e.g., "privacy-preserving" technologies that increase data collection without providing substantive privacy benefits). Like other compliance practices, technological designs may set legal precedent. Algorithmic decoupling helps to explain how perverse practices may become institutional standards not only through sociolegal mechanisms but also through sociotechnical mechanisms, primarily cloud platform-dependence (Narayan, 2022; Cutolo & Kenney, 2021)and open source innovation (West & Gallagher, 2006).

2.4 Privacy-Preserving Analytics

This study explores organizational decoupling in the context of a burgeoning area of IS technology for social performance: privacy-preserving analytics (PPA). Privacy technology is not new—we define PPA techniques as a particular subset of privacy enhancing technologies (PETs),[1] a variety of tools used by consumers, regulators, and organizations to negotiate information privacy issues for over three decades (Goldberg, 2007). Privacy is a multifaceted, context-dependent social concept associated with a wide range of attitudes and behaviors (Dinev et al., 2015; Acquisti et al., 2015; Bélanger & James, 2020; Belanger & Crossler, 2011). Likewise, while all the systems referred to as "privacy-preserving" in our study were used in practice to govern user data processing, the methods—and the precise definition of privacy preserved—varied (McGuigan et al., 2023). This study focuses on techniques and standards used to preserve privacy in both the inputs to and the outputs of data analysis, such as secure multiparty computation (Goldreich, 2009)(which describes cryptographic protocols for distributed computing designed not to reveal private inputs), differential privacy (a formal guarantee that outputs of analysis are not sensitive to the inclusion of any one individual's information, usually accomplished by noise injection (Dwork et al., 2006)), and federated learning (which describes techniques for training machine learning models without transferring raw data

[1] 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.

off client devices (McMahan et al., 2017)). Table 1 lists all the PPA practices adopted by our participants.

Organizations have used differential privacy, for example, to send COVID-19 exposure notifications and auto-complete text or emojis (Apple, Google), collect telemetry (Microsoft Windows), share data with clients and researchers (Meta, LinkedIn, Microsoft, Google, U.S. Census Bureau), and more (Desfontaines, 2021). In fact, probably spurred by regulatory initiatives in the U.S. and around the world, the number of private and public sector organizations adopting and deploying PPAs has significantly increased in recent years. New startups such as Tumult Labs are offering PPA consulting services and building open-source software. And cryptographic and federated methods—such as Google's Privacy Sandbox (Goel, 2022)—may soon replace key aspects of online advertising.

With respect to privacy practices in general, there is evidence of decoupling in existing research: Waldman (2018) 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, and several studies critique the claims to privacy made by public PPA proposals (McGuigan et al., 2023; Tang et al., 2017; Berke & Calacci, 2022; Martin et al., 2023). But the organizational processes behind PPA adoption specifically—and the technological aspects of decoupling in general—are less understood. And the resulting impacts—on digital privacy as well as data science, social science research, policymaking, and other data-dependent processes (Abowd & Schmutte, 2019; Hotz et al., 2022)—are largely untested outside primarily theoretical research in computer science and statistics (Acquisti & Steed, 2023).

A few interview studies explore practitioners' challenges with differential privacy adoption specifically, but these studies mostly focus on usability (Dwork et al., 2019; Munilla Garrido et al., 2023; Sarathy et al., 2023; Ngong et al., 2024; Rosenblatt et al., 2024). Some critical studies evaluate PPA proposals on technical or philosophical grounds (Tang et al., 2017; Berke & Calacci, 2022; McGuigan et al., 2023; Martin et al., 2023; Smart et al., 2022) and explore challenges with communication and participation during adoption (boyd & Sarathy, 2022; Abdu et al., 2024). Other studies investigate data ethics practices (Hirsch et al., 2020) and privacy practices for artificial intelligence (AI) products (Lee et al., 2024). But little research examines organizational aspects of these technologies. By investigating this question, our study contributes to the ongoing project of documenting and describing the social impact of PPA technologies.

3 Methods

This research is based on a seven-month qualitative study of PPA adoption through semi-structured interviews with practitioners—including engineers, lawyers, managers, researchers, policy experts,

and executives—at technology firms, privacy-focused startups, non-profits, and government agencies. These organizations include many of the most prominent deployments of PPA to date in the United States. Research on organizational adoption of PPA is scant, but qualitative methods have a long, impactful history of helping researchers theorize about emerging phenomena in IS and management (Monteiro et al., 2022; Wiesche et al., 2017; Edmondson & Mcmanus, 2007).

3.1 Data Collection

Our data are comprised primarily of IRB-approved interviews with 28 individuals responsible for helping their organizations decide whether and how to implement PPA systems. Table 1) describes their roles and technologies used. We contacted practitioners working on PPA products or services at organizations that had considered deploying PPA, though not all have actually deployed a PPA system. (All had made it at least as far as prototyping.) We sourced interview candidates either from professional networks (18 contacted, $N = 9$ interviewed) or known to both authors through public PPA work (7 contacted, $N = 5$ interviewed). We also asked those candidates to recommend one or two others at their organization (23 contacted, $N = 14$ interviewed). SM Appendix C contains additional details about our recruitment strategy.

We designed our sample to explore theoretical variation between different adoption settings, aiming for analytical generalization from case studies to theory (Eisenhardt, 1989; Lee & Baskerville, 2003). Our sample includes 21 organizations: eight technology firms ($N = 10$), six in the Fortune 500 ($N = 8$); five privacy-focused startups ($N = 6$), organizations with privacy-branded products or offering PPA as a service; four non-profits ($N = 5$); and representatives from three U.S. government agencies ($N = 3$), two responsible for federal data collection and public statistics and one 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.

After we analyzed this first round of data, we conducted a second round of interviews between July and August 2023 with three practitioners in legal and policy roles to validate our understanding of how private firms interact with regulators about PPA adoption. We stopped data collection when our categories reached theoretical saturation, such that further interviews would spark no new insights (Charmaz, 2014). Each participant completed a short demographic survey and participated in a 50–100 minute interview (57 minutes, median) through video conferencing, under the condition that their identities were kept confidential. Interviews centered on open-ended questions about organizational processes involved with 1) the decision to adopt PPA, including motivations and trade-offs; 2) design and deployment, especially communication, common challenges faced, and future trends. Interviews were semi-structured, co-constructed by the interviewer and the participant to allow flexibility to explore new phenomena (Charmaz, 2014). We piloted our initial interview

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: statistical disclosure limitation (Matthews & Harel, 2011)(e.g., suppression, data swapping).

Table 1: Practitioners interviewed. Participants were given differential privacy (DP) and federated learning (FL) as examples but were allowed to name any practices they used for PPA.

guide using two think-aloud interviews (Willis & Artino, 2013)with colleagues and three practice interviews with volunteer junior practitioners. While we asked about PPA adoption in every interview, we adjusted the protocol to explore different areas of theoretical interest over the course of the study. (SM Appendix B provides our interview guide.) We supplemented interview transcripts with internal documentation provided by participants, press releases, white papers, blogs, news articles, and other archival documents.

3.2 Analysis

As is common in inductive research and grounded theory (Gioia, Corley, & Hamilton, 2013; Charmaz, 2014), qualitative analysis alternated continuously between 1) first-order, primarily inductive creation of analytic codes, 2) second-order aggregation and abductive theoretical analysis, and 3) written and visual presentation of our emergent theoretical model, grounded in first-order quotations. In first-order analysis, the first author annotated transcripts with short, precise descriptions—over 2,500 unique codes—staying grounded in the participants' language and focusing on actions and processes to avoid preconceived framing (Charmaz, 2014). Early on, the second author re-coded a sample of four interviews and provided critical feedback to calibrate our coding.

In second-order analysis, we critically sorted and synthesized initial codes to draw out hypotheses and narratives (axial & theoretical coding). We began to define tentative concepts by comparing first-order codes and excerpts and by comparing the accounts of different participants. At this stage, we adopted a theoretically agnostic stance, permitting extant concepts (such as "scaling up") only when they fit our data and first-order analysis (Charmaz, 2014). We also compared and generalized

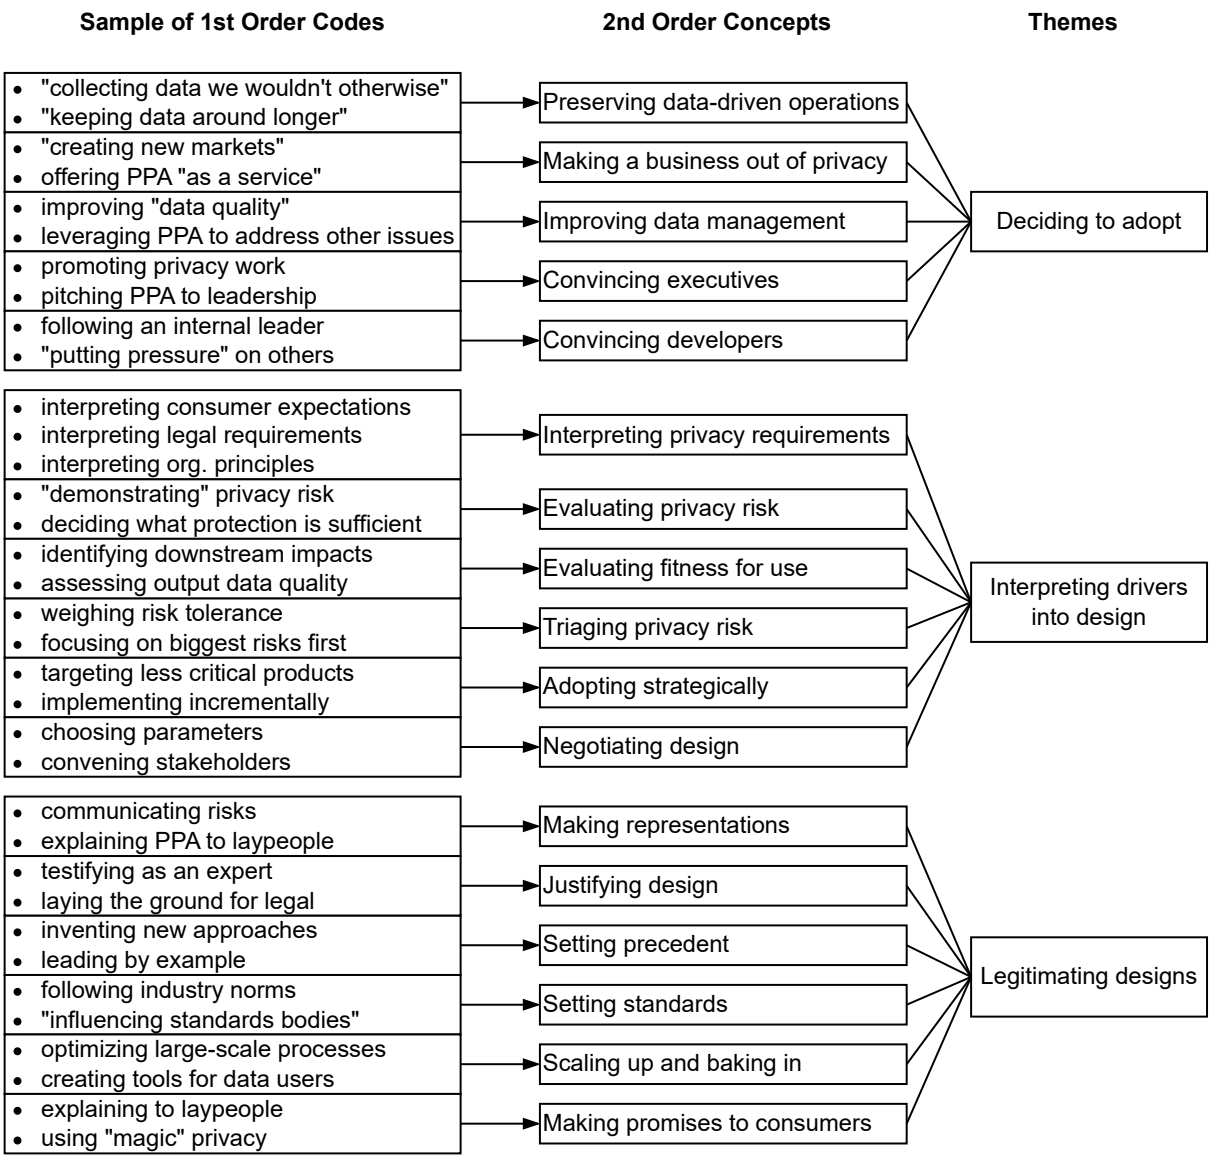

Figure 1: Themes and second-order concepts used in our process model, with an illustrative sample of first-order codes. Figure inspired by Gioia, Corley, and Hamilton (2013).

across types of organizations and industries, similar to a case study design (Eisenhardt, 1989). We gradually arranged concepts in multi-layer hierarchies and eventually three themes describing the overarching processes involved with PPA adoption (Figure 1) and mapped the relationships between them with written notes and iterative process diagramming (Figures D.1–D.3). As in the first-order analysis, the authors discussed and exchanged notes and diagrams describing the concepts and their relationships until a consistent process model emerged (Gioia, Corley, & Hamilton, 2013). Analysis occurred alongside data collection, and we continuously adjusted our interview guide to follow up on topics of theoretical interest (Charmaz, 2014).

4 Case Study: Adoption of Privacy-Preserving Analytics

In the following analysis, we trace the couplings between formal policy, its implementation, and its outcomes, and highlight where these couplings are likely to break. First, we describe the ways organizations constructed PPA adoption as an appropriate response to external privacy expectations (§4.1). Second, we describe the processes by which those narratives were interpreted into specific technological design choices—choices potentially decoupled from policies or outcomes (§4.2). Third, we describe the ways that organizations justified PPA adoption to satisfy stakeholder expectations, setting a precedent for future adoption (§4.3).

4.1 Deciding to Adopt

Organizational investment in PPA technologies specifically has increased sharply in the last decade following "loud and furious" (P4) public and regulatory pressure embodied by sweeping data privacy regulations, including the General Data Protection Regulation (GDPR) and the California Consumer Privacy Act (CCPA), coupled with publicized data privacy scandals—for example, the news that Cambridge Analytica had deceptively amassed the personal data of millions of Facebook users (Confessore, 2018).

Across the array of organizations we studied, practitioners consistently pointed to these external privacy expectations—embodied particularly by regulators, the media, and consumer advocates—as the root of their motivation to develop and deploy PPA systems. All but three practitioners named the threat of regulation—including fines, lawsuits, and, for government officials, criminal penalties—as contributing to their organizations' decision to adopt, and sixteen (especially those at for-profits and privacy startups) mentioned specific legal requirements or agreements with regulators. Seven—especially those working in policy or legal roles outside of government—also mentioned the possibility of negative media coverage and public backlash that could sour relations with external stakeholders and make it more difficult to recruit talented employees: "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). Faced with these expectations and the increasing inadequacy of existing data governance methods against modern reconstruction, linkage, and other attacks, many organizations—such as the U.S. Census Bureau (Abowd & Hawes, 2023)—looked for new solutions.

4.1.1 Motivational narratives.

These expectations provided broad motivation for organizations to change privacy practices but did not prescribe PPA adoption specifically—none of our participants described a regulation or public campaign that favored new PPA technologies over long-standing, non-algorithmic alternatives such as data minimization. In the organizations we studied, PPA adoption came to be viewed as an

appropriate solution through three different narratives. Practitioners and decision-makers used one or more of these narratives to convince decision-makers and justify their investment.

Preserving data-driven operations. Most commonly, organizations developed PPA to preserve a data-dependent model of operation—for example, behavioral advertising based on precise tracking and measurement. As one employee of a technology firm 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).

Instead, organizations developed their own, overwhelmingly algorithmic, solutions. In online advertising, for example, Google developed plans 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 invested in research on "novel privacy-preserving technologies" for ad measurement (Meta Research, 2019). Both organizations framed PPA adoption as a compromise between business interests and privacy concerns. In a blog post titled "A Path Forward for Privacy and Online Advertising", Meta's Chief Privacy Officer wrote: "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" (Egan, 2020).

PPA adoption provided a path for organizations in our study to "collect data we wouldn't otherwise collect" (P3), "monetize data more properly" (P17), and "keep data around longer" (P26)—for example, by anonymizing data to bypass legal limitations on storage. Practitioners in both private and public organizations said they tended to "collect the data first and then figure out what it's useful for later" (P25). Some practitioners 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).

Making a business out of privacy. While all but one organization we studied adopted to preserve operations, thirteen—disproportionately startups—also welcomed PPA as an opportunity to access new, privacy-conscious market segments, develop more competitive marketing, or develop new services. A director recalled, "Before... we would appeal to the principles of the company. Now... you can start saying things like, 'It probably will appeal to this market' " (P10). Web browsers like DuckDuckGo and Brave have grown market shares around privacy-preserving branding. Startups like Tumult Labs or Leap Year Technologies, recently acquired by cloud data giant Snowflake, offer PPA services to businesses, non-profits, and government agencies. And even within less privacy-branded technology companies in our sample, we observed some teams relying on different narratives than others, depending on the extent to which privacy differentiated the product they worked on. Non-profit and government organizations in particular used PPA systems to share new

sources of data with researchers—for example, the U.S. Census Bureau first used differential privacy to release new data on commuting patterns (Machanavajjhala et al., 2008).

Improving data management. To combat the view that additional privacy infrastructure was "costing resources with no clear benefits" (P17), some practitioners also argued that PPA would have side benefits beyond compliance. Several practitioners argued internally that the improvements to data management required to adopt PPA would also reduce "bad science" and ultimately make for better products. PPA adoption sometimes 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). However, no organization we studied adopted PPA solely to improve data management.

4.1.2 Convincing executives.

The primary audience for practitioners' adoption narratives was the key decision-makers within the organization—including the chief privacy officer, CEO, board members, and other "legal and risk" executives (P13). To convince executives, practitioners translated external privacy expectations into concrete business costs—one privacy engineer would "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). Executive buy-in helped push forward adoption on a case-by-case basis and drove the formation of internal policies—one large technology firm in our sample, for example, integrated PPA adoption into its existing privacy review process for new features.

4.1.3 Convincing developers.

Participants at government agencies and more hierarchical organizations relied mostly on these top-down policies to convince developers and other employees to contribute: "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). Participants at other organizations, though, discussed the importance of a less formal "privacy culture," especially for organizations that relied on product teams to "self-forward" (P8) relevant cases for privacy or volunteer for PPA adoption.

While internal policies were often based on cost calculus, the manner of adoption was often tied up with employees' moral judgments. Six of our participants—all at for-profit technology firms and start-ups—said they considered adopting PPA because it was the "right thing" to do for users. One privacy engineer at a technology firm 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)

This internal advocacy often depended on the leadership of influential privacy practitioners in centralized teams. Two participants at large technology firms observed peers "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 when an influential leader left, internal adoption of PPA dwindled. At least two organizations in our sample were deconstructing these central teams by the end of our study, distributing privacy professionals to product and infrastructure teams across the organization.

4.1.4 Adopting strategically.

Restricted to limited time and resources, most organizations we studied did not apply PPA uniformly across products and features—in fact, many adopted PPA for only one or two products or features. Instead, the rollout strategy depended on triage. For example, larger organizations had systems for prioritizing data deemed more "sensitive" or risky. Though most executives and managers agreed that it was better to start 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 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). Another perceived a fundamental 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).

4.2 Interpreting Policies into Designs

How did these external and internal drivers of adoption inform choices made during deployment? Adopting PPA—and negotiating its design—is not yet as simple as choosing a vendor or product off the shelf. Organizations made many specific design decisions to make their new systems "privacy-preserving" and align them with internal policies and external expectations. PPA practitioners were responsible for interpreting privacy requirements and guarantees, evaluating trade-offs in proposed designs, and negotiating with product teams, lawyers, and executives to triage privacy requirements, define scope, and settle on appropriate designs.

4.2.1 Interpreting privacy requirements.

In private firms and public organizations alike, PPA practitioners described "interpreting" or "translating" legal requirements and internal policies into technical specifications for algorithmic systems. From a former director at a technology firm:

> The way [the organization] did GDPR was: we had this privacy legal department, 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)

In rare cases practitioners could reference regulatory guidance—from the European Courts, for example (Data Protection Working Party, 2014)—or "hints" from regulators about which practices would be considered unacceptable (P23). Most relied on internal policies written by executives and legal teams: "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). These internal policies were designed to satisfy both legal requirements and the social expectations of "key opinion leaders" in policy and regulation, as one policy researcher describes: "What are their expectations, and what do we need to do to meet them?" (P12) Less commonly, participants referenced users' expectations—but nearly all of our participants did not interact with data subjects.

The process of interpreting these "external mandates" (P11), as one participant called them, was not straightforward. Around half of our participants mentioned conflicting, deficient, or "unreasonable" expectations in external regulations 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. One executive at a technology firm said, "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).

4.2.2 Triaging privacy risks.

When organizations did adopt PPA, simply evaluating the reduction of privacy risk provided by a particular algorithm was not straightforward. No single design met all requirements, especially for the large organizations we studied. As one executive put it, "you want to ask yourself what risk are you willing to take, how much uncertainty are you willing to live with" (P14). Another manager at a large technology firm 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 [on] 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)

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).

The ϵ parameter in differential privacy, for example, theoretically bounds the amount an individual's inclusion increases their risk of unwanted disclosure, but difficulty interpreting its value has contributed to heated epistemic disagreements (boyd & Sarathy, 2022; Nanayakkara & Hullman, 2022). Practitioners disagreed on what designs and parameter settings are meaningfully "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 guarantees were still meaningful after common relaxations. Others doubted further whether differential privacy is even an appropriate technical conception of privacy (see, e.g., Hotz et al., 2022; Seeman & Susser, 2023). Some fell back on "experimental" or "practical" guarantees to convince stakeholders (Dwork et al., 2019).

Without a clear conception of privacy risk, tuning the strength of privacy protections was more art than science. Legal requirements that data be "anonymous" or "confidential", for example, have been satisficed by successively stronger technical standards in just the past decade, including the use of k-anonymity to satisfy the Health Insurance Portability and Accountability Act (HIPAA) Privacy Rule (Malin et al., 2011) and the use of differential privacy in the 2020 Decennial Census (Abowd & Hawes, 2023). As one privacy engineer at a large technology firm 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, stakeholders with different incentives interpreted privacy risk differently. For example, Social Science One, an association of researchers partnering with Facebook to release a large dataset, argued that less stringent privacy standards satisfied the General Data Protection Regulation and Facebook's FTC consent decree, but Facebook disagreed, releasing the dataset with differential privacy (King & Persily, 2020).

4.2.3 Negotiating design.

These interpretations guided design negotiations between privacy practitioners, the teams responsible for implementing the PPA system, and other internal (and rarely external) stakeholders. A privacy engineer said,

> [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)

For six participants, all at for-profit firms, the relationship between business operations and PPA adoption was adversarial. A former employee of a large technology firm said:

> 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).

As a result, design negotiations usually centered managerial and technical concerns—most participants framed the design process as maximizing "fitness for use"—including data quality, cost, efficiency, usability, and interoperability—subject to a minimum "privacy bar". Minimum standards helped reinforce boundaries during design negotiations—as one privacy engineer at a large technology firm said, "The [product] team would come in and say okay well how about this ϵ. So a lot of the time, we can really easily just point to, 'No we have a standard, it's at this ϵ, use it'" (P19). Practitioners in at least four organizations leveraged formal processes like privacy review to resist opposing pressure from other executives: "everything goes through privacy, so we can just tell the leadership 'Hey, we're not letting this launch unless you do this'" (P11).

Because the interpretations of privacy risk and privacy guarantees could vary, standards were often flexible: "a lot of this we've just had to come up with out of thin air" (P23). A privacy engineer put it 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 k-anonymity, for example, another 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 [usually we] start the discussion from $k = 1000$ or $k = 100$ at least and then go down from there" (P8). A lawyer for a technology firm 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." Larger organizations had teams of privacy experts who helped product teams set parameters case-by-case. But some are now curtailing the authority of these teams. Recently, Meta reportedly permitted product teams, rather than privacy teams, to make the final decision about what privacy risks are acceptable (Huang, 2025).

4.3 Legitimating Designs

Once designs were set and PPA systems deployed, organizations still needed to convince regulators, consumer advocates, and other stakeholders that their design choices were acceptable to gain the social, economic, and legal benefits of adoption. Organizations relied again on privacy experts to decide how to represent the privacy properties of their PPA designs in legal arguments to regulators

or marketing promises to consumers. And organizations scaled up their PPA practices, setting standards internally and promoting those standards to the rest of their industry, creating precedent for rational response to privacy pressures.

4.3.1 Making representations.

In our study, all the organizations that deployed a PPA system translated its properties into some external representation, usually a legal defense, a promise to consumers, and in some cases a policy campaign.

Justifying designs. Several practitioners—particularly those who adopted to preserve existing operations—designed PPA systems with legal justification in mind. As a policy researcher asked, "Can you back that [public statement] up with your systems in an investigation?" (P12) This judgment fell to practitioners' legal and technological expertise.

> At some point the lawyer and I just 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 laws 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 [differential privacy] 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... probably you're going to be fine.' (P18)

Practitioners were prepared to rely heavily on their own expertise. A researcher at another large technology firm recounted defending 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)

Making promises to consumers. Practitioners also contributed to public representations their organizations made about privacy, including public privacy policies and marketing. Most privacy engineers did not have direct contact with data subjects—they explained their work only to legal and policy teams. Those that did described the difficulty of accurately representing algorithmic protections to laypeople: "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).

Some organizations dealt with this opacity by choosing "easy-to-describe" systems (P15). Others, particularly non-profits and government agencies, sought to increase transparency with open source software and detailed public descriptions. Several reported making changes in response to feedback from the public, especially at non-profit organizations. The U.S. Census Bureau, for example, submitted their differentially private 2020 disclosure avoidance system to multiple rounds of public comment, commissioned reviews from organizations like JASON and MITRE, and adjusted parameters and post-processing in response (boyd & Sarathy, 2022). Even when code was public, though, our participants—and independent researchers (Dwork et al., 2019; Gong, 2022)—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.

Others were less concerned about the need to comprehensively educate consumers: "[PPA is] 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). At least three participants, in both government and industry, worried that marketing new PPA protections would "muddy the waters" (P21) by revealing flaws in previous practices. A privacy director at a technology firm pointed out a practical upside to operating PPA "in the background": firms would no longer have to ask for consent and risk "creep[ing] people out" (P7).

Several practitioners had concerns about misrepresentation. Six of our participants, mostly in private industry, feared adoption that amounted to "magic privacy" or privacy "pixie dust"—black box techniques that, when invoked in marketing or legal copy, symbolically assure consumers and regulators of strong privacy and foreclose further inspection. An executive at a technology firm said bluntly: "most of [PPA] 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 disclosing limitations: "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 at for-profit firms—brought up the possibility that their PPA efforts were just "good theater" (P28), "privacy whitewashing" (P12), or "adoption for show" (P5). One participant mentioned Google's Privacy Sandbox as an example of this kind of proposal.

4.3.2 Setting precedent.

Substantive or otherwise, organizations' PPA practices became a model for others, particularly in the absence of clear industry standards. From a lawyer at a technology firm: "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). One privacy engineer at a large technology firm explained:

> 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)

As a result, internal standard-setting efforts were developed in anticipation of future regulation: "Let's try to fix this two years before they make it mandatory" (P3).

Leading by example. Some organizations—particularly privacy-focused startups—aimed to actively *guide* future regulation. Executives at two privacy boutiques agreed: "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). Practitioners in legal and policy teams at large private firms explicitly advocated to policymakers for regulation that would provide "safe harbor" (P2) from regulatory requirements for organizations which implemented their preferred PPA techniques. Practitioners at two large technology firms and a privacy-focused startup described steps their organizations took to influence regulation and develop relationships with policymakers—as that director put it, "glad handing" (P2)—outside the normal course of fact-sharing. Other organizations in our study worked to "influence standards bodies" (P14) such as the International Organization for Standardization (ISO), the National Institute of Standards and Technology (NIST) or the World Wide Web Consortium (W3C).

These efforts influenced other organizations. At least seven practitioners in non-profits and smaller organizations modeled their PPA systems after others' prominent deployments, particularly the use of differential privacy in the Decennial Census.

Scaling up & baking in. Second, organizations spread internal standards through developer tools and documentation. One manager said, "We bake [PPA] into infrastructure. We make it so it can't be screwed up" (P23). All but five practitioners described building or updating software infrastructure more generally in the course of scaling up their PPA systems. And most practitioners, especially those using PPA to create new products, built or updated tools for developers—mostly software libraries and internal platforms—to increase the capacity of a limited number of PPA experts. They aimed to construct a "well-lit path" for developers. As one corporate executive 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 could lay "well-lit paths" for whole industries by sharing their infrastructure. Amazon and other cloud service companies, for example, have included PPA tooling and features in their analytics products (AWS, 2023). Twelve participants, disproportionately from smaller non-profit and government organizations, mentioned integrating external PPA software libraries or other infrastructure—built by large technology firms such as IBM or by open source communities such as OpenDP—to deploy their own PPA systems, though these

tools still required expertise to use correctly.

A few of industry practitioners had concerns that these kinds of infrastructure could help firms box out competitors and control PPA development. Several mentioned Google's Privacy Sandbox and a new Apple feature that allows users to opt out of some in-app tracking used by ad brokers including Meta and Google (Morrison, 2022). Both projects drew anti-competition criticism from advertisers in France and the United Kingdom (Lomas, 2021). As one researcher at a large technology firm 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

We observed multiple points of decoupling between external privacy expectations on the one hand and the properties of deployed "privacy-preserving" systems on the other (Table 2). Some of these points of decoupling are consistent with existing organizational theory. When internal constituents failed to reinforce external pressure, leaders and developers often abstained from adopting PPA for core products—a commonly studied type of policy-practice decoupling (Bromley & Powell, 2012). And the trends in adoption we observed parallel the initial stages of legal endogeneity (Edelman, 2016, p. 27-41): organizations encountered ambiguous or absent privacy regulation, constructed PPA as a relevant solution, designed & implemented PPA systems to prioritize managerial concerns, and diffused those systems across industry. It is not yet clear the extent to which PPA systems will be endorsed by courts and administrative agencies, but practitioners expected it: "A lot of [adoption] is not so much 'this is what the law says' as 'differential privacy seems to make the regulators happy'" (P19). Technology companies such as Google are already advocating for regulatory exceptions for the PPA techniques they use (Google, 2022).

However, the material, technological aspect of these information systems complicates existing theory in two key ways. First, we observed less obtrusive points of decoupling arising from the algorithmic aspects of PPA systems, a new dimension of decoupling that we term *algorithmic decoupling*; second, our analysis suggests that instances of algorithmic decoupling impact law and society through new mechanisms. We highlight two important mediators of algorithmic decoupling based on cross-sectional analysis.

5.1 Algorithmic Decoupling

Technological infrastructure—the code bases and cloud services that industry "runs on"—is inextricably linked with practice (Lampland & Star, 2009). Just as routines executed by people can become decoupled from policies and expectations, so can routines executed by algorithmic systems (Lessig,

Stage	Description of decoupling	*Representative quotation*
Deciding to adopt	PPA adoption justifies additional data collection, undercutting privacy benefits	*[There is] market demand for something that will get you legal cover for hoovering up as much information as possible… the business model isn't going to change at all. (P15)*
	Reliance on voluntary adoption results in low take-up	*We depend on people to in the company to understand the privacy context of the company and self forward to triage. (P8)*
	Leadership exempts core products from 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)*
Interpreting policies into designs	Product teams negotiate for weak privacy parameters	*The [product] team is trying to do the least amount of privacy that we will approve... So it's really about how much political capital [they] have within the organization... (P19)*
	Designers misinterpret privacy guarantees	*To make [open source DP libraries] work you still have to know what you're doing... Worst case scenario you'll think you're using DP when really you're actually not. (P26)*
	PPA designs are ambiguously related to privacy requirements	*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)*
Legitimating designs	Marketing makes overly simplified or deceptive claims	*Most of [PPA] marketing is bullshit... they're writing checks that their tech cannot cash. (P7)*

Table 2: Points of decoupling in PPA adoption.

1999). Organizational research describes how employees form an interpretation of their organization's legal environment and mobilize that "legal reading" in their everyday work (Fuller et al., 2000). Our findings add a new dimension: PPA experts also employed their own *technological readings*, interpreting technical properties alongside the social and legal 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 design, legal and technological readings were jointly consequential. Particularly for recently developed techniques, modern privacy laws rarely admit to straightforward, specific translations to technical implementation (Nissim & Wood, 2018; Balebako et al., 2014).

Our study suggests that this additional layer of technological interpretation can also contribute to decoupling. We call this dimension *algorithmic decoupling*—a gap between formal policy (e.g., an organizations' public promises about privacy) and the material aspects of organizational practice

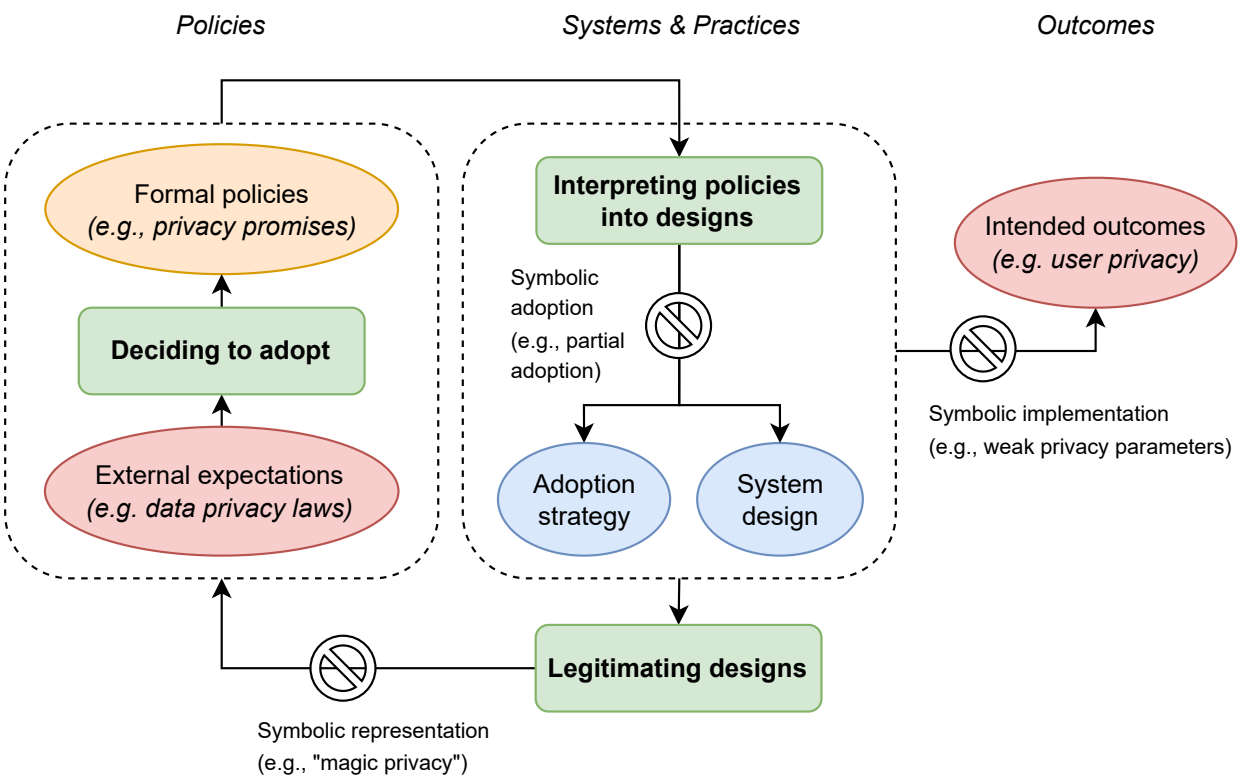

Figure 2: Algorithmic decoupling in our emergent process theory. Figures D.1–D.3 detail the bolded subprocesses.

embedded in algorithmic processes (e.g., the ways personal data are processed) (Volkoff et al., 2007). Our study focuses on the professionals experts assigned to navigate this gap, translating between the ostensive and the material aspects of their organizations' privacy practices.

Compared to other organizational practices, the operation of algorithmic systems may be more easily obfuscated, subject to less scrutiny, and therefore more easily and permanently decoupled from policies and outcomes. Information technologies generally gain scale through translation and loose coupling between many layers of digital devices, networks, algorithms, and services (Faik et al., 2020; Yoo et al., 2010). The relative invisibility and complexity of these layers can make algorithmic systems inscrutable to non-experts (Burrell, 2016; Metcalf et al., 2023; Selbst et al., 2023; Jin & Salehi, 2024). For example, a major 2002 privacy transparency requirement for federal agencies was undermined by "the inaccessible idiom of technology", which impeded public participation and oversight (Bamberger & Mulligan, 2008). Moreover, algorithmic systems increase the *scale* of organizational practices involving data processing and data-driven decision-making; designers' decisions have an outsized influence on the outcomes of these processes, compared to other routines.

Algorithmic decoupling complicates prior research on organizational decoupling (Bromley &

Powell, 2012; Boxenbaum & Jonsson, 2017). Prior research separates decoupling into policy-practice decoupling (symbolic adoption) and means-ends decoupling (symbolic implementation) (Bromley & Powell, 2012). Algorithmic decoupling adds an additional dimension—practitioners identified instances of algorithmic decoupling both between policies and practices (policy-practice) and between practices and outcomes (means-ends) (Figure 2).

Algorithmic decoupling shares certain properties with means-ends decoupling in particular: because of its inscrutability, it may be more durable than traditional policy-practice decoupling, which some argue do not withstand public scrutiny for long (Bromley & Powell, 2012). And while extensive research suggests that policy-practice decoupling may be reversed when employees reinforce societal expectations based on professional standards, moral commitments, or personal identity (Edelman, 2016; Haack et al., 2012; Turco, 2012; Gioia, Patvardhan, et al., 2013), algorithmic decoupling (and means-ends decoupling) may occur nonetheless (Bromley & Powell, 2012)—when non-expert developers misinterpret privacy guarantees, for example.

Means-ends decoupling, however, is more likely in opaque institutional fields where the effect of practices on outcomes, such as sustainability, is hard to precisely measure (Wijen, 2014). This is only partially the case for modern algorithmic systems; the material aspects of "black box" algorithmic systems are often opaque, but it sometimes possible for technical experts with access to a system to precisely measure its impacts on well-defined outcomes of concern such as discrimination and disinformation (Kroll, 2018). Crucially, algorithmic decoupling involves not only disconnect between process and outcomes (as in means-ends decoupling) but also between policies and the properties of algorithmic systems used to process data. This ambiguity is especially salient when regulations focus more on procedures (e.g., standards for data processing) than outcomes, requiring practitioners to justify choices in terms system properties (e.g., the ability to "de-identify" personal data).

Studies disagree on whether decoupling is more likely in early in the adoption process, when implementation expectations are flexible for early adopters (Bromley & Powell, 2012), or later, when established symbolic cues may suffice (Tolbert & Zucker, 1983; Kennedy & Fiss, 2009). Like policy-practice decoupling, algorithmic decoupling seems more likely when technical interpretations of the underlying algorithms are uncertain or contested, such as early on in adoption. Policy-practice decoupling is not always intentional; it may arise as managers "muddle through" competing stakeholder expectations (Crilly et al., 2012). In our study, technologists muddled through technical ambiguities in addition to competing requirements. Unlike policy-practice decoupling, which is inversely related to pressure from market stakeholders, regulators, and other external constituents (Stevens et al., 2005; Okhmatovskiy & David, 2012; Bromley & Powell, 2012; Marquis & Qian, 2014) and resistance from internal constituents (Turco, 2012), algorithmic decoupling may persist

in spite of scrutiny if oversight lacks technical specificity or there is no clear consensus on the acceptable algorithmic translation of social values such as privacy.

5.2 Algorithm-Mediated Legal Endogeneity

Algorithmic decoupling also entails new mechanisms for initial "guesses" at compliance to endogenously alter industry practices and institutional norms (Edelman, 2016). Information systems represent material standards and guidelines for the ethical treatment of personal data (Verbeek, 2006; Lampland & Star, 2009)—fundamental matters of privacy in IS, for example, are delegated to technologists and standard setting bodies (Waldman, 2018; Doty & Mulligan, 2013).

The construction of shared IS infrastructure, then, has potential for lasting influence on the institutional environment (Faik et al., 2020). In our study, for example, less-resourced developers' tendency to rely on only a few entrenched tool libraries meant that early adopters had even more influence over practices in the rest of industry. This influence grows as more developers use PPA products built into dominant cloud computing platforms such as Amazon Web Services (Narayan, 2022; Cutolo & Kenney, 2021; AWS, 2023). And algorithmic systems receive additional legal deference as courts often treat their design as a technical inevitability rather than question the design choices that lead to their creation (Selbst et al., 2023; Metcalf et al., 2023; Jin & Salehi, 2024).

5.3 Important Mediators of Algorithmic Decoupling

5.3.1 Privacy Champions.

The technical experts able to penetrate the veil of technological idiom hold special leverage over algorithmic decoupling. They had varying amounts of influence over the decision to adopt—like the success of enterprise IS adoption in general (Liang et al., 2007), decoupling depends heavily on the support or opposition of influential executives and managers and their social ties to other decoupling or non-decoupling organizations (Westphal & Zajac, 2001; Fiss & Zajac, 2004, 2006; Weaver et al., 1999). But when executives decided to adopt, the process of interpreting policies into designs was heavily dependent on PPA experts, often only one or two individuals. Nine of the organizations we studied—disproportionately smaller organizations—mentioned lack of experience as a barrier to PPA adoption and six consulted with external experts; one non-profit employee said that their PPA deployment would not have "gotten off the ground" (P19) without an external consultant from a large tech firm. Moreover, organizations in our study relied heavily on performances of expertise to justify their PPA systems to the public. The HIPAA Privacy Rule, for example, permits de-identification methods certified by a statistical expert (U.S. Department of Health and Human Services, 2012; Malin et al., 2011).

Some practitioners leveraged their central role in adoption to promote strong privacy standards

and prevent decoupling. Recent work chronicles the role of privacy "champions": institutional entrepreneurs who advance privacy through informal education and daily work when official policies are missing or insufficient (Tahaei et al., 2021). Many of our participants viewed their work as aligned with external calls for privacy. The approaches these participants took to advocate for adoption mirrored stages of emergent moral leadership (Solinger et al., 2020), including a precipitating scandal, internal coalition building, negotiation with other stakeholders, and establishment of formal structures (e.g., privacy review). They promoted PPA adoption through technological defaults—the "well-lit paths" our participants mentioned—and formed special interest groups, often centralized teams, to set and maintain privacy standards. But they also struggled to prioritize and organize around ethics in corporate environments (Lee et al., 2024), as do software engineers more generally (Widder et al., 2023; Ali et al., 2023)

5.3.2 Motivational Narratives.

Cross-organization analysis of our sample suggests that algorithmic decoupling is also mediated by the dominant narrative motivating adoption.

Decoupling was more difficult for organizations making a business out of privacy. Our participants reported decoupling most commonly when adopting to preserve existing data-driven operations, disproportionately at large technology firms. For example, the strategy of adopting only for less "mission critical" use cases (§4.1) was reported disproportionately by practitioners in private industry, as were concerns about deceptive invocations of "magic privacy" algorithmic techniques (§4.3). At large firms adopting mostly to comply with existing regulations, employee-led adoption efforts were more likely to falter as key privacy leaders left the company and the company broke up central privacy groups. But privacy-motivated practitioners at privacy-branded startups were more positive about PPA standards at their companies. Indeed, prior research suggests that when employees' identities are tied up with their organizations' external representations, decoupling is less likely to succeed (Turco, 2012).

Large organizations had more influence on the institutional environment. Institutionally endogenous processes such as participation in standards bodies were disproportionately discussed by practitioners at large firms. They used legal justifications and "standardization work" to convince regulators that PPA adoption was sufficient to keep existing practices compliant. Industry leaders had an advantage in influencing the institutional environment—both technology adoption and legitimacy spread through networks of influence (DiMaggio & Powell, 1983; Rogers, 2003).

5.4 Boundary Conditions

In our study, we analyze the adoption of a unique class of algorithmic systems for privacy-preserving analytics. Besides providing thick descriptions of important phenomena like PPA adoption (Lee

& Baskerville, 2003), unconventional contexts can be useful for developing newly insightful theory (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. Our model applies particularly to technological practices that 1) are implemented to improve social performance, 2) involve routines enacted by algorithms as well as people, and 3) have complex properties or impacts that require expertise and access 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, especially as calls for AI regulation proliferate (Hirsch et al., 2020; Lee et al., 2024). Thus, many of the mechanisms we identify in this unconventional context—for example, algorithmic decoupling—could also help explain the interaction between regulation and technology adoption in fields such as AI ethics, 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. 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. One practitioner, for example, described a future of "automatically invokable" mathematical techniques that assure users their privacy is preserved. It may be that the PPA systems deployed today *are* a substantive step towards privacy, and many of our participants were optimistic about their organizations' efforts. Scholars hope, for example, that the use of differential privacy could discourage dubious statistical practices such as p-hacking (Oberski & Kreuter, 2020).

But left to self-regulate, prominent early adopters may influence adoption in ways contrary to the public interest. For example, nearly all organizations in our study framed privacy risks as invasions or exploitation by state actors or other "attackers," shifting focus away from internal threats to privacy (Seeman & Susser, 2023). Privacy scholars argue that by focusing on specific technical properties such as individual anonymity or local processing, commercial PPA proposals legitimize invasive practices and foreclose more expansive privacy norms (Barocas & Nissenbaum, 2014; McGuigan et al., 2023; Martin et al., 2023; Yew et al., 2024)—a case of perverse innovation (Burk, 2016). 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). And scholars may worry that even as PPA adoption addresses real privacy concerns, it also legitimizes asymmetric economic and social relations (McGuigan et al., 2023; Veale, 2023; Viljoen,

2021). As one participant said, PPA "can't make [data collection] ethical just because it makes it private" (P15).

One remedy to managerial mediation is closer scrutiny (Edelman, 2016). 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 legal accountability in algorithmic systems (Selbst, 2021; Metcalf et al., 2023), particularly in 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 Meta, Apple, and Google (Tang et al., 2017; Berke & Calacci, 2022; McGuigan et al., 2023; Martin et al., 2023). Some of our participants—echoed by differential privacy scholars (Cummings et al., 2024; Gong, 2022; Dwork et al., 2019)—called for additional efforts to make PPA techniques more transparent, such as a registry of PPA parameter choices (Dwork et al., 2019).

Scrutiny and transparency are not unalloyed goods—they must be accompanied by formal mechanisms for accountability (Han, 2015; Birchall, 2014). For example, Google's Privacy Sandbox included a publicly documented plan to automatically cluster users into interest groups based on their behavior. The plan drew opposition from competing browsers Firefox, Brave, and Vivaldi (Bohn, 2021) and from the Electronic Frontier Foundation, which called it "the opposite of privacy-preserving technology" (Cyphers, 2019). Google eventually proposed a new design which grouped users into fewer, purportedly less sensitive clusters (Goel, 2022). Here, scrutiny from competitors and advocates helped to make technical reforms to an unpalatable design. But the Topics API still gave advertisers the power to target users based on their behavior, leaving Google's core business model relatively untouched while hampering its competition (Lomas, 2021)—and Google has since reverted its plans to phase out third-party cookies entirely. Waldman (2018) suggests that strong regulatory interventions, such as consent orders and weighty fines, can shock organizations into more integrated privacy practices. One PPA practitioner in our study described how a "push from the outside" (P19)—for their organization, a court ruling—initiated centralized privacy review and empowered PPA practitioners to maintain higher standards.

Judges, policymakers, investigators, and managers tasked with holding PPA operators accountable should look beyond the invocation of a technique like differential privacy to examine its full technical manifestation—including parameters, definitions of sensitivity, and post-processing steps. Privacy and algorithm impact assessments optimistically could provide structure for internal advocacy and provide valuable information to outside observers (Selbst, 2021), though such assessments have had mixed effects in practice (Bamberger & Mulligan, 2008; Brandtner, 2021; Smart et al., 2022). But even as policymakers seek to advance the development of PPA technologies (e.g., National Science and Technology Council, 2023; Office of Science and Technology Policy, 2022), they

should also take care not to endorse PPA techniques outside of context and to scrutinize PPA systems implemented within "safe harbor" frameworks. Data privacy regulation has long struggled with decoupling—reviews of a "gold standard" 2000 U.S. Safe Harbor agreement for processing personal data from E.U. citizens, for example, revealed that most participating organizations failed to implement basic requirements and others made false claims about certification (Connolly, 2008). Moreover, policymakers should not assume that a solely technological solution is sufficient (Green & Viljoen, 2020)or that "magic privacy" systems even function as advertised (Raji et al., 2022). In addition to scrutiny, policymakers and foundations could offer funding and other support to help less-resourced organizations contribute to PPA practice.

7 Conclusion, Limitations, and Future Research

Our study analyzes 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 expectations and technology design—interactions that have great implications for digital privacy. We conceptualize new technological dimensions to theories of decoupling (algorithmic decoupling) and legal endogeneity.

Our work has limitations. First, 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. Second, our participants' responses may be prone to social desirability bias or crafted to serve professional motives, though we do not repeat their views uncritically. Third, this study focuses on a set of formal, technical, and algorithmic approaches to privacy, and our analysis is influenced by our own backgrounds in computer science, economics, and public policy (see SM Appendix A for further reflection). Research from disciplines outside of computer science and statistics could further elucidate the social trade-offs facing practitioners and the downstream consequences of adoption for privacy standards. Moreover, our sample—like the population of PPA practitioners—is predominantly American, white, and cisgender male. Our study thus examines dominant perspectives in PPA work, but those perspectives do not encompass all possible paths for PPA development—for example, organizations' PPA plans rarely accounted for documented inequities in access to privacy (Skinner-Thompson, 2020; Allen, 2022; Madden et al., 2017). Researchers and policymakers could imagine alternative trajectories for PPA development that elevate public interests—for example, as a means to facilitate algorithm auditing (Xu & Zhang, 2021).

As large technology firms promote PPA techniques as the "future of personalized advertising" (Egan, 2020) and the digital economy, we hope our study will inspire further rigorous empirical investigation and theorization about the ways regulators and consumer advocates may shape the

development of socially-motivated innovation in the public interest.

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, Carrie Leana, 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 (IRB) at Carnegie Mellon University in 2021, #00000327.

Funding. This material is based upon work supported by the National Science Foundation under Grant No. 2319919.

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.

Abdu, A. A., Chambers, L. M., Mulligan, D. K., & Jacobs, A. Z. (2024). Algorithmic Transparency and Participation through the Handoff Lens: Lessons Learned from the U.S. Census Bureau's Adoption of Differential Privacy. *The 2024 ACM Conference on Fairness, Accountability, and Transparency*, 1150–1162.

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., Brandimarte, L., & Loewenstein, G. (2015). Privacy and human behavior in the age of information. *Science*, *347*(6221), 509–514.

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

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.

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.

AWS. (2023). *Differential Privacy - AWS Clean Rooms*. Amazon Web Services, Inc. https://aws.amazon.com/clean-rooms/differential-privacy/

Balebako, R., Marsh, A., Lin, J., Hong, J., & Cranor, L. F. (2014). The Privacy and Security Behaviors of Smartphone App Developers. *NDSS Symposium*.

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.

Baptista, J., Wilson, A. D., & Galliers, R. D. (2021). Instantiation: Reconceptualising the role of technology as a carrier of organisational strategising. *Journal of Information Technology*, *36*(2), 109–127.

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.

Belanger, F., & Crossler, R. (2011). Privacy in the Digital Age: A Review of Information Privacy Research in Information Systems. *Management Information Systems Quarterly*, *35*(4), 1017–1041.

Bélanger, F., & James, T. L. (2020). A Theory of Multilevel Information Privacy Management for the Digital Era. *Information Systems Research*, *31*(2), 510–536.

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*. Microsoft AI Blog for Business & Tech. https://blogs.microsoft.com/ai-for-business/differential-privacy/

Bohn, D. (2021). Nobody is flying to join Google's FLoC [magazine]. *The Verge*.

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

Boxenbaum, E., & Jonsson, S. (2017). Isomorphism, Diffusion and Decoupling: Concept Evolution and Theoretical Challenges. In *The SAGE Handbook of Organizational Institutionalism* (pp. 77–97). SAGE Publications Ltd.

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.

Brandtner, C. (2021). Decoupling Under Scrutiny: Consistency of Managerial Talk and Action in the Age of Nonprofit Accountability. *Nonprofit and Voluntary Sector Quarterly*, *50*(5), 1053–1078.

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.

Brunsson, N. (2003). *The Organization of Hypocrisy: Talk, Decisions and Actions in Organizations* (2nd edition). Copenhagen Business School Pr.

Burk, D. (2016). Perverse Innovation. *William & Mary Law Review*, *58*(1), 1.

Burrell, J. (2016). How the machine 'thinks': Understanding opacity in machine learning algorithms. *Big Data & Society*, *3*(1), 2053951715622512.

Butler, T., Gozman, D., & Lyytinen, K. (2023). The regulation of and through information technology: Towards a conceptual ontology for IS research. *Journal of Information Technology*, *38*(2), 86–107.

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.

Carroll, A. B. (1979). A Three-Dimensional Conceptual Model of Corporate Performance. *Academy of Management Review*, *4*(4), 497–505.

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

Chen, Y.-D., Brown, S. A., Hu, P. J.-H., King, C.-C., & Chen, H. (2011). Managing Emerging Infectious Diseases with Information Systems: Reconceptualizing Outbreak Management Through the Lens of Loose Coupling. *Information Systems Research*, *22*(3), 447–468.

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

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

Connolly, C. (2008). *The US Safe Harbor - Fact or Fiction?* Galexia. Pyrmont, Australia.

Crilly, D., Zollo, M., & Hansen, M. T. (2012). Faking It or Muddling Through? Understanding Decoupling in Response to Stakeholder Pressures. *Academy of Management Journal*, *55*(6), 1429–1448.

Cummings, R., Desfontaines, D., Evans, D., Geambasu, R., Huang, Y., Jagielski, M., 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. (2024). Advancing Differential Privacy: Where We Are Now and Future Directions for Real-World Deployment. *Harvard Data Science Review*.

Cutolo, D., & Kenney, M. (2021). Platform-Dependent Entrepreneurs: Power Asymmetries, Risks, and Strategies in the Platform Economy. *Academy of Management Perspectives*, *35*(4), 584–605.

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

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

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

de Vaujany, F.-X., Fomin, V. V., Haefliger, S., & Lyytinen, K. (2018). Rules, Practices, and Information Technology: A Trifecta of Organizational Regulation. *Information Systems Research*, *29*(3), 755–773.

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

Dinev, T., McConnell, A. R., & Smith, H. J. (2015). Research Commentary—Informing Privacy Research Through Information Systems, Psychology, and Behavioral Economics: Thinking Outside the "APCO" Box. *Information Systems Research*, *26*(4), 639–655.

Dobbin, F., & Kalev, A. (2022). *Getting to Diversity: What Works and What Doesn't*. Harvard University Press.

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.

Doty, N., & Mulligan, D. K. (2013). Internet Multistakeholder Processes and Techno-Policy Standards. *Journal on Telecommunications and High Technology Law*, *11*, 135–182.

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. *Proceedings of the Third Conference on Theory of Cryptography*, 265–284.

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.

Edmondson, A. C., & Mcmanus, S. E. (2007). Methodological fit in management field research. *Academy of Management Review*, *32*(4), 1246–1264.

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.

Feldman, M. S., & Pentland, B. T. (2003). Reconceptualizing Organizational Routines as a Source of Flexibility and Change. *Administrative Science Quarterly*, *48*(1), 94–118.

Fichman, R. G. (2004). Going Beyond the Dominant Paradigm for Information Technology Innovation Research: Emerging Concepts and Methods. *Journal of the Association for Information Systems*, *5*(8).

Fiss, P. C., & Zajac, E. J. (2004). The Diffusion of Ideas over Contested Terrain: The (Non)adoption of a Shareholder Value Orientation among German Firms. *Administrative Science Quarterly*, *49*(4), 501–534.

Fiss, P. C., & Zajac, E. J. (2006). The Symbolic Management of Strategic Change: Sensegiving Via Framing and Decoupling. *Academy of Management Journal*, *49*(6), 1173–1193.

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.

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

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

Gioia, D. A., Patvardhan, S. D., Hamilton, A. L., & Corley, K. G. (2013). Organizational Identity Formation and Change. *Academy of Management Annals*, *7*(1), 123–193.

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* (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.

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.

Hallett, T., & Hawbaker, A. (2021). The case for an inhabited institutionalism in organizational research: Interaction, coupling, and change reconsidered. *Theory and Society*, *50*(1), 1–32.

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

Hanelt, A., Busse, S., & Kolbe, L. M. (2017). Driving business transformation toward sustainability: Exploring the impact of supporting IS on the performance contribution of eco-innovations. *Information Systems Journal*, *27*(4), 463–502.

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

Hennigsson, S., & Eaton, B. D. (2024). Governmental regulation and digital infrastructure innovation: The mediating role of modular architecture. *Journal of Information Technology*, *38*(2), 126–143.

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* (Research Paper No. 628).

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.

Hu, P. J.-H., Hu, H.-f., Wei, C.-P., & Hsu, P.-F. (2016). Examining Firms' Green Information Technology Practices: A Hierarchical View of Key Drivers and Their Effects. *Journal of Management Information Systems*, *33*(4), 1149–1179.

Huang, K. (2025). *Meta Curbs Privacy Teams' Sway Over Product Releases*. The Information. https://www.theinformation.com/briefings/meta-curbs-privacy-teams-sway-over-product-releases

Jin, A., & Salehi, N. (2024). (Beyond) Reasonable Doubt: Challenges that Public Defenders Face in Scrutinizing AI in Court. *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, 1–19.

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

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

Keller, R., Ollig, P., & Fridgen, G. (2019). Decoupling, Information Technology, and the Tradeoff between Organizational Reliability and Organizational Agility.

Kennedy, M. T., & Fiss, P. C. (2009). Institutionalization, framing, and diffusion: The logic of TQM adoption and implementation decisions among U.S. Hospitals. *Academy of Management Journal*, *52*(5), 897–918.

Ketter, W., Schroer, K., & Valogianni, K. (2023). Information Systems Research for Smart Sustainable Mobility: A Framework and Call for Action. *Information Systems Research*, *34*(3), 1045–1065.

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

Kroll, J. A. (2018). The fallacy of inscrutability. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, *376*(2133), 20180084.

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*.

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

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.

Li, J., & Wu, D. ((2020). Do Corporate Social Responsibility Engagements Lead to Real Environmental, Social, and Governance Impact? *Management Science*, *66*(6), 2564–2588.

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.

Lim, A., & Tsutsui, K. (2011). Globalization and commitment in corporate social responsibility: Cross-national analyses of institutional and political-economy effect. *American Sociological Review*, *77*(1), 69–98.

Loeser, F., Recker, J., Brocke, J. vom, Molla, A., & Zarnekow, R. (2017). How IT executives create organizational benefits by translating environmental strategies into Green IS initiatives. *Information Systems Journal*, *27*(4), 503–553.

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). Privacy, Poverty, and Big Data: A Matrix of Vulnerabilities for Poor Americans. *Washington University Law Review*, *95*(1), 53–126.

Malhotra, A., Melville, N. P., & Watson, R. T. (2013). Spurring Impactful Research on Information Systems and Environmental Sustainability. *MIS Quarterly*, *37*(4), 1265–1274.

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.

Marquis, C., & Qian, C. (2014). Corporate Social Responsibility Reporting in China: Symbol or Substance? *Organization Science*, *25*(1), 127–148.

Martin, K., Nissenbaum, H., & Shmatikov, V. (2023). *No Cookies For You!: Evaluating The Promises Of Big Tech's 'Privacy-Enhancing' Techniques.* https://papers.ssrn.com/abstract=4655228

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

McGuigan, L., Sivan-Sevilla, I., Parham, P., & Shvartzshnaider, Y. (2023). Private attributes: The meanings and mechanisms of "privacy-preserving" adtech. *New Media & Society*.

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.

Meta Research. (2019). *New Research Award in Privacy Preserving Tech*. https://research.fb.com/blog/2019/11/facebook-announces-new-research-awards-in-privacy-preserving-tech-at-ccs/

Metcalf, J., Moss, E., & boyd, d. (2019). Owning Ethics: Corporate Logics, Silicon Valley, and the Institutionalization of Ethics. *Social Research: An International Quarterly*, *86*(2), 449–476.

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.

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. *MIS Quarterly*, *46*(4), iii–xix.

Morozov, E. (2013). *To Save Everything, Click Here: The Folly of Technological Solutionism* (45555th edition). PublicAffairs.

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

Munilla Garrido, G., Liu, X., Matthes, F., & Song, D. (2023). Lessons Learned: Surveying the Practicality of Differential Privacy in the Industry. *Proceedings on Privacy Enhancing Technologies*, *2023*, 151–170.

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

Nanayakkara, P., & Hullman, J. (2022). What's Driving Conflicts Around Differential Privacy for the U.S. Census. *IEEE Security & Privacy*, 2–11.

Narayan, D. (2022). Platform capitalism and cloud infrastructure: Theorizing a hyper-scalable computing regime. *Environment and Planning A: Economy and Space*, *54*(5), 911–929.

National Science and Technology Council. (2023). *National Strategy to Advance Privacy-Preserving Data Sharing and Analytics*. Executive Office of the President.

Ngong, I. C., Stenger, B., Near, J. P., & Feng, Y. (2024). Evaluating the usability of differential privacy tools with data practitioners. *Twentieth Symposium on Usable Privacy and Security (SOUPS 2024)*, 21–40.

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

Oberski, D., & Kreuter, F. (2020). Differential Privacy and Social Science: An Urgent Puzzle. *Harvard Data Science Review*, *2*(1).

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

Okhmatovskiy, I., & David, R. J. (2012). Setting Your Own Standards: Internal Corporate Governance Codes as a Response to Institutional Pressure. *Organization Science*, *23*(1), 155–176.

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.

Orton, J. D., & Weick, K. E. (1990). Loosely Coupled Systems: A Reconceptualization. *Academy of Management Review*, *15*(2), 203–223.

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

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.

Powell, W. W., & DiMaggio, P. J. (2023). The Iron Cage Redux: Looking Back and Forward. *Organization Theory*, *4*(4), 26317877231221550.

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.

Rosenblatt, L., Howe, B., & Stoyanovich, J. (2024). *Are Data Experts Buying into Differentially Private Synthetic Data? Gathering Community Perspectives*. arXiv: 2412.13030 [cs].

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. *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, 1–19.

Schoeneborn, D., Morsing, M., & Crane, A. (2020). Formative Perspectives on the Relation Between CSR Communication and CSR Practices: Pathways for Walking, Talking, and T(w)alking. *Business & Society*, *59*(1), 5–33.

Scott, W. R. (2007). *Institutions and Organizations: Ideas and Interests* (3rd edition). SAGE Publications, Inc.

Seeman, J., & Susser, D. (2023). Between Privacy and Utility: On Differential Privacy in Theory and Practice. *ACM Journal on Responsible Computing*.

Seidel, S., Recker, J., & vom Brocke, J. (2013). Sensemaking and Sustainable Practicing: Functional Affordances of Information Systems in Green Transformations. *MIS Quarterly*, *37*(4), 1275–1299.

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

Selbst, A., 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. (2020). *Privacy at the Margins*. Cambridge University Press.

Smart, M. A., Sood, D., & Vaccaro, K. (2022). Understanding Risks of Privacy Theater with Differential Privacy. *Proc. ACM Hum.-Comput. Interact.*, *6*, 342:1–342:24.

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.

Stevens, J. M., Steensma, H. K., Harrison, D. A., & Cochran, P. L. (2005). Symbolic or substantive document? The influence of ethics codes on financial executives' decisions. *Strategic Management Journal*, *26*(2), 181–195.

Strong & Volkoff. (2010). Understanding Organization—Enterprise System Fit: A Path to Theorizing the Information Technology Artifact. *MIS Quarterly*, *34*(4), 731.

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].

Tolbert, P. S., & Zucker, L. G. (1983). Institutional Sources of Change in the Formal Structure of Organizations: The Diffusion of Civil Service Reform, 1880-1935. *Administrative Science Quarterly*, *28*(1), 22–39.

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

U.S. Department of Health and Human Services. (2012). *Guidance on De-identification of Protected Health Information*.

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!* To appear in: Hans-Wolfgang Micklitz and Giussepe Vettori (eds.), The Person and the Future of Private Law (Hart, forthcoming).

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

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

Volkoff, O., Strong, D. M., & Elmes, M. B. (2007). Technological Embeddedness and Organizational Change. *Organization Science*, *18*(5), 832–848.

Waldman, A. E. (2018). Designing Without Privacy. *Houston Law Review*, *55*(3).

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.

West, J., & Gallagher, S. (2006). Challenges of open innovation: The paradox of firm investment in open-source software. *R&D Management*, *36*(3), 319–331.

Westphal, J. D., & Zajac, E. J. (2001). Decoupling Policy from Practice: The Case of Stock Repurchase Programs. *Administrative Science Quarterly*, *46*(2), 202–228.

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. *MIS Quarterly*, *41*(3), 685–701.

Wijen, F. (2014). Means versus Ends in Opaque Institutional Fields: Trading off Compliance and Achievement in Sustainability Standard Adoption. *Academy of Management Review*, *39*(3), 302–323.

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.

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*.

Yew, R.-J., Qin, L., & Venkatasubramanian, S. (2024). You Still See Me: How Data Protection Supports the Architecture of AI Surveillance. *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, *7*, 1709–1722.

Yoo, Y., Henfridsson, O., & Lyytinen, K. (2010). Research Commentary: The New Organizing Logic of Digital Innovation: An Agenda for Information Systems Research. *Information Systems Research*, *21*(4), 724–735.

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, we tried to remain aware of our own cultural and epistemic perspectives in our work. Reflexivity is not the same as accounting and correcting for bias—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 based in the United States and work at elite Western universities. 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—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 toward certain framings of our results 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 toward certain topics (e.g. economic trade-offs or privacy scholarship) more than others. Though we made an effort to recruit and consider perspectives other than our own, we primarily leverage our backgrounds to "study up" (Nader, 1972) and critically analyze culturally hegemonic institutions close to ourselves.

B Interview Guide

The final version of our semi-structured interview guide can be viewed here. Participants were compensated with a $30 gift card or donation to a charity. Each interview was recorded (with participant consent and IRB approval) and transcribed verbatim by the first author.

C Recruitment

In the first period of data collection between July 2021 and January 2022, we calibrated each successive wave of recruitment ($N = 9$ contacted July–August, $N = 11$ September–October, $N = 3$ November) to examine adoption settings and other theoretical interests we had yet to explain with previous data (e.g., in the second wave, we included privacy-focused startups). (e.g., the process of interpreting adoption drivers into specific designs)

Because privacy is acutely important to marginalized groups (Skinner-Thompson, 2020), we deliberately aimed to include those perspectives in our sample and explicitly requested referrals to participants from underrepresented backgrounds. Still, Our sample was predominantly American (90%), white (70%), non-Hispanic (85%), heterosexual (65%) and cisgender male (55%), based on participants who chose to self-identify for each category (22 for race, ethnicity, and gender; 17 for sexuality). Our sample is similar in racial diversity to the U.S. high-technology workforce, but more diverse in gender, sexuality, and Hispanic origin (U.S. Equal Employment Opportunity Commission, 2016).

D Additional Figures

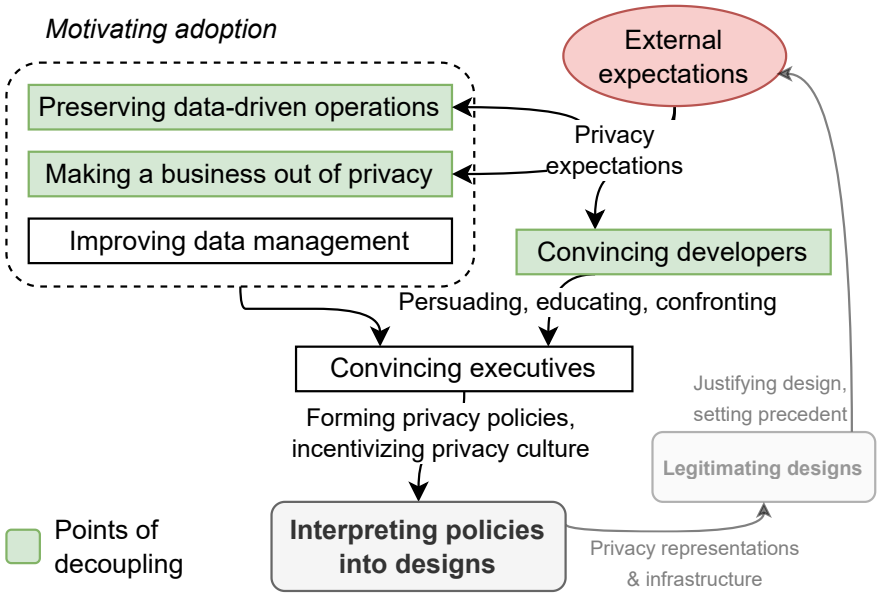

Figure D.1: Driving adoption. Second-order concepts are boxed. Key processes associated with managerial mediation (and possibly decoupling) or expert mediation are highlighted.

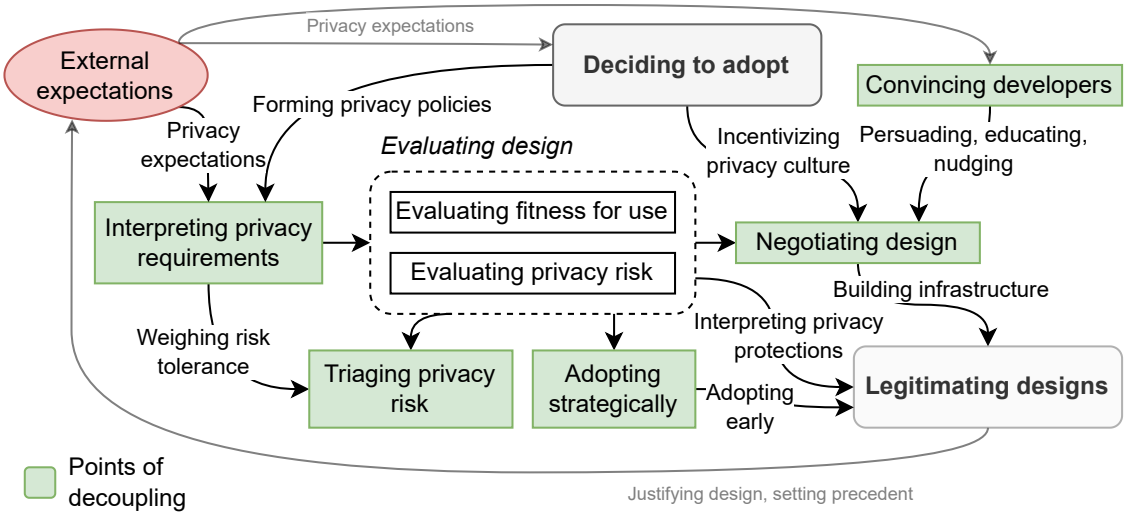

Figure D.2: Interpreting drivers into designs. Second-order concepts are boxed. Key processes associated with managerial mediation (and possibly decoupling) or expert mediation are highlighted.

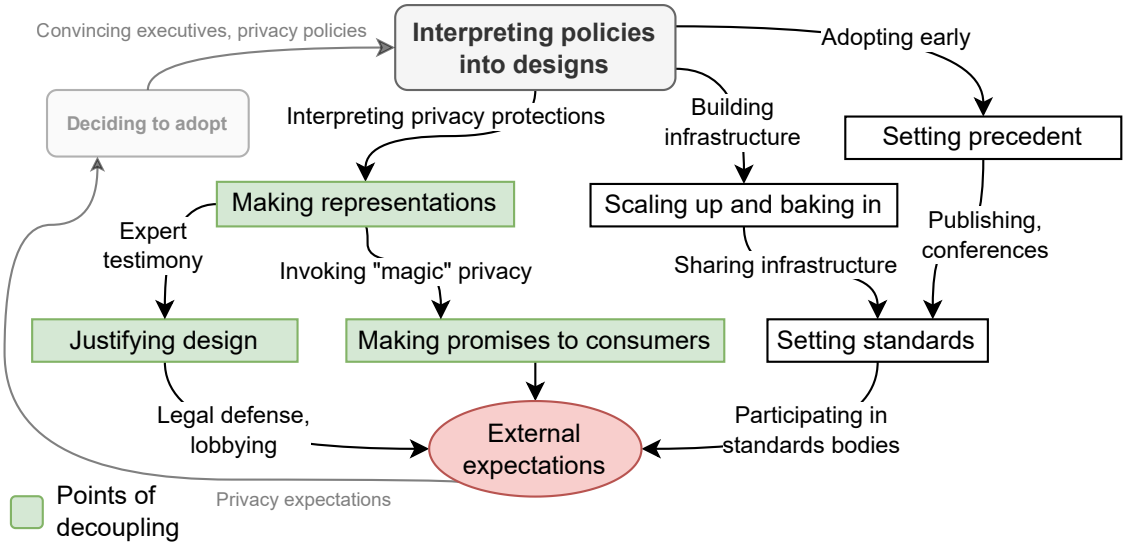

Figure D.3: Legitimating design. Second-order concepts are boxed. Key processes associated with managerial mediation (and possibly decoupling) or expert mediation are highlighted.