NO COOKIES FOR YOU!: EVALUATING THE PROMISES OF BIG TECH'S 'PRIVACY-ENHANCING' TECHNIQUES

DECEMBER 9, 2024

Kirsten Martin,[1] *Helen Nissenbaum,*[2] *Vitaly Shmatikov*[3]

ABSTRACT

We examine three common principles underlying a slew of 'privacy-enhancing' techniques recently deployed or scheduled for deployment by big tech companies: (1) limiting access to personal data by third-parties, (2) using inferences and minimizing use and retention of raw data, and (3) ensuring personal data never leaving users' device. Our Article challenges these principles, but not on the grounds that techniques offered to implement them fail to achieve their stated goals. Instead, we argue that the principles themselves fall short when the privacy-enhancing technique does not address privacy-violating behavior. Through philosophical analysis and technical scrutiny, we reveal the misalignment between the principles and a sound conception of privacy. We reinforce our findings empirically with a series of factorial vignette user studies, which demonstrate a surprising gap between the principles and users' actual privacy expectations. The most general conclusion that can be derived from our findings is that any effort to create successful privacy-enhancing systems must start with the explicit adoption of a meaningful conception of privacy.

[1] William P. and Hazel B. White Center Professor of Technology Ethics, Mendoza College of Business at the University of Notre Dame.

[2] Andrew H. and Ann R. Tisch Professor of Information Science, Cornell Tech at Cornell University.

[3] Professor of Computer Science, Cornell Tech at Cornell University. We would like to thank the participants at the Privacy Law Scholars Conference (2023), Digital Life Initiative Workshop at Cornell Tech, and the SPILab at the University of Michigan School of Information for their constructive feedback.

NO COOKIES FOR YOU!: EVALUATING THE PROMISES OF BIG TECH'S 'PRIVACY-ENHANCING' TECHNIQUES

INTRODUCTION

A. The Times They Are A-Changin'

Privacy abusers of yesteryear are some of the biggest privacy champions today. Following the exposure of their data collection and data sharing practices by privacy activists and popular media, and under pressure from policy imperatives (e.g., "privacy by design" in the General Data Protection Regulation's (GDPR) Article 25 in the E.U. and the growing patchwork of privacy laws and regulations in the U.S.), "platform" tech companies are increasingly reaching for answers in privacy-enhancing technologies. Industry- and public-facing media rattle off a plethora of technologies such as, differential privacy, encryption, and federated learning.[4] Although genuine interest in promoting the value of privacy for its own sake may drive adoption of privacy-enhancing techniques (PETs), understandably, adoption may also hold instrumental value to firms in softening public opinion and side-stepping the strictures of government regulation.

Digital advertising, and especially behavioral advertising, is still by far the most important revenue source for many platforms and content publishers. They are interested in privacy protection that (1) keeps platforms' own advertising-based business models intact, and (2) facilitates, or at least does not obstruct, key new technologies, such as machine learning trained on users' data. Promoters of PETs claim that companies can continue to enjoy the "utility" of users' data while satisfying the demands of public policy and public opinion.[5][6]

[4] Others in the market include encryption in transit (e.g. SSL), encryption at rest, de-identification, anti-tracking protections, secure multi-party computation, and homomorphic encryption.

[5] *See* Omar Ali Fdal, *What Are Privacy-Enhancing Technologies (Pets) and How You Can Choose the Right One(s)*, CPO MAGAZINE (July 1, 2022), https://www.cpomagazine.com/data-privacy/what-are-privacy-enhancing-technologies-pets-and-how-you-can-choose-the-right-ones/ [https://perma.cc/3CQY-LSLK].

[6] It is worth noting, this emphasis on preserving the viability of online behavioral advertising does not imply that all stakeholders in the economic landscape will flourish equally and it might even result in the virtual elimination of entire business models of non-platform players (e.g., certain types of ad-tech companies). For example, Web and mobile tracking techniques (such as third-party Web cookies and mobile advertising identifiers) endowed so-called "third parties," including thousands of ad-tech companies and data brokers, with the technical ability to track users and create their behavioral profiles. Insofar as privacy technologies deny access to these third parties, "first parties," (in particular, platform companies like Google) with direct relationships with users may be able to sustain or even increase business interests at the expense of other actors in the space. One stakeholder who would benefit from any decrease in behaviorally targeted advertising would be consumers, since consumer welfare is found to be lower with behaviorally targeted ads as compared to ads based solely on search terms. "...targeted ads in both studies are more likely to be associated with lower quality vendors, as well as higher prices (when comparing identical products), relative to competing alternatives found in search results." *See* EDUARDO A. S. MUTSTRI, IDRIS ADJERID & ALESSANDRO ACQUISTI, BEHAVIORAL ADVERTISING AND CONSUMER WELFARE: AN EMPIRICAL INVESTIGATION 4 (2023).

B. You Keep Using That Word Privacy

Adopting PETs—technical systems for enhancing and protecting privacy—seems like a step in the right direction for a variety of reasons. The adoption of PETs arguably implements the idea of privacy by design enshrined in the GDPR. By reflecting that sense of baked-in privacy, not easily revoked, this approach also implies that the creators of formerly privacy-violating technology are taking responsibility for binding their own hands.

At the same time, committing to PETs poses an essential question: what makes any given technical system (technology) *privacy*-enhancing? Selectively preventing certain flows of personal data while allowing others does not, by itself, qualify a system as privacy-enhancing; only certain patterns of constraints would so qualify. For example, a company may impose constraints on data in order to protect trade secrets, market power, or security, each of which may implicate privacy but not necessarily be fully aligned.[7]

Claiming that a given system is *privacy*-enhancing requires both that it demonstrably achieves certain flow patterns *and* that the patterns embody a sound conception of privacy. No matter how closely a system in question aligns with a conception of privacy held by the data processor, if that conception is flawed, the mission of building a privacy-enhancing system will fail. As an analogy, consider the sights of a firearm: no matter how perfectly aligned it is, its efficacy depends, simultaneously, on whether the target to which it is aligned is the right target.[8] Accordingly, we ask of some of the common techniques proposed as privacy enhancements whether what they enhance is, in fact, privacy. The conception of privacy that guides the authors of this Article is contextual integrity (CI), basically asserting that privacy is the appropriate flow of information.[9]

In Part I, we examine three popular privacy principles espoused by the tech industry as solutions that address privacy: (1) limiting access by third-parties to users' personal information, (2) minimizing the use and retention of raw data, and (3) ensuring personal data does not leave users' devices.. We explore how these principles work in practice and examine whether the approach to privacy implicit in these principles actually addresses users' privacy interests. We also explain how these principles manifest in judicial decisions and regulations. To argue that the principles fall short we apply philosophical analysis and technical scrutiny. These

[7] *See* Rory Van Loo, *Privacy Pretexts*, 108 CORNELL L. REV. 101, 111 (2022), cornelllawreview.org/wp-content/uploads/2023/03/2884.pdf [https://perma.cc/S6FJ-DCG7].

[8] *See* Brian C. Smith, *The Limits of Correctness*, 14, 15 ACM SIGCAS COMPUT. & SOC'Y 18 (1985) ("Just because a program is 'proven correct', in other words, you cannot be sure that it will do what you intend.").

[9] *See generally* HELEN NISSENBAUM, PRIVACY IN CONTEXT: TECHNOLOGY, POLICY, AND THE INTEGRITY OF SOCIAL LIFE (2020).

draw on a proposed set of three benchmarks for a sound conception of privacy that reveal the misalignment between the principles, on the one hand, and a sound conception of privacy, on the other. These benchmarks are precision (is the conception clear, rigorous, and internally consistent), ethics (does the conception explain privacy's ethical significance), and fidelity (is the conception roughly faithful to common usage).

In Part II, we reinforce our findings empirically with a series of factorial vignette user studies, which demonstrate a surprising gap between the three principles and users' privacy expectations. These empirical studies test whether the principles abide by the third benchmark, namely, rough fidelity to common usage. Structured by a conception of privacy as CI, our vignette study involves a contextual actor collecting user data and using it for a variety of purposes. Vignettes systematically vary (1) the actor, (2) whether inferences are drawn, and/or (3) whether a third-party is involved in receiving or using the data in question. Varying the values for actor, information type, and purpose of data collection within the CI template, we seek to uncover whether it made a difference that raw data was used rather than inferences, that first parties or third parties were involved in collection or use of data, and the nature of purposes for which data was used.

In Parts III-V, we discuss the results of our studies. In sum, we find:

1. A distinction between first and third parties is not fundamental to privacy preferences. Respondents judged the creation and use of inferences by a third-party data broker to be a privacy violation to the same degree as when performed by the first party (search, news, etc.) for use in behavioral advertising. In fact, respondents judged a first party *selling* their inferences to be *worse* than the data broker selling their inferences. For search in particular, respondents rated a third party broker the same or better than the search firm using inferences to subsequently place online ads.
2. Committing to not use "raw data" without at the same time expressing commitments about inferences does not assuage privacy concerns. Respondents judged that restricting access to raw data and, instead, creating, storing, and using inferences was not a privacy solution when used for advertising, selling access to the knowledge, or improving services. Respondents judged the use of inferences compared to raw data the same or worse at meeting their privacy expectations. This finding holds across types of data (location versus inferences based on location, search terms versus inferences based on search terms, etc.) and across purposes and uses. In fact, the use of raw data appears to better meet privacy expectations than use of inferences when improving services and

(in some circumstances) when placing ads on a website. We discuss this apparent paradox in Section II.A.

3. Purpose matters. Respondents' judgments were sensitive to the purposes for which information is used. Specifically, when the flows of information from one party to another serve contextual ends and values (e.g., to improve services), they were consistently found to be appropriate, across contextual actors (news sites, search engines, social networks). However, the same information used to promote non-contextual purposes, such as, online ad targeting, did not meet privacy expectations and was judged more negatively as a privacy violation.
4. Respondents rate selling or using inferences more negatively than selling or using raw data. When judging whether a third-party data broker should have access to raw data versus inferences based on that raw data, respondents judged selling inferences to be a privacy violation of greater magnitude than selling raw data. Respondents slightly preferred data brokers to buy raw data rather than the inferences based on that same data (however, both scenarios were negative). In addition, our studies showed that respondents evaluated the creation and use of inferences by a data broker *without the raw data leaving the original company* as a greater privacy violation than either (1) selling raw data directly to a data broker or (2) allowing trackers to collect the same data for a data broker.
5. The proposed "Sandbox" privacy solutions do not fully address users' privacy expectations or provide a solution over the alternative of third-party trackers and ad networks placing personalized ads. More specifically, respondents' ratings did not favor the collection and use of raw data over inferences for personalized ads; neither did it make a positive difference that only first parties were privy to the collection and use of personal data for purposes of targeted advertising.

The most general conclusion that can be derived from our findings is that any effort to create successful privacy-enhancing systems must start with the explicit adoption of a meaningful conception of privacy. Accordingly, any claims to having created a privacy-enhancing system need to make explicit the nature of the underlying conception of privacy it embodies.

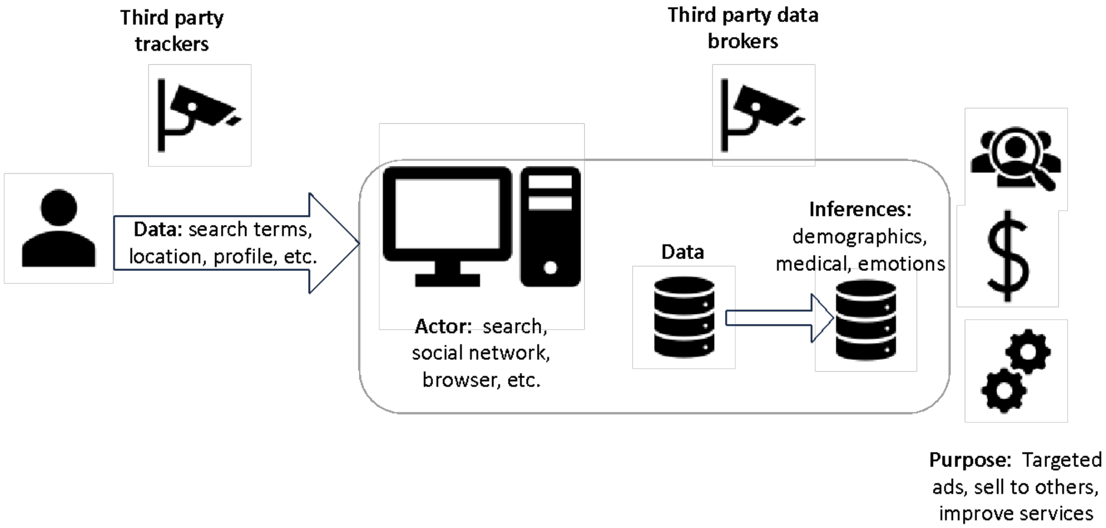
Third party trackers
Third party data brokers
Data: search terms, location, profile, etc.
Actor: search, social network, browser, etc.
Data
Inferences: demographics, medical, emotions
Purpose: Targeted ads, sell to others, improve services

I. THREE PRINCIPLES: CHALLENGED

The domain of privacy-enhancing techniques is indefinitely spurred by an ever-growing array of mechanisms, tools, and techniques that afford highly sophisticated data flow patterns.[10] In the ideal, these mechanisms, tools, and techniques would be artfully configured and orchestrated to produce privacy-enhancing versions of existing systems—for example, end-to-end encryption enhancing privacy in a messaging service. Accordingly, the quality or success of PETs rests not only on proving and verifying advanced mathematical and technical properties but, further, on demonstrating their success in achieving privacy aims in practical application.

In this Article, we focus on a subset of these applications that, recently, have been promoted as enhancements to the signature systems and services of prominent ad tech incumbents. We have organized them around three high-level principles that capture their key aims: (1) limit access by third- parties to users' personal information, (2) minimize the use and retention of raw data, and (3) ensure personal data does not leave users' devices. [11] The three principles are simplified abstractions, i.e., descriptive expressions aimed at non-experts in respective technologies, including regulators and present and future customers, to convey the essence of privacy enhancements of functional systems that employ these PETs.

Associated with these three principles are a range of privacy-enhancing techniques aimed at achieving and enforcing them. This Article questions whether modifying a system to meet the requirements of one or more of the three principles inevitably results in meaningful privacy enhancements. Even if these modifications significantly reroute data flows, the resulting flow patterns, even as they adhere to the principles, may not enhance privacy, under a sound conception of the term. To ascertain whether a conception of privacy is sound, we adopt three benchmarks:

(1) It is clear, rigorous, and internally consistent;
(2) It explains privacy's ethical significance; and
(3) It is roughly faithful to common usage.

[10] While we focus here on limiting access to third parties and minimizing the use and retention of raw data, others include secure multiparty computation, zero-knowledge proofs, and differential privacy. *See* Alessandro Acquisti & Ryan Steed, *Learning to Live with Privacy-Preserving Analytics*, 66 COMMC'NS ACM 24 (2023), https://doi.org/:10.1145/3597173; Ryan Steed & Alessandro Acquisti, *Presentation Video: Privacy-Preserving Analytics on the Ground* (2023), https://www.usenix.org/conference/pepr23/presentation/steed [https://perma.cc/47TM-JQ9L]; YUN SHEN & SIANI PEARSON, PRIVACY ENHANCING TECHNOLOGIES: A REVIEW (2011), https://citeseerx.ist.psu.edu/document?repid=rep1&type=pdf&doi=6bf9f0a288dd496de6bca96f360702b028fa0b58 [https://perma.cc/574J-BBWE]; Xuefei Yin, Yanming Zhu & Jiankun Hu, *A Comprehensive Survey of Privacy-preserving Federated Learning: A Taxonomy, Review, and Future Directions*, 54 ACM COMPUTING SURVS. 1, 1 (2021), https://doi.org/10.1145/3460427.

[11] Our argument does not assert that these three principles, and their associated technologies, fill the space of PETs being pursued.

In adopting these benchmarks, we are following a general approach in analytic philosophy to evaluating a proposed definition, or characterization of a complex concept, either moral (e.g., justice, responsibility, security) or nonmoral (e.g., knowledge, causation, action.) This involves checking the proposed conceptual analysis against commonsense beliefs, in an iterative back-and-forth dialog between intuitively held beliefs and a formally proposed definition.[12] Because, over time, natural-language terms accumulate a richness of meaning, which may result in ambiguity, contestation, or even internal inconsistency, a formal definition may need to excise some of these elements. At the same time, it must remain in balance with the term's natural meaning and significance. Here, the benchmarks provide heuristic criteria for evaluating a proposed conception of *privacy* – formally rigorous and true to common usage. The latter ensures that a definition retains meaning and relevance for natural language speakers, while the former ensures the degree of coherence and clarity demanded of a concept that is foundational for scientific, philosophical, regulatory, or legal analysis.

In the case of moral concepts, a definition needs to explain the moral weight it carries. A meaningful conception of justice, for example, would need to correctly identify unjust actions and policies and it should account for why they are wrong and should be revised. The third benchmark, accordingly, requires a meaningful concept of privacy to explain its normative force, why it is wrong to violate it, and why privacy deserves protection.

The benchmarks – precision, ethics, and fidelity – guide our evaluation of the three principles: (1) limiting third parties' access to personal information, (2) minimizing the use of raw data (while continuing to use inferences[13]), and (3)

[12] A version of this general method is famously known as "reflective equilibrium," a dialogical process for testing the coherence of a given belief (moral or other) against other beliefs, or for testing the coherence of a proposed account of a general concept (moral or other) in relation to a set of beliefs. We are aware that reflective equilibrium is itself a contested concept and have sought to be roughly accurate in adapting it to privacy. *See generally* DANIELS NORMAN, REFLECTIVE EQUILIBRIUM (2003).

[13] We are not the first to critically examine the creation and use of inferences as a privacy violation. For example, the creation of inferences from data violates a form of quantitative privacy. *See* David Gray & Danielle Citron, *The Right to Quantitative Privacy*, 98 MINN. L. REV. 62 (2013). As can inferences created about an individual based on similar people's inferences. *See generally* Solon Barocas & Karen Levy, *Privacy Dependencies*, 95 WASH. L. REV. 555 (2020); Brent Mittelstadt, *Protecting Health Privacy through Reasonable Inferences*, 22 Am. J. BIOETHICS 65 (2022); Alicia Solow-Niederman, *Information Privacy and the Inference Economy*, 117 NW. UNIV. L. REV. 357 (2022); Jacob L. Kröger, Leon Gellrich, Sebastian Pape, Saba R. Brause & Stefan Ullrich, *Personal Information Inference from Voice Recordings: User Awareness and Privacy Concerns*, 1 PROC. PRIV. ENHANCING TECH. 6 (2022); Sandra Wachter, *The Theory of Artificial Immutability: Protecting Algorithmic Groups under Anti-discrimination Law*, 97 TUL. L. REV. 149 (2022); Sandra Wachter & Brent Mittelstadt, *A Right to Reasonable Inferences: Re-thinking Data Protection Law in the Age of Big Data and AI*, 2019 COLUM. BUS. L. REV. 494 (2019).

maintaining users' data only on their devices. For each of these three principles, we check how well the underlying conception of privacy meets the benchmarks.

Throughout the paper, we draw on the conception of privacy as contextual integrity (CI), which defines privacy as the *appropriate flow* of personal information. Appropriateness, at first approximation, is conformance with entrenched contextual informational norms, where the contexts in question include familiar social domains, or institutions, such as healthcare, education, commerce, etc. Contextual informational norms, abbreviated to privacy norms, are explicit or implicit rules expressed in terms of five parameters: senders, recipients, subjects, attribute type, and transmission principle —the "CI-tuple." The CI-tuple may also serve a descriptive function, capturing information flows (e.g., resulting from technical systems or information practices) in terms of dimension that are essential for ascertaining their privacy properties. When evaluating either norms or practices in terms of their *ethical* standing, CI prescribes a three-layered approach, one that scrutinizes harms and benefits to stakeholders, a second scrutinizing how ethical and political principles, such as justice, equality, etc., are affected. The third, a distinctive contribution of the theory of CI, considers the extent to which respective norms or practices promote the functions, purposes, and values of a given context.[14]

A. Limiting Access to Personal Information by Third Parties

The actor who has access to information (or the individual) matters for privacy expectations. Throughout the centuries, delineating who can see a person,[15] who can take a picture,[16] who can hear a secret,[17] and so forth., pervade our stories about privacy. And theories of privacy have provided guidance as to how to think about who can share or collect information within a given context. In the theory of CI, this insight is embodied in the data recipient parameter. It is no surprise, therefore, to find this insight in the principle of limiting access by third parties in, for example, Apple's Intelligent Tracking Prevention, which allows users to refuse cross-site tracking.[18]

[14] CI defines privacy as the appropriate flow of information to and from particular actors, based on specific transmission principles, and towards ends, purposes, and goals defined by the context. NISSENBAUM, *supra* note 9 (see especially, Chapter 8).

[15] *See* Ellis Davidson & Hilda Roderick, *The Legend of Lady Godiva*, 80 FOLKLORE 107 (1969). Lady Godiva is the 11th century story of Tom the Tailor ("peeping tom") looking at Lady Godiva when he was not supposed to while she was riding a horse in a public square.

[16] *See generally* Samuel Warren & Louis Brandeis, *The Right to Privacy*, *in* KILLING THE MESSENGER: 100 YEARS OF MEDIA CRITICISM 1 (Tom Goldstein ed., 1989).

[17] *See generally* ALEXANDRE DUMAS, THE THREE MUSKETEERS (Mary C. Waldrep & Suzanne E. Johnson eds., 2007). Throughout the book *The Three Musketeers*, the characters say "it's not my story to tell" when not wanting to gossip about others.

[18] *Privacy*, APPLE (2023), https://www.apple.com/privacy/features/ [https://perma.cc/6CAJ-75SZ] (last visited Nov. 8, 2024) ("You may have noticed that when you look at something to buy

Drawing a general distinction between first and third- parties sets apart actors with whom an individual interacts directly and those who are indirectly implicated. On the Web, this distinction is associated with a distinction between domains to which individual users intentionally navigate, with which they directly engage, versus unsolicited, potentially unknown, domains that nevertheless gain access to information about them, both within a given Web session and across many sessions. An early set of Web technologies that enabled this access by third parties was aptly named the third-party cookie.

The design of Web browsers operationalized the first-party, third-party distinction by displaying the domain of the top-level, first-party webpage in the browser's URL bar while content of third-party domains is silently embedded in frames within the first party's webpage. This distinction is thought to reflect the intuitive difference between the content that the user solicits via an intentional action (e.g., typing a URL into the browser or clicking on a link) and the unsolicited content that is rendered in the user's browser without the user intentionally requesting it (e.g., ads and analytics scripts). To protect cookie security, Web browsers enforce the same-origin policy which prevents a Web actor (e.g., a website) from retrieving cookies other than those set by actors from the same domain.

Controversial from the time they were introduced as a Web standard, third-party cookies have long been a thorn in the side of privacy advocates.[19] Whereas the functionality of first-party cookies was seen, potentially, to enrich the relationship between individuals and first parties, allowing the latter to maintain ongoing "relationships" with individuals (in the context of a Web session), third-party cookies were seen by detractors as ways for unsolicited parties to enter the fray. The logic behind the different attitudes to first and third parties is that it is reasonable and appropriate for information to flow to parties with whom individuals intentionally and directly engage, while hackles are raised when data flows to unknown and unsolicited others, who are conceived as lurkers and unsought interlopers.

The lens of CI reveals another reason for placing third parties under greater scrutiny, namely, in order to ascertain whether they are appropriate recipients for the data they capture through third-party cookies and other trackers. In our own work and the work of others, there is consistent evidence that people consider flows

online, you suddenly start seeing it everywhere else you go on the web. This happens when a third party tracks cookies and other website data to show you ads across various websites. Intelligent Tracking Prevention uses the latest in machine learning and on-device intelligence to fight this cross-site tracking. It hides your IP address from trackers so what you look at on the web remains your business, not an advertiser's. And you don't have to change any settings for these protections because Intelligent Tracking Prevention is on by default.").

[19] *See* Helen Nissenbaum, *From Preemption to Circumvention: If Technology Regulates, Why Do We Need Regulation (and Vice Versa)?*, 26 BERKELEY TECH. L. J. 1367, 1382 (2011).

to third-party actors from the ad ecosystem, for example, to be illegitimate.[20] They also find unacceptable other flows from first parties to certain third parties, for example, data brokers.[21] Here and in Part II we use the term "non-contextual" actors as a shorthand way to refer to data recipients in data flow occurrences that are inappropriate because of the recipient, holding fixed the values for the other four parameters. The term "contextual" actor is used for the reverse of this and could be the first party or even a third-party, for example, a delivery service, if appropriate.

Image 1: First versus Third Parties and Privacy Theory - Search.

	First Party	Third Party
Contextual Actor	Search engine	Payment processor; analytics firm seeking to improve services
Non-Contextual Actor	Ad network owned by same company	Ad trackers owned by different company

A turning point was reached when media outlets took an interest in the behavioral advertising landscape. Most striking was the Wall Street Journal article *What They Know About You*,[22] exposing to the public that there were innumerable actors, mostly unseen, with full or partial access to people's Web activities. Third-party cookies emerged as the focus of public attention and regulation's public enemy number one. Since then, the momentum has grown to put a stop to these murky surveillance practices and third-party Web cookies. I n the mobile domain, third-party ad libraries emerged as relatively easy scapegoats. Responding to public pressure, major services and platforms, such as Google, Mozilla, and Apple, loudly declared support for PETs aimed at restricting data collection by third parties, particularly throttling the powers of third-party cookies in Web browsers; Facebook, on its platform, restricted apps' access to user profiles.

Before explaining our reasons for questioning this principle as a privacy solution, it should be noted that advancing technical measures, which throttle the

[20] *See* Kirsten Martin & Helen Nissenbaum, *Measuring Privacy: An Empirical Test Using Context to Expose Confounding Variables* , 18 COLUM. SCI. & TECH. L. REV. 176, 214 (2016).

[21] Kirsten Martin & Helen Nissenbaum, *Privacy Interests in Public Records: An Empirical Investigation*, 31 HARV. J.L. & TECH. 111, 141 (2017); Ido Sivan-Sevilla, Helen Nissenbaum & Patrick Parham, Public Comment on FTC's Commercial Surveillance ANPR (Nov. 21, 2022), https://nissenbaum.tech.cornell.edu/papers/FTC_Public_Comment%20_Nov_25.pdf [https://perma.cc/4D9Y-GD8U].

[22] Jennifer Valentino-DeVries, *What They Know about You - Personal Information Tracked Online*, WALL ST. J., July 31, 2010, https://www.wsj.com/articles/SB10001424052748703999304575399041849931612 [https://perma.cc/G3Y2-9U8S].

spread of personal data to third parties, may significantly reduce the propagation of personal data. Removal of third-party cookies, for example, makes it difficult for Web trackers to link visits by the same browser to different websites and consequently makes it difficult to create comprehensive profiles of users' browsing activities. Requiring apps to ask for permission to access the phone's identifier makes it more difficult for advertising libraries (SDKs) to link different apps installed on the same phone and consequently makes it difficult for them to create comprehensive profiles of users' phone activities.[23] As we explain below, limiting the collection of user data to these actors as a privacy solution has less to do with their third-party status and more to do with being a non-contextual actor and limiting the possibility of the use of that data for non-contextual goals and purposes.

1. *Clarity and Morality: First versus Third Parties*

There are reasons to question whether this family of PETs achieves its eponymous aim of enhancing privacy. To begin, the first-party, third- party distinction strays far from the ideals of clarity and rigor; furthermore, it does not provide a consistent normative foundation for disparate privileges to personal data. Where the first-party, third-party dichotomy is given operational precision in formal technical terms, thereby approaching the first benchmark, it ultimately runs afoul of the third benchmark—fidelity to common usage—as demonstrated in the empirical studies (see Part IV. A.).

In 2012, the first- party, third- party distinction was already called into question as a foundation for policy or technology design. According to Mayer and Mitchell,[24] developments in browser technologies and the evolution of Web-based business models within the broader political economy of data had blurred the distinction in significant ways. For one, major Internet companies act as both a first and a third party.[25] Platforms such as Facebook assert a first-party privilege with users interacting directly on its website.[26] At the same time, it acts technically as a third party with users accessing independent websites that contain the "Like" button. Similarly, Google is a first party when users interact with, say, Google Maps, and a third party when users interact with a website, for example *The Washington Post,* that includes a Google-provided analytics script. Yet, as long as these services operate as metaphorical federations under the same top-level

[23] Patience Haggin, Keach Hagey & Sam Schechner, *Apple's Privacy Change Will Hit Facebook's Core Ad Business. Here's How*, WALL ST. J., Jan. 29, 2021, https://www.wsj.com/articles/apples-privacy-change-will-hit-facebooks-core-ad-business-heres-how-11611938750 [https://perma.cc/JTQ6-XNNR].

[24] Jonathan R. Mayer & John C. Mitchell, *Third-party Web Tracking: Policy and Technology*, *in* IEEE SYMP. ON SEC. AND PRIV. 413, 413 (2012).

[25] *Id.* at 415.

[26] *Id.*

domains, the technical interpretation would allow respective companies to assert a first-party relationship with users across all these services, even when it's acting as a third-party behind the scenes. This set up flies in the face of an intuitive definition of first party as the party the user intends to interact with. If the ethical justification for prioritizing first-party access is users' intention, this reasoning, by which third parties are granted first party status merely due to corporate ownership structure, is flawed.

Finally, the contemporary political economy of commercial data industries has allowed an aggressive pursuit of mergers and acquisitions by corporate titans ("big tech"), such as Google, Microsoft, Amazon, and Apple. As a result, these firms control vast and enormously diverse data holdings over which (for the most part) they assert first-party entitlements. These entitlements allow big tech firms to combine personal information from these disparate sources, irrespective of whether they share a common, top-level Web domain.[27] A privacy pledge to throttle third-party access rings hollow in a world of first parties with access to ever-increasing aspects of individuals' lives.[28] Not only does the first principle -- keeping out third-parties -- fail the benchmarks of clarity and common usage, it yields a concept with questionable moral substance. Through the lens of CI, first parties can be non-contextual actors when an ownership designation trumps moral guidance for privacy expectations.[29]

With social media, idiosyncrasies of the contemporary data economy exacerbate the disjuncture between the technical first- and third-party distinction and historically based, normative underpinnings of data privilege. According to the former, platforms, such as Facebook, Instagram, Gmail, and YouTube, would qualify as first parties because users seek these sites explicitly and intentionally. On this basis, social media companies could claim far-reaching rights over user-generated data and, as we know, have asserted these claims, only recently constrained by general privacy regulations in Europe and the U.S.

Traditional communication providers, including postal services and telecommunication companies, by contrast, are bound by explicit limits on access to content that users create. This makes sense. Even when the first action a user performs is picking up a phone, it is the party on the other end who is the intended

[27] *Cf.* Helen Nissenbaum, *Respect for Context as a Benchmark for Privacy Online: What It Is and Isn't*, *in* CAHIER DE PROSPECTIVE 19, 19–28 (2014).

[28] Although it lies outside the scope of this Article to develop this point, we wish to draw attention to important literature that has deeply influenced our thinking. *See generally* Reuben Binns & Elettra Bietti, *Dissolving Privacy, One Merger at a Time: Competition, Data and Third-Party Tracking*, 36 COMPUT. L. & SEC. REV. 1 (2020); JULIE COHEN, BETWEEN TRUTH AND POWER (2019); Amy Kapczynski, *The Law of Informational Capitalism*, 129 YALE L. J. 1460 (2020).

[29] *See* Danielle Keats Citron & Daniel J. Solove, *Privacy Harms*, 102 B.U. L. REV 793, 797 (2022) ("[P]eople's expectations may be betrayed, resulting in their data being shared with third parties that may use it in detrimental ways—although precisely when and how is unknown.").

recipient of communications and not the service intermediary. Recognizing the technical access that telecommunications providers could exercise, Congress has restricted companies' ability to listen to user communications unless the interception is incident to providing the service or necessary to aid law enforcement.[30] This has resulted in a striking inconsistency between the regulatory treatment of telecommunications and edge providers, respectively, which could be explained as an artifact of the disparate historical jurisdictions of the Federal Trade Commission and the Federal Communications Commission. Although pursuing the regulatory issues is beyond the scope of this Article, the intuitive parallel we wish to draw is that people do not conceive of themselves as communicating with platforms (as first parties); rather, platforms serve as intermediary third parties in communications between friends, groups, and communities.[31]

The technical distinction between first parties and third parties, which maps onto Web domains, misfires in the other direction, too. Large data companies, such as Google, Meta, and Amazon, host properties and assets they own or manage at a variety of domains. Content belonging to major Web platforms —including videos and images—are hosted at multiple domains that may be different from their first-party domains (e.g., facebook.com and fbcdn.net), relying on cross-origin mechanisms to orchestrate their Web applications.[32] As a result, enforcement by Web browsers of domain-based separation between "first" and "third" parties does not map neatly onto modern Web-based systems. Despite vocal commitments to block third-party cookies, among a suite of privacy enhancing techniques, it is unlikely that browser companies would go so far as to break these existing Web-based distributed systems. Nor should they, from the perspective of CI, because third-party domains can be legitimate recipients of data from and about website visitors; in other words, such domains can be contextual actors.

Other instances of legitimate third-party, contextual actors may be outside businesses that receive consumer (visitor, customer, client) data in order to provide necessary functionality, improve services, or protect the interests of businesses and

[30] *See* 18 U.S.C. § 2511(2).

[31] The popularity of end-to-end encrypted messaging services may be a revealing indicator of public sensibilities about whom the first parties are. *See generally* James B. Speta, *A Common Carrier Approach to Internet Interconnection*, 54 FED. COMM. L.J. 225 (2001); Adam Candeub, *The Common Carrier Privacy Model UCDL*, 51 U.C. DAVIS L. REV. 805 (2018); Jack M. Balkin, *How to Regulate (and Not Regulate) Social Media*, 71 J. FREE SPEECH L. 1 (2021).

[32] For example, Google owns the Chrome Web browser, the Android mobile OS, multiple first parties (e.g., Google search engine, Gmail, Google maps), and multiple third parties (e.g., Doubleclick). To give one concrete example, location is an especially valuable piece of information about users, uniquely available on mobile phones. Location is available to "conventional" first and third parties via standard APIs controlled by standard access-control permissions. Yet Google's and Apple's mobile operating systems have privileged access to fine-grained location information collected by the device, beyond what is available to the apps (i.e., first parties) and not controlled by the standard permission mechanism.

their customers. These could include payment processors, cybersecurity companies, or analytics companies, as long as their practices demonstrably serve respective contextual ends and purposes. Critics have also pointed out that under the guise of protecting data subjects from third parties, firms may hide their opportunistic behavior. According to Rory Van Loo, a legal scholar, Amazon (ab)uses the principle favoring first parties over third parties to deny consumer data to Sonos, a third-party manufacturer of speakers used with Amazon's digital assistant. When Sonos requested anonymous error-rate data available through Alexa in order to improve the quality of its speakers when used with Alexa,[33] Amazon denied this request by citing user privacy, despite users benefitting from Sonos having access to this data and improving their services.[34] Van Loo offers an alternative explanation for withholding the data—namely, to give Amazon's own smart-speaker devices a competitive advantage.[35]

In sum, to demonstrate their earnest commitment to privacy, data processors promise to install PETs that would restrict data from being accessed by or sold to third parties. One would expect that defining what they mean by first and third parties in formal or technical terms would add clarity and rigor. Instead, the distinction appears to introduce conceptual inconsistency and fails the criteria of the first benchmark. The second benchmark seems also to pose challenges to this principle. It is perplexing, too, when they keep out those who ought to be let in, and vice versa. At worst, this principle gives convenient political cover to dominant actors in the data industry to pursue their prior interests. At best, it highlights a genuine aspect of privacy, which merely needs sharpening here and there. Setting skepticism aside, in Part III.A we take up the question of how closely this principle tracks the meaning of privacy to nonexpert users whose attitudes we have studied.

2. Law and Regulation

Law and regulation have adopted versions of the first-party, third-party distinction in regulating data flows.[36] For example, regulations have directly

[33] Van Loo, *supra* note 7, at 24.

[34] *Id.* The sharing and use of data for the benefit of the consumer or more generalized benefits (e.g., public health) is found to engender trust and be within privacy expectations. Kirsten Martin, *Privacy Governance for Institutional Trust (Or Are Privacy Violations Akin to Insider Trading?)*, 96 WASH. U. L REV. 1367, 1382 (2018).

[35] Van Loo, *supra* note 7, at 24 ("After all, Amazon itself recorded people's conversations in their homes without users' permission or even awareness. Moreover, Amazon shared actual recordings of consumers' in-home conversations with independent consultants it had hired—thereby handing over much more sensitive data to third parties than what Sonos requested. Amazon's broader behavior with respect to data thus suggests Amazon may have been using privacy as a pretext to keep anonymized voice data from Sonos.").

[36] Within criminal law, the oft-cited third-party doctrine explicitly uses the designation of third-party to mean a lack of privacy expectations. Third parties are any actor collecting or receiving data other than the individual themselves. *See* Orin S. Kerr, *The Case for the Third-Party Doctrine*, 107 MICH.

addressed the conditions under which companies can share consumer information with third parties. The California Consumer Privacy Act (CCPA) affords California consumers the right to request a business that sells their information to a third-party to disclose what information was collected and sold and, under certain circumstances, direct the business not to sell that information to third parties.[37]

Where CCPA incorporates third parties in determining when notification or consent is necessary, the Federal Trade Commission's jurisprudence incorporates sharing data with third parties as a form of a broken promise.[38] In particular, Solove and Hartzog note that the FTC also has enforced promises to consumers not to share data with third parties during bankruptcy proceedings.[39]

Courts have also addressed the role of third parties and privacy interests of individuals within consumer law. For example, in *hiQ Labs v. LinkedIn*, the court rejected LinkedIn's claims that their users had a privacy interest in the data that was shared and collected while on the platform.[40] HiQ identified any LinkedIn users whose activities suggested they were looking for employment. HiQ could then sell that knowledge to the users' employer. The courts found little evidence of LinkedIn users' privacy interest in the information disclosed on LinkedIn.[41] Alternatively, in a case where plaintiffs challenged Facebook's tracking of Facebook users on third

L. REV 561, 569 (2009); Neil Richards, *The Third-Party Doctrine and the Future of the Cloud*, 94 WASH. U. L. REV. 1441, 1441 (2017); Monu Bedi, *Facebook and Interpersonal Privacy: Why the Third-Party Doctrine Should Not Apply*, 54 B.C. L. REV. 1, 12 (2013); Paul Ohm, *The Many Revolutions of Carpenter*, 32 HARV. J.L. & TECH. 357, 362 (2018).

[37] Inferences that are drawn about the consumer are included as a type of personal information covered by the CCPA. *See* Jordan M. Blanke, *Protection for 'Inferences Drawn': A Comparison Between the General Data Protection Regulation and the California Consumer Privacy Act*, 2 GLOB. PRIV. L. REV. 81, 81 (2020).

[38] Daniel J. Solove & Woodrow Hartzog, *The FTC and the New Common Law of Privacy*, 114 COLUM. L. REV. 583, 628–29 (2014) ("Much of the FTC's privacy jurisprudence is based upon a deception theory of broken promises. Some of these promises are explicit and clear, such as when a company violates its own privacy policy, so the determination of a violation requires little interpretation. The types of broken promises cases include…Promises to maintain confidentiality or to refrain from disclosing information to third parties.").

[39] First Amended Complaint for Permanent Injunction and Other Equitable Relief at 2–3, FTC v. Toysmart.com, LLC, No. 00-11341-RGS (D. Mass. July 21, 2000) (describing privacy policy not to disclose personal information to third parties); *see also In re* Toysmart.com, FTC File No. X00 0075, No. 00-11341 RGS (F.T.C. July 21, 2000) (Swindle, Comm'r, dissenting), available at https://www.ftc.gov/sites/default/files/documents/cases/toysmartswindlestatement_0.htm [https://perma.cc/APM4-UXMX] ("Toysmart promised its customers that their personal information would *never* be sold to a third-party, but the Bankruptcy Order in fact would allow a sale to a third-party. In my view, such a sale should not be permitted because 'never' really means never."). *See* Solove & Hartzog, *supra* note 38, at 629.

[40] *See* HiQ Labs, Inc. v. LinkedIn Corp., 31 F.4th 1180 (9th Cir. 2022); Erika M. Douglas, *The New Antitrust/Data Privacy Law Interface*, 130 YALE L.J. 647, 663 (noting "the courts were skeptical of LinkedIn's claim of user privacy protection, finding little concrete evidence of the privacy harm LinkedIn claimed would occur to users from HiQ's continued access to their profile information.").

[41] *See* HiQ Labs, Inc., 31 F.4th 1180; Douglas, *supra* note 40 (noting "the courts were skeptical of LinkedIn's claim of user privacy protection, finding little concrete evidence of the privacy harm LinkedIn claimed would occur to users from HiQ's continued access to their profile information.").

party sites while Facebook users are logged out, the court found that a user did have reasonable expectations of privacy when browsing the Internet and logged out of Facebook.[42]

In a more recent case, iPhone users claim that Facebook still tracks them for advertising purposes even when they opt out of tracking.[43] This practice is contrary to the settings in the Apple App Store, which recently defaulted to asking users to opt in or opt out of cross-app tracking by third parties for advertising.[44] And in the EU, the European Data Protection Board ruled that Meta's practice of tracking users for behaviorally targeted ads on Facebook and Instagram is not considered a legitimate business practice and that users must consent to such practices.[45] The advertising network was not seen as a part of the same "service" as the social network, and Meta was fined 380 million euros.[46] In this decision, the EU did not privilege the entitlements of the first party (Meta) above those of third parties.[47]

B. Minimizing the Use and Retention of Raw Data

> *"When we say raw data, we typically refer to data that is readily available but cannot be easily used speaking. Raw data is compiled from multiple sources, and different sources can often mean that information is displayed in various formats. "*[48]

> *"What do vegetables and data have in common? They both bring more benefits in their raw form. While standard Google Analytics reports can quickly satisfy your hunger, raw data lets you cook something unique and get fresh insights."*[49]

[42] *See In re* Facebook, Inc. Internet Tracking Litig., 956 F.3d 589 (9th Cir. 2020). The Ninth Circuit found a reasonable expectation of privacy when users were logged out.

[43] Taylor Hatmaker, *Facebook Users Sue Meta, Accusing the Company of Tracking on IOS through a Loophole*, TECHCRUNCH (Sept. 22, 2022, 12:17 PM), https://techcrunch.com/2022/09/22/meta-lawsuit-ios-privacy/ [https://perma.cc/C8F2-QFKV].

[44] *Id.*

[45] *Breaking: Meta Prohibited from Use of Personal Data for Advertising*, NOYB.EU (Jan. 4, 2023), https://noyb.eu/en/breaking-meta-prohibited-use-personal-data-advertising [https://perma.cc/Q6CH-ATPP].

[46] *Id.*

[47] *Id.*

[48] Indrė Jankutė-Carmaciu, *What is Raw Data and How it's Used*, WHATAGRAPH (Aug. 20, 2020), https://whatagraph.com/blog/articles/what-is-raw-data [https://perma.cc/XQ72-9T4D].

[49] Vlada Malysheva, *What is Raw Data and How to Use it*, OWOX (Jan. 25, 2024), https://www.owox.com/blog/articles/what-is-raw-data/ [https://perma.cc/AE62-E44R]. *See also* Shubha Ghosh, *Commercializing Data*, 3 ELON L. REV. 195, 201 & n.31 (2012) (citing GOMULKIEWICZ, NGUYEN & CONWAY-JONES, LICENSING INTELLECTUAL PROPERTY: LAW AND APPLICATION 418–19 (2008)) (proposing a distinction between raw data as "individual facts" and cooked data as "original arrangements of facts") ("Data production can best be divided into raw data and cooked data.").

> *"Raw data, also known as primary data, are data (e.g., numbers, instrument readings, figures, etc.) collected from a source . . . .".*

In practical terms, the second principle, like the first, promises to restrict access to or use of data for the sake of privacy, but instead of restricting it on the basis of the recipient, viz. "third parties," this principle restricts on the basis of whether the data is *raw.* And the type of information or knowledge at issue has consistently been an important facet of privacy theory. For some, the type of information is so important as to dictate privacy norms, as with information labeled "sensitive." [50] However, promises surrounding raw data do not necessarily extend to data that is not raw—presumably processed in some way—such as inferred or derived data, including models or profiles drawn therefrom. And placing protections around raw data, but not inferences, is not a clear privacy-enhancing solution.[51]

Instances of systems and practices that are put forward as privacy-enhancing under this principle include Google's policy of allowing users to delete stored search data attributable to the user while still retaining and using inferences drawn from that search and other behavioral data through custom audience segments. These custom audience segments are inferences about customers based on recent search queries as well as other behavioral data.[52] Advertisers can then target users later based on these inferences. In other words, though Google allows search history to be deleted, inferences based on them still may be used for personalized advertising.

Another example of restrictions on raw-data access is the joint Meta/Mozilla proposal for "interoperable private attribution" to measure ad conversion in a privacy-preserving fashion;[53] still another is the effort at all major tech companies

[50] Paul Ohm, *Sensitive Information*, 88 S. CAL. L. REV. 1125 (2015). However, the privacy norms of information—even that deemed "sensitive" in a survey question—is better understood through privacy as CI and dependent upon the context, actors, and transmission principles. *See* Martin & Nissenbaum, *supra* note 20, at 202-10.

[51] The ability to create new knowledge with widely known raw data is a long-standing issue. *See, e.g.*, Alessandro Acquisti & Ralph Gross, *Predicting Social Security Numbers from Public Data*, 106 PROCS. OF THE NAT'L ACAD. OF SCIS. 10975 (2009) (exploring how Social Security numbers could be inferred from birth data and readily available information from data brokers and social network profiles); Ilaria Liccardi, Alfie Abdul-Rahman & Min Chen, *I Know Where You Live: Inferring Details of People's Lives by Visualizing Publicly Shared Location Data*, *in* PROCEEDINGS OF THE 2016 CHI CONFERENCE ON HUMAN FACTORS IN COMPUTING SYSTEMS 1–12 (2016), https://doi.org/10.1145/2858036.2858272 (exploring how publicly available geographic information from Tweets could accurately infer "average income based on one's neighborhood, average housing cost, debt, and other demographic information, such as political views"); Charles Duhigg, *How Companies Learn Your Secrets*, N.Y. TIMES, Feb. 16, 2012, https://www.nytimes.com/2012/02/19/magazine/shopping-habits.html [https://perma.cc/3H4D-M9DF] (identifying someone as pregnant from purchase history data).

[52] *About Custom Segments*, GOOGLE ADS HELP, https://support.google.com/google-ads/answer/9805516?hl=en (last visited December 2, 2023) [https://perma.cc/3JRL-UK8K].

[53] *IPA End to End Protocol*, GITHUB (July 28, 2022), https://github.com/patcg-individual-drafts/ipa/blob/main/IPA-End-to-End.md [https://perma.cc/A7DV-KKDX]; Martin Thomson, *Privacy*

to develop and deploy federated learning, a family of decentralized machine learning technologies whose full spectrum of uses is yet to be worked out.[54] Here too, announcing that raw data spread is throttled, with no parallel commitment around data inferred from it.

A conception of privacy that would support a prescriptive distinction between raw data and inferences, in our view, suffers similar shortcomings to those of the first principle, namely, a failure to meet Benchmark (1) (rigor and clarity) and Benchmark (2) (solid moral footing). In Part III.A, our empirical studies demonstrate that the raw-data principle is discordant with meaning and significance ascribed to privacy in common usage, too.[55]

1. *Clarity and Morality: Raw Data versus Inferences*

It is clear that *raw* as a quality of data is a suggestive metaphor. The absence of a universally adopted formal definition of raw data is not, itself, a problem, except when ambiguities pull policies, practices, and technical design in different, or even incompatible directions. One account of raw data refers to data that is directly given by users, or is directly collected or captured from them as they engage in data-

Preserving Attribution for Advertising, MOZILLA (Feb. 8, 2022), https://blog.mozilla.org/en/mozilla/privacy-preserving-attribution-for-advertising/ [https://perma.cc/7GTV-A39W].

[54] *See e.g.*, Matthias Paulik, Matt Seigel, Henry Mason, Dominic Telaar, Joris Kluivers, Rogier van Dalen, Chi Wai Lau, Luke Carlson, Filip Granqvist, Chris Vandevelde, Sudeep Agarwal, Julien Freudiger, Andrew Byde, Abhishek Bhowmick, Gaurav Kapoor, Si Beaumont, Áine Cahill, Dominic Hughes, Omid Javidbakht, Fei Dong, Rehan Rishi & Stanley Hung, *Federated Evaluation and Tuning for On-Device Personalization: System Design & Applications*, APPLE (Feb. 2022), https://machinelearning.apple.com/research/federated-personalization [https://perma.cc/F5G4-4WGE]; *Applying Federated Learning to Protect Data on Mobile Devices*, META (June 14, 2022), https://engineering.fb.com/2022/06/14/production-engineering/federated-learning-differential-privacy/ [https://perma.cc/6HDN-EMCM]; Jeff Omhover, *Federated Learning with Azure Machine Learning: Powering Privacy-Preserving Innovation in AI*, MICROSOFT (May 30, 2023), https://techcommunity.microsoft.com/t5/ai-machine-learning-blog/federated-learning-with-azure-machine-learning-powering-privacy/ba-p/3824720 [https://perma.cc/J6DV-S7Q9]; *Federated Learning,* GOOGLEAI, https://federated.withgoogle.com/#about (last visited Oct. 25, 2024) [https://perma.cc/2WM7-32YX].

[55] Here we question the delineation between raw data and created inferences as imprecise and lacking moral weight in terms of privacy judgments. However, others have raised other ethical issues with the creation and use of inferences such as around manipulation, discrimination & "unfair bias," or inability to contest decisions about you and causing due process issues. *See, e.g.*, Ryan Calo, *Digital Market Manipulation*, 82 GEO. WASH. L. REV. 995 (2013); Ido Kilovaty, *Legally Cognizable Manipulation*, 34 BERKELEY TECH. L.J. 449 (2019); Tal Z. Zarsky, *Privacy and Manipulation in the Digital Age*, 20 THEORETICAL INQUIRIES IN LAW 157 (2019); Solon Barocas & Andrew D. Selbst, *Big Data's Disparate Impact*, 104 CAL. L. REV. 671 (2016); Brent Daniel Mittelstadt, Patrick Allo, Mariarosaria Taddeo, Sandra Wachter & Luciano Floridi, *The Ethics of Algorithms: Mapping the Debate*, 3 BIG DATA & SOCIETY (2016); Sandra Wachter, *Affinity Profiling and Discrimination by Association in Online Behavioral Advertising*, 35 BERKELEY TECH. L.J. 367 (2020); Sandra Wachter, *The Theory of Artificial Immutability: Protecting Algorithmic Groups Under Anti-discrimination Law*, 97 TUL. L. REV. 149 (2022); Danielle Keats Citron & Frank Pasquale, *The Scored Society: Due Process for Automated Predictions*, 89 WASH. L. REV. 1 (2014).

generating activities with platforms and services.[56] Data directly provided may[57] include any data that subjects enter into online forms, such as a search term, address, age, religion, and order selections; or any content they may post on social media, such as comments, friends, birthdate, activities, preferences —to name a fraction of the possibilities. Directly captured data could include products of user activities, such as photos, typed texts, and visited URLs, as well as captured sensor data, engagement data, and the myriad types of machine-machine data such as time spent on a website, phone numbers dialed, areas clicked, geolocation positions and paths, heart-rate measurements, and data generated by smartphone operating systems, as well as other types of "metadata."

The second strain of meaning also refers to data directly captured but more narrowly scopes the data conceived as *raw* to data without independent semantics, as it were, machine-interpretable but not human-interpretable data. These could include URLs, IP addresses, geolocation coordinates, Bluetooth and WiFi signals, and a slew of the data generated by smartphones, including data provided to apps that they are not obligated to mention to users. Raw data is used with terms such as *atomic* and *primitive*. The second strain would not count as raw the vast category of semantically rich data directly shared by data subjects with platforms and services, such as religion, search terms, music playlists, photographs, etc.

Setting aside differences between the two interpretations of raw, it remains unclear what exactly are data processors'[58] practices with respect to data that are *not* raw, at best, declared in notoriously vague and ambiguous privacy policies.[59] They might even brag of unfettered publication of processed data, claiming that the

[56] *See e.g.*, Gavin Wright, *Raw Data (Source Data or Atomic Data)*, TECHTARGET (May 6, 2021), https://www.techtarget.com/searchdatamanagement/definition/raw-data. ("Raw data (sometimes called source data, atomic data or primary data) is data that has not been processed for use. A distinction is sometimes made between data and information to the effect that information is the end product of data processing. Raw data that has undergone processing is sometimes referred to as cooked data") [https://perma.cc/RJ4G-B9WB]; Jack Vaughan, *Data*, TECHTARGET (July 31, 2019),https://www.techtarget.com/searchdatamanagement/definition/data [https://perma.cc/3MVM-TK9S];
Robert Sheldon, *Information*, TECHTARGET (May 18, 2021), https://www.techtarget.com/searchdatamanagement/definition/information [https://perma.cc/SGT7-Z9RV].

[57] A website's privacy policy might describe limits on what a company does with such data.

[58] Using this term as defined, for example in the GDPR. Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC, 2016 O.J. (L 119) 1.

[59] Joel R. Reidenberg, Travis Breaux, Lorrie Faith Cranor, Brian French, Amanda Grannis, James T. Graves, Fei Liu, Aleecia McDonald, Thomas B. Norton & Rohan Ramanath, *Disagreeable Privacy Policies: Mismatches Between Meaning and Users' Understand*ing, 30 BERKELEY TECH. L.J. 39 (2015).

data in question are merely aggregate statistics, or even "privacy-preserving" statistics derived under constraints of differential privacy.[60]

Whatever the differences between different interpretations of "raw", they have in common that they serve as the basis for a distinction—between *raw data* and *processed* ("cooked") data. Processed data may cover a range of alternatives, including what data scientists might call derived or inferred data, whether from raw or processed data. It could also include data, such as profiles, proclivities, indices, and so forth, inferred using sophisticated machine learning techniques. Defenders of the second principle might argue that pointing to different interpretations is a mere quibble compared to what these interpretations have in common, namely, a shared commitment to different treatment for data belonging in the underlying categories, respectively—raw vs. processed.

In our view, by contrast, the practical import of the two interpretations of raw versus processed (inferred, derived, "cooked") is significant since firms make different commitments based on this demarcation between raw and processed data. There are broad swaths of data that either would or would not be assured protection under the respective interpretations.[61] For example, all the data with semantic meaning that constantly flows to data controllers in the course of everyday activities is either covered by or not covered by the second principle, depending on interpretation. With assured restrictions yielding allowable practices miles apart under respective meanings of raw, the conception of privacy resting on these is disturbingly indeterminate in meaning and application, thereby failing to meet the second benchmark requiring clarity and rigor. Our empirical studies, discussed in Part III.A, further reinforce this failure.

For privacy as contextual integrity, the appropriateness of the data flow is judged based partly on the type of information shared, collected, known, or generated. For example, grades and previous coursework are appropriate types of information for the education context but income or height would not be appropriate to collect and use within the education context. Similarly, symptoms and diagnoses are appropriate for the health context, but not for the education

[60] For a more detailed analysis of how differential privacy assuages privacy needs, *see* Jeremy Seeman & Daniel Susser, *Between Privacy and Utility: On Differential Privacy in Theory and Practice* (Oct. 31, 2022) (unpublished manuscript), https://ssrn.com/abstract=4283836.

[61] Scholars have previously identified how simple definitions of privacy can miss the privacy violations involved in creating and using inferences. For how a rights approach misses privacy issues with inferences, see Ari Ezra Waldman, *Privacy's Rights Trap*, 117 NW. U. L. REV. ONLINE 88, 94 (2022); Daniel J. Solove, *The Limitations of Privacy Rights*, 98 NOTRE DAME L. REV. 975, 986 (2022). For how traditional approaches miss inferences, see IGNACIO COFONE, THE PRIVACY FALLACY: HARM AND POWER IN THE INFORMATION ECONOMY (2023). For how sensitive data labels miss harm of inferences, see Daniel J. Solove, *Data is What Data Does: Regulating Based on Harm and Risk Instead of Sensitive Data*, 118 NW. U. L. REV. 1081, 1081 (2024).

context, so a firm attempting to ensure that the type of information is appropriate for the context in which they are engaging with a user would be lauded.

However, in terms of the second benchmark, the delineation between raw data versus inferences does not hold normative force in dictating privacy norms. The appropriate information type for a given context may be raw or inferred. For example, a medical diagnosis is inferred from raw data and other inferred data. Yet such information is completely appropriate for a doctor or nurse to create and use for the purpose of treatment.

On the other hand, inferences may be the source of privacy violations. Harms and other disturbing practices may undoubtedly be traced to direct access to, and use and retention of, raw data, such as identity theft due to data breaches, ads based on search terms, surveillance triggered by sensor data, unfair treatment based on directly shared data like race, religion or health status, and inappropriate exposure through location data.[62] Many inappropriate data practices such as these that have drawn the attention and condemnation of privacy researchers, advocates, and regulators are attributable to processed or derived data, including inferences, models, profiles, indices, and scores. Virtually the entire edifice of behavioral advertising rests on inference about users' mental states (e.g., purchasing intent, vulnerability to manipulation), vast abuses of unfair decision systems (e.g., housing, jobs, parole, insurance, credit, price discrimination, etc.) rest on data-derived profiles, and a never-ending stream of shocking stories reveals people's vulnerability to inferences drawn about them from unscrupulous mobile apps, such as mental health, fertility tracking, religious observance, and life milestones.[63]

[62] When Michael Hayden, head of the National Security Agency from1999 to 2005, famously said, "We kill people based on metadata," he affirmed the power to infer reliable incriminating evidence through social network activity. *See* Johns Hopkins University, *The Johns Hopkins Foreign Affairs Symposium Presents: The Price of Privacy: Re-Evaluating the NSA*, YOUTUBE (Apr. 7, 2014), https://www.youtube.com/watch?v=kV2HDM86XgI [https://perma.cc/33MC-KDY8].

[63] Danielle Keats Citron & Daniel J. Solove, *Privacy Harms*, 102 B.U. L. REV. 793 (2022); Drew Harwell*, Is Your Pregnancy App Sharing Your Intimate Data with Your Boss?,* WASH. POST, Apr. 10, 2019, https://www.washingtonpost.com/technology/2019/04/10/tracking-your-pregnancy-an-app-may-be-more-public-than-you-think/ [https://perma.cc/53F9-RNVX]; *FTC Enforcement Action to Bar GoodRX from Sharing Consumers' Sensitive Health Info for Advertising*, FTC (Feb. 1, 2023), https://www.ftc.gov/news-events/news/press-releases/2023/02/ftc-enforcement-action-bar-goodrx-sharing-consumers-sensitive-health-info-advertising [https://perma.cc/4NSJ-NBY8]; *FTC to Ban Betterhelp from Revealing Consumers' Data, Including Sensitive Mental Health Information, to Facebook and Others for Targeted Advertising*, FTC (Mar. 14, 2023), https://www.ftc.gov/news-events/news/press-releases/2023/03/ftc-ban-betterhelp-revealing-consumers-data-including-sensitive-mental-health-information-facebook [https://perma.cc/F3XU-9NF7]; Khadeeja Safdar, *Churches Target New Members, with Help from Big Data*, WALL ST. J., Dec. 26, 2021, https://www.wsj.com/articles/churches-new-members-personal-online-data-analytics-gloo-11640310982 [https://perma.cc/BFK7-7737].

The point is not that the principle itself is ethically problematic; it is that swearing to abstemious practices in the treatment of raw data, while leaving inferred (processed, derived) data out in the cold, as it were, does little to meet the demands of the third benchmark—the ability to explain privacy's ethical force. Our empirical research, reported in Section III.A. reinforces this point.

2. Law and Regulation

The delineation of raw data versus inferences is addressed in privacy regulation as well. As Blanke notes, these laws and regulations are attempting to acknowledge that inferences drawn from data about the individual can become more dangerous to privacy than the vast collection and storage of the data itself.[64] For example, the CCPA specifically includes "inferences drawn" as part of its definition of personal information[65] and also includes the right to know inferences drawn about oneself. [66] In general, the CCPA allows individuals in California to find out what information is collected about them and to opt out of the transfer or sale of that information. The definition of information about the consumer includes "inferences drawn from any of the information identified in this subdivision to create a profile about a consumer reflecting the consumer's preferences, characteristics, psychological trends, predispositions, behavior, attitudes, intelligence, abilities, and aptitudes."[67] Inferences drawn about an individual are then treated as personal information within the CCPA.

Alternatively, the GDPR is more limited in addressing the creation of new knowledge through inferences, but does focus more on protecting the raw data collected and stored —including "the right to access; the right to rectification; the right to erasure (the 'right to be forgotten'); the right to restriction of processing; and the right to data portability."[68] Within the GDPR, data is defined as either

[64] *See generally* Jordan M. Blanke, *Protection for 'Inferences Drawn': A Comparison between the General Data Protection Regulation and the California Consumer Privacy Act*, 2 GLOB. PRIV. L. REV. 81 (2020).

[65] *Id.*

[66] CAL. ATT'Y GEN. OP. NO. 20-303 (March 10, 2022), https://oag.ca.gov/system/files/opinions/pdfs/20-303.pdf ("under the California Consumer Privacy Act, a consumer has the right to know internally generated inferences about that consumer, unless a business can demonstrate that a statutory exception to the Act applies") [https://perma.cc/2PPD-PK95].

[67] CAL. CIV. CODE § 1798.140(m) (2021). The Attorney General's opinion above was the first to consider the statute and found that "for purposes of the CCPA, 'inference' means 'the derivation of information, data, assumptions, or conclusions from facts, evidence, or another source of information or data.' An inference is essentially a characteristic deduced about a consumer (such as 'married,' 'homeowner,' 'online shopper,' or 'likely voter') that is based on other information a business has collected (such as online transactions, social network posts, or public records)." The inference provisions are now codified at CAL. CIV. CODE § 1798.140(r), (K) (2024).

[68] Blanke, *supra* note 64, at 85.

provided by the individual or observed about the individual. In contrast, inferences are considered derived data where new knowledge is developed from the observed or provided data.[69] According to Wachter and Mittlestadt, inferences are not afforded the same protections under the GDPR as for provided or observed data, including rights to notification and deletion. In order for data to be protected by the GDPR, that data must be deemed "personal."[70] And inferences are not as clearly defined or protected under the GDPR as they are under the CCPA.[71]

In 2022, the Federal Trade Commission stated their intent to enforce laws against the illegal use and sharing of consumer data, including the use of data to create inferences.[72] The FTC then filed a complaint against the location data broker Kochava and included the possible harmful creation of inferences as to consumers' LGBTQ+ identity and visits to medical facilities.[73] For laws and regulations that include a consumer harm component, a violation of privacy can constitute a consumer harm. For example, in antitrust regulations, the abuse of power may be through data governance and privacy policies harming an individual,[74] or with the FTC's unfairness doctrine, consumer harm can include privacy violations. For example, a search engine can abuse their power in a noncompetitive market by offering inefficient and low-quality search results and they can also abuse their power through privacy practices that benefit the firm. However, only recently have courts turned their attention to whether companies abuse their market power

[69] Sandra Wachter & Brent Mittelstadt, *A Right to Reasonable Inferences: Re-Thinking Data Protection Law in the Age of Big Data and AI*, 2019 COLUM. BUS. L. REV. 494, 516 (2019) ("derived (e.g. country of residency derived from the subject's postcode) and inferred data (e.g. credit score, outcome of a health assessment, results of a personalization or recommendation process) are not 'provided by' the data subject actively or passively, but rather created by a data controller or third-party from data provided by the data subject and, in some cases, other background data").

[70] *Id.*

[71] Blanke, *supra* note 64, at 86.

[72] *See, e.g.,* Kristin Cohen, *Location, Health, and Other Sensitive Information: FTC Committed to Fully Enforcing the Law against Illegal Use and Sharing of Highly Sensitive Data*, FTC (July 11, 2022), https://www.ftc.gov/business-guidance/blog/2022/07/location-health-and-other-sensitive-information-ftc-committed-fully-enforcing-law-against-illegal,
("data aggregators and brokers – companies that collect information from multiple sources and then sell access to it (or analyses derived from it) to marketers, researchers, and even government agencies. These companies often build profiles about consumers and draw inferences about them based on the places they have visited.") [https://perma.cc/EK7Q-8FBF].

[73] FTC v. Kochava Inc., No. 2:22-cv-00377-DCN, at 6, 9 (D. Idaho August 29, 2022) ("Precise geolocation data associated with MAIDs, such as the data sold by Kochava, may be used to track consumers to sensitive locations, including places of religious worship, places that may be used to infer an LGBTQ+ identification, domestic abuse shelters, medical facilities, and welfare and homeless shelters.") ("Consumers have no insight into how this data is used – they do not, for example, typically know or understand that the information collected about them can be used to track and map their past movements and that inferences about them and their behaviors will be drawn from this information.").

[74] *See generally* Kirsten Martin, *Platforms, Privacy & The Honeypot Problem*, 37 HARV. J. L. & TECH. 1087 (2024).

through the creation and use of *inferences*—and whether the creation and use of inferences is considered a privacy violation. In fact, platforms have been able to "circumvent a narrower interpretation of special category data" within the EU by focusing on proxies and inferences, which had been thought to be outside protected data categories.[75] Recent rulings in the EU, however, have placed special protections required on any data that could lead to an inference that would be within a "special category" of data. For example, having information about one's spouse is not a special category; yet sexual orientation is a special category which could be inferred from spousal information.[76]

C. Leaving Personal Data on the Device

> *"Federated learning is a privacy-enhancing technology that we use to improve models on device without sending users' raw data to Google servers. Google Assistant uses federated learning to improve "Hey Google."*[77]

> *"Face recognition and scene and object detection are done completely on your device rather than in the cloud. This allows Apple to provide you with these advanced features without accessing your photos. And apps can access your photos only with your permission."*[78]

The first principle derogates third parties; the second elevates raw data; the third adds a twist, holding even first parties at bay. Apple may be credited for popularizing the idea that "your data never leaves your device," or "your data stays on your device," [79] but Google, which has spearheaded the development of

[75] Natasha Lomas, *Sensitive Data Ruling by Europe's Top Court Could Force Broad Privacy Reboot*, TECHCRUNCH (Aug. 2, 2022), https://techcrunch.com/2022/08/02/cjeu-sensitive-data-case/ [https://perma.cc/DGW3-H9VG].

[76] *Id.* ("The relevant bit of the case referral to the CJEU related to whether the publication of the name of a spouse or partner amounted to the processing of sensitive data because it could reveal sexual orientation. The court decided that it does. And, by implication, that the same rule applies to inferences connected to other types of special category data.") ("Examples of inferences could include using the fact a person has liked Fox News' page to infer they hold right-wing political views; or linking membership of an online Bible study group to holding Christian beliefs; or the purchase of a stroller and cot, or a trip to a certain type of shop, to deduce a pregnancy; or inferring that a user of the Grindr app is gay or queer." Lomas, "Sensitive Data Ruling")("Examples of inferences could include using the fact a person has liked Fox News' page to infer they hold right-wing political views; or linking membership of an online Bible study group to holding Christian beliefs; or the purchase of a stroller and cot, or a trip to a certain type of shop, to deduce a pregnancy; or inferring that a user of the Grindr app is gay or queer.").

[77] GOOGLE, https://support.google.com/assistant/answer/10176224?hl=en [https://perma.cc/GNA6-UHQD].

[78] APPLE, https://www.apple.com/privacy/features/ [https://perma.cc/6AD9-S6QE].

[79] *Id.* Apple planned to deploy a technology that would scan photos on user's phone for illegal child abuse material (CSAM). The only information that would leave the device is whether any of the photos matched a secret CSAM database. Deployment of this technology was postponed indefinitely

federated learning, has been a vocal proponent, too. In framing its Privacy Sandbox suite of technologies, Google promises to maintain Chrome usage data within the user's instance of the browser and on "your" device, whether mobile, laptop, or desktop.[80] The key principle is that users should have no privacy concerns regarding the vast swaths of data generated about themselves in their daily, minute-by-minute intimate uses of their mobile devices and browsers because this data never leaves their device. For the Privacy Sandbox, some inferences derived from raw data remain on a consumer's device and are accessed only by the browser (on the device).[81] However, other inferences such as special topics can be queried by third parties.

As a concrete example, Google's Privacy Sandbox exposes limited information about the user's interests (inferred from their browsing activities tracked by the browser) via a Topics API, which is a curated list of categories.[82] This API can be accessed by Web parties and therefore provides an avenue for inferences about the user to leave the device. Another way in which the Privacy Sandbox indirectly exposes inferences about users is via the Protected Audience API.[83] This technology enables website s to create custom categories and assign users to them based on their interactions, with the site acting as the first party (note the connection with the first principle). Custom categories created via this API are not curated by Google and enable behavioral advertising to target a much broader set of inferences than the Topics API. Unlike the Topics API, these inferences are not linkable to specific users and cannot be accessed by Web parties via the browser (and, therefore, stay on the device).[84]

In offering these assurances, tech companies seem to be supporting a version of "privacy as secrecy.' Firms are claiming to keep the raw data away from prying eyes and may even ensure that the first party—the firm that manufactured the device—does not have access to that raw data.[85] In this way, the third principle is a combination of the first two: limiting who has access to the knowledge and focusing

due to multiple controversies, including the possibility of false matches and the risk that this technology could be re-purposed to scan for other material (e.g., censorship by oppressive governments).

[80] *Privacy Sandbox on Android*, GOOGLE, https://privacysandbox.com/intl/en_us/android/ (last visited October 30, 2024) ("Topics are selected entirely on your device, so the information about the apps you use isn't shared with external parties.") [https://perma.cc/EB92-62WU].

[81] *Id.*

[82] *Topics API: Relevant Ads Without Cookies*, GOOGLE, https://privacysandbox.com/proposals/topics/ (last visited December 2, 2023) [https://perma.cc/BME9-2377].

[83] *Protected Audience API Origin Trial and AdSense*. GOOGLE ADSENSE HELP, https://support.google.com/adsense/answer/12570693?hl=en (last visited December 2, 2023) [https://perma.cc/G9HQ-8Y5H].

[84] *Id.*

[85] *Privacy Sandbox on Android*, GOOGLE, https://privacysandbox.com/intl/en_us/android/ (last visited October 28, 2024) ("Topics are selected entirely on your device, so the information about the apps you use isn't shared with external parties.") [https://perma.cc/RJN2-89J4].

on the use of inferences rather than raw data. Nevertheless, in order for the device to provide the range of functionalities that make it useful, profitable, and even necessary for modern living, it must support a vast spectrum of bidirectional data flows.

The critical question is: what data actually leaves your device and is sent to central servers owned by tech corporations, where it is processed and used to provide functionality and generate value for those companies? While we may be assured that raw data—such as the list of URLs we visit, daily step counts, or heart rates—does not make this journey, inferences drawn from that data presumably often do. Although tech companies have been forthright in detailing instances of privacy-enhancing applications of federated learning, such as word prediction in "smart" keyboards, they have not disclosed how it might be used in other instances, such as profiling user behavior. For Google's Sandbox, the inferences drawn about the user remain on the device—but this is not necessarily true for all claims of 'the data stays on your device.' This lack of precision troubles the second privacy benchmark, for we are left wondering what data does and does not flow, to whom, and for what purpose?

The second benchmark, requiring privacy to have moral clout, also challenges the conception underlying this principle. Drawing this conclusion from a parallel point that we made in the preceding section, we argue that if *derived* data (models, inference, profiles, etc.) is at least as powerful and valuable as the raw data from which it is derived, then whether the latter sits on the device does not lessen the impact of the former on the individual. Harms in question may still involve unfair discrimination, manipulation, and exploitation of individual users, and challenge the values and purposes of important societal institutions.[86]

II. SHEDDING LIGHT THROUGH EMPIRICAL STUDY

A. Study Design

Our studies are structured around the conception of privacy as CI. Readers unfamiliar with CI may find fuller accounts[87] useful, however, it should be possible

[86] HELEN NISSENBAUM, *Breaking Rules for Good*, *in* PRIVACY IN CONTEXT: TECHNOLOGY, POLICY, AND THE INTEGRITY OF SOCIAL LIFE 158 (2010).

[87] *See generally id.*; Kirsten Martin & Helen Nissenbaum*, Privacy Interests in Public Records: An Empirical Investigation*, 31 HARV. J.L. & TECH. 111 (2017); Martin & Nissenbaum, *supra* note 20; Kirsten Martin & Helen Nissenbaum, *What Is It About Location?*, 35 BERKELEY TECH. L.J. 251 (2020).

to grasp key findings without consulting other sources. The scenarios we have developed for testing respondents' answers are situated in intuitively recognizable social domains through social role s and attributes that are typical of these respective domains. Further, in certain of the test conditions, to trigger respondents' evaluations, or contextual expectations, the scenarios refer to the purposes served by the information flows in question. In the latter, we ascertain users' assessment of appropriateness of data flows according to contextual functions, purposes, and values either served or disserved by them.

To learn whether the principles of these PETs embody a conception of privacy with rough fidelity to common usage—one of the benchmarks of a sound conception of privacy—through systematic empirical inquiry, we have sought to characterize these concerns in terms of attitudes to data flows described in terms the parameters of CI to learn what factors that are significant. We note, however, that the two design patterns that inform our study designs, (1) denying access to data to third parties and (2) using inferences rather than raw data do not have anything to say about (contextual) ends, purposes, and values. In fact, many of these technologies are explicitly designed so that purposes do not change. The goal is still to enable behavioral advertising and train machine learning models on users' data. Even though we have narrowed the scope of our studies to claims of membership in the class of PETs on the grounds of (1) and (2), and even though the proponents of these approaches tend not to discuss "use for a purpose," we have included purposes as a factor in our studies because of the important role they play in CI as ground for the ethical standing of a given information practice.

B. Methods

We applied the factorial vignette methodology[88] to assess how well PETs, corresponding to the three principles, meet the privacy expectations of respondents, and in turn the benchmarks laid out above, with particular focus on the first and second. Factorial vignette surveys present respondents with a series of twenty to forty vignettes in which multiple factors are systematically varied in order to test their relative importance to respondents' judgments. These factors constitute the independent variables of our study.

The factors chosen for our study correspond to a subset of the contextual factors (or parameters) of CI. For each vignette, values for the parameters are varied. After seeing each vignette, respondents are asked to complete a simple rating task—the degree to which a scenario is appropriate or "okay"—from which we later extract the statistical relevance of each of the factors.

These vignettes systematically and simultaneously varied in the type of data collected, the contextual actor collecting data, the type of inference drawn using the collected data, and how the knowledge about the individual was used. We used the following factors:

- Contextual Actor: We varied the contextual actor that the individual in the vignette would interact with. We included a search engine, a browser, two social networks (one focused on photos/videos and one focused on friends/acquaintances), and a news site. Note that the term "contextual actor" refers to individuals acting in context relevant capacities, like physician, student, and web searcher.

[88] Guillermina Jasso, *Factorial Survey Methods for Studying Beliefs and Judgments*, 34 SOCIO. METHODS & RSCH. 334, 340 (2006) ("In the factorial survey approach, each respondent is asked to rate the level of a specified outcome variable (such as healthiness or wage attainment or just prison sentence) corresponding to a fictitious unit (a person, say, or a family), which is described in terms of potentially relevant characteristics such as age, gender, study or eating habits, access to medical care or housing, and the like. The respondent is presented a large set of these fictitious units, termed *vignettes*. Statistical techniques are used to retrieve the equation implicitly used by each respondent in assigning the level of the outcome variable to each vignette.").

- Data: We varied the type of data that was collected initially and included search terms,[89] location,[90] likes/clicks/engagement,[91] profile information, and Web activity.[92]

- Inference: Some of the vignettes, referred to inferences that were derived from information collected. These inferences included demographics (age, income, family status, etc.),[93] emotions (or mood),[94] friends and activities (types of bars, reading preferences, LGBTQ+ friends, etc.),[95] interests (diabetes concerns, retirement planning, babies, etc.), and medical inferences (recent medical procedures or doctor visits).[96]

- Purpose/Use: The purpose of the data collection or derived inference varied from improving services and placing ads on the site to allowing others to later place ads when online and selling the knowledge to others. We refer to this with the phrase "use for a purpose."

[89] *See* JACOB LEON KRÖGER, ROGUE APPS, HIDDEN WEB TRACKING AND UBIQUITOUS SENSORS 2 (2022).

[90] *See* Kirsten Martin & Helen Nissenbaum, *What Is It About Location?*, 35 BERKELEY TECH. L.J. 251, 259 (2020); Han Bo, Paul Cook & Timothy Baldwin, *Geolocation Prediction in Social Media Data by Finding Location Indicative Words*, *in* PROCEEDINGS OF COLING 2012: TECHNICAL PAPERS 1045, 1046 (2012).

[91] *See* Sandra Wachter & Brent Mittelstadt, *A Right to Reasonable Inferences: Re-Thinking Data Protection Law in the Age of Big Data and AI*, 2019 COLUM. BUS. L. REV. 494, 516 (2019).

[92] *See* KRÖGER, *supra* note 88, at 155-56.

[93] *See* Zijian Wang, Scott A. Hale, David Adelani, Przemyslaw A. Grabowicz, Timo Hartmann, Fabian Flöck & David Jurgens, *Demographic Inference and Representative Population Estimates from Multilingual Social Media Data*, *in* WORLD WIDE WEB CONFERENCE 2056 (2019); Aron Culotta, Nirmal Kumar Ravi & Jennifer Cutler, *Predicting the Demographics of Twitter Users from Website Traffic Data*, *in* 29 PROCEEDINGS OF THE AAAI CONFERENCE ON ARTIFICIAL INTELLIGENCE 72, 72-76 (2015).

[94] *See* Kat Roemmich, Florian Schaub & Nazanin Andalibi, *Emotion AI at Work: Implications for Workplace Surveillance, Emotional Labor, and Emotional Privacy*, *in* PROCEEDINGS OF THE 2023 CHI CONFERENCE ON HUMAN FACTORS IN COMPUTING SYSTEMS 1, 1 (2023); KRÖGER, *supra* note 88, at 120-21; Robert Booth, *Facebook Reveals News Feed Experiment to Control Emotions*, GUARDIAN (June 29, 2014), https://www.theguardian.com/technology/2014/jun/29/facebook-users-emotions-news-feeds [https://perma.cc/W9WK-KJSU]; MEGAN A. MORENO, LAUREN A. JELENCHICK, KATIE G. EGAN, ELIZABETH COX, HENRY YOUNG, KERRY E. GANNON & TARA BECKER, FEELING BAD ON FACEBOOK: DEPRESSION DISCLOSURES BY COLLEGE STUDENTS ON A SOCIAL NETWORKING SITE 2 (2011); Luke Stark & Jesse Hoey, *The Ethics of Emotion in Artificial Intelligence Systems*, *in* PROCEEDINGS OF THE 2021 ACM CONFERENCE ON FAIRNESS, ACCOUNTABILITY, AND TRANSPARENCY 782, 782 (2021).

[95] *See* Sandra Wachter, *Affinity Profiling and Discrimination by Association in Online Behavioral Advertising*, 35 BERKELEY TECH. L.J. 367, 376, 380, 383 (2020).

[96] *See* BEDI GILLINDER, FACUNDO CARRILLO, GUILLERMO A. CECCHI, DIEGO FERNÁNDEZ SLEZAK, MARIANO SIGMAN, NATÁLIA B. MOTA, SIDARTA RIBEIRO, DANIEL C. JAVITT, MAURO COPELLI & CHERYL M. CORCORAN, AUTOMATED ANALYSIS OF FREE SPEECH PREDICTS PSYCHOSIS ONSET IN HIGH-RISK YOUTHS 6 (2015).

For example, one vignette in the survey could describe a search engine collecting location data and using the data to place ads on the site. The next vignette could describe a browser collecting search terms to improve the functionality of the site. For any given vignette, the respondent judged whether the described data was "OK." See Figure 1 for the different factors and values as well as an example vignette.

FIGURE 1: Vignette Template and Example

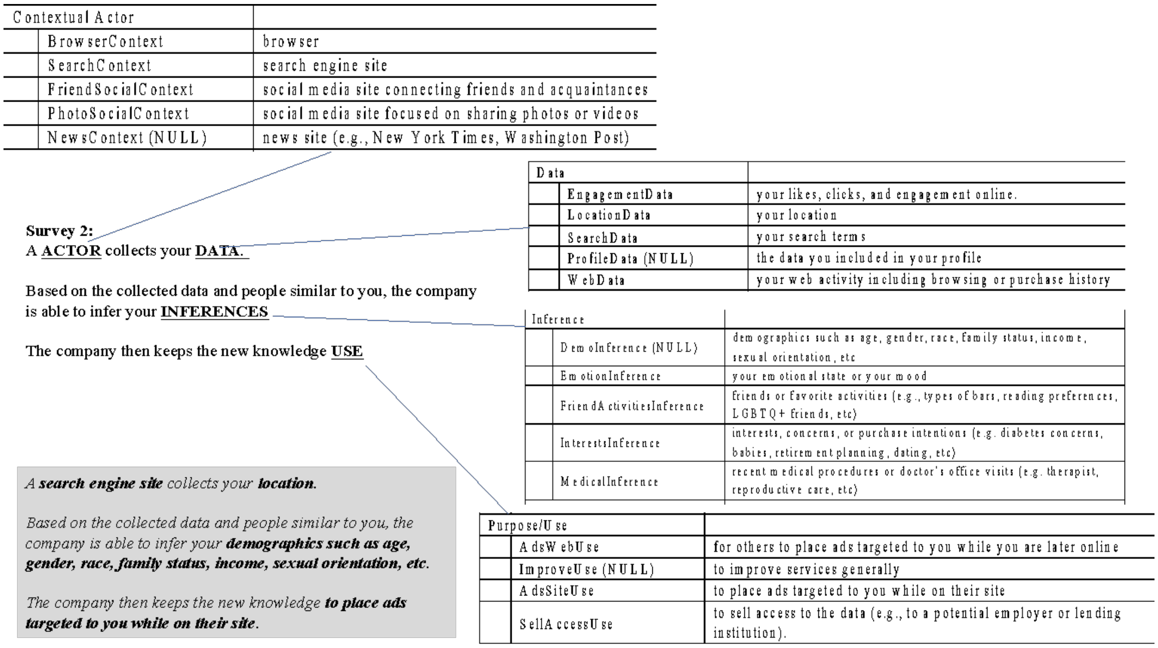

Contextual Actor		
	BrowserContext	browser
	SearchContext	search engine site
	FriendSocialContext	social media site connecting friends and acquaintances
	PhotoSocialContext	social media site focused on sharing photos or videos
	NewsContext (NULL)	news site (e.g., New York Times, Washington Post)

Data		
	EngagementData	your likes, clicks, and engagement online.
	LocationData	your location
	SearchData	your search terms
	ProfileData (NULL)	the data you included in your profile
	WebData	your web activity including browsing or purchase history

Survey 2:
A **ACTOR** collects your **DATA**.

Based on the collected data and people similar to you, the company is able to infer your **INFERENCES**

The company then keeps the new knowledge **USE**

Inference		
	DemoInference (NULL)	demographics such as age, gender, race, family status, income, sexual orientation, etc
	EmotionInference	your emotional state or your mood
	FriendActivitiesInference	friends or favorite activities (e.g., types of bars, reading preferences, LGBTQ+ friends, etc)
	InterestsInference	interests, concerns, or purchase intentions (e.g. diabetes concerns, babies, retirement planning, dating, etc)
	MedicalInference	recent medical procedures or doctor's office visits (e.g. therapist, reproductive care, etc)

*A **search engine site** collects your **location**.*

*Based on the collected data and people similar to you, the company is able to infer your **demographics such as age, gender, race, family status, income, sexual orientation, etc**.*

*The company then keeps the new knowledge **to place ads targeted to you while on their site**.*

Purpose/Use		
	AdsWebUse	for others to place ads targeted to you while you are later online
	ImproveUse (NULL)	to improve services generally
	AdsSiteUse	to place ads targeted to you while on their site
	SellAccessUse	to sell access to the data (e.g., to a potential employer or lending institution).

1. Data Collection

We ran six studies. In each, the general template was similar and the factors/values were consistent. In other words, all six survey conditions included the data types of user engagement, location, search terms, profile data, and Web history. However, we did vary whether or not inferences were used (Surveys 2, 3, and 4) or not (Surveys 1 and 6). In addition, to capture realistic data collection and use, we included both a primary actor (browser, search engine, etc.) and a third-party (data broker, etc.) as the organization that created and/or used the inference. Figure 2 illustrates how the inclusion of inferences and the involvement of a third-party were systematically varied for Surveys 1 through 6.

Our analysis is in two parts. First, we examine whether the use of inferences, rather than raw data, better meets the privacy expectations of consumers. When a browser or search engine argues that it uses inferences, but not raw data, for ads, is this a privacy solution? Second, we examine whether precluding third parties from collecting, accessing, or using consumer data is a privacy solution. Were a search engine or social network, for example, to preclude third parties from accessing user data, but creates, uses, and sells access to inferences drawn from it, have they honored privacy?

FIGURE 2: Survey Design for All Six Survey Runs

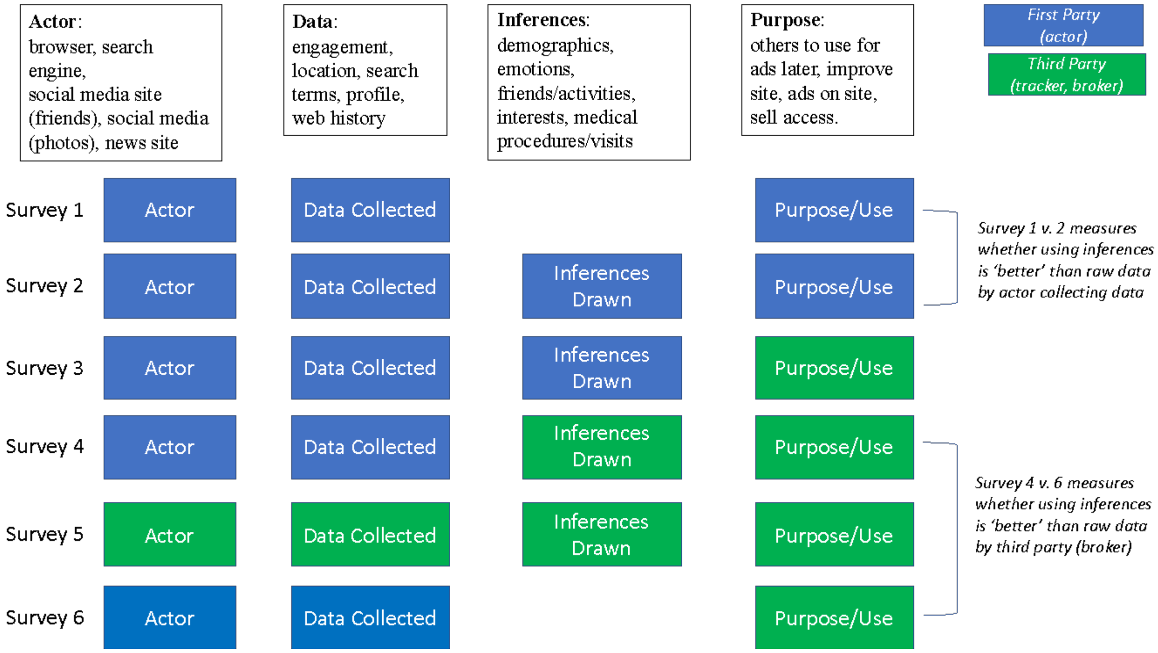

Actor: browser, search engine, social media site (friends), social media (photos), news site

Data: engagement, location, search terms, profile, web history

Inferences: demographics, emotions, friends/activities, interests, medical procedures/visits

Purpose: others to use for ads later, improve site, ads on site, sell access.

🟦 First Party (actor)

🟩 Third Party (tracker, broker)

Survey 1	🟦 Actor	🟦 Data Collected		🟦 Purpose/Use
Survey 2	🟦 Actor	🟦 Data Collected	🟦 Inferences Drawn	🟦 Purpose/Use
Survey 3	🟦 Actor	🟦 Data Collected	🟦 Inferences Drawn	🟩 Purpose/Use
Survey 4	🟦 Actor	🟦 Data Collected	🟩 Inferences Drawn	🟩 Purpose/Use
Survey 5	🟩 Actor	🟩 Data Collected	🟩 Inferences Drawn	🟩 Purpose/Use
Survey 6	🟦 Actor	🟦 Data Collected		🟩 Purpose/Use

Survey 1 v. 2 measures whether using inferences is 'better' than raw data by actor collecting data

Survey 4 v. 6 measures whether using inferences is 'better' than raw data by third party (broker)

2. Sample

The surveys were run in March and April of 2023 using Prolific, which is a platform for researchers to recruit survey respondents.[97] For each condition, respondents each rated thirty vignettes with a single rating task—the degree to which the described scenario was OK. Table 1 includes the summary statistics for the six survey conditions used in this paper. Respondents were paid 1.60 GBP for an hourly rate of 11.50 GBP, which was above average for the site. Each condition had unique respondents—i.e., respondents for Survey 1 were precluded from Surveys 2– 6.

[97] *See What Is Prolific and How Does It Work?*, PROLIFIC (Aug. 20, 2024), https://participant-help.prolific.com/en/article/dc132c [https://perma.cc/X78N-3HBN].

TABLE 1: Summary Statistics Surveys 1– 6

	Survey 1		Survey 2		Survey 3		Survey 4		Survey 5		Survey 6	
# of Respondents	488		586		592		573		581		588	
# of Vignettes	14640		17580		17760		17190		17430		17640	
Average Rating	-23.41		-30.88		-29.92		-34.67		-32.74		-26.9	
S.D. Rating	66.23		64.64		62.8		61.8		64.18		64.28	
	Count	**%**	**Count**	**%**	**Count**	**%**	**Count**	**%**	**Count**	**%**	**Count**	**%**
Age												
18-24	77	15.7%	78	13.3%	61	10.3%	86	15.1%	121	20.8%	79	13.4%
25-34	189	38.7%	203	34.6%	224	37.8%	228	39.7%	228	39.2%	196	33.3%
35-44	134	27.4%	159	27.1%	147	24.8%	134	23.3%	119	20.4%	139	23.6%
45-54	45	9.2%	61	10.4%	86	14.5%	53	9.2%	62	10.6%	92	15.6%
55-64	33	6.7%	54	9.2%	51	8.6%	50	8.7%	38	6.5%	58	9.8%
65+	9	1.8%	31	5.2%	23	3.8%	22	3.8%	12	2.0%	24	4.0%
total	487	99.8%	586	100.0%	592	100.0%	573	100.0%	580	99.8%	588	100.0%
Gender												
Male	260	53.2%	314	53.5%	326	55.0%	282	49.2%	299	51.4%	310	52.7%
Female	216	44.2%	261	44.5%	256	43.2%	276	48.1%	269	46.3%	269	45.7%
Non-binary / third gender	7	1.4%	10	1.7%	6	1.0%	11	1.9%	13	2.2%	3	0.5%
Prefer not to say	4	0.8%			3	0.5%	4	0.7%			6	1.0%
Prefer to self-describe	1	0.2%	1	0.1%	1	0.12%						
total	488	100.0%	586	100.0%	592	100.00%	573	100.0%	581	100.0%	588	100.0%

III. RESULTS 1: INFERENCES VERSUS RAW DATA AS PRIVACY SOLUTION

In order to measure whether creating inferences rather than raw data better met consumers privacy expectations, we compared the results of survey conditions with and without the use of inferences:

1. We tested whether first parties using raw data was a greater privacy violation than when they use inferences based on that raw data.
2. We tested whether third parties (data brokers) buying and using raw data was a greater privacy violation than (1) buying and using inferences or (2) creating and using inferences based on that raw data.

Table 2 includes the results of regressing the rating task—the degree to which a given vignette is rated "OK"—on the vignette factors. In general, the average rating task for all conditions is negative, meaning that individuals do not find the vignettes describing the collection and use of their data to be OK on average. In addition, the results show that the *purpose for which knowledge is used* is generally more important than other types of factors, such as the type of data, the actor that collects the data, and inferences drawn. The coefficients for use/purpose are greater than the coefficients for the type of data: changing how knowledge is used had a greater impact than changing the type of information collected, all else being equal.

TABLE 2: Regression results for Surveys 1– 6.

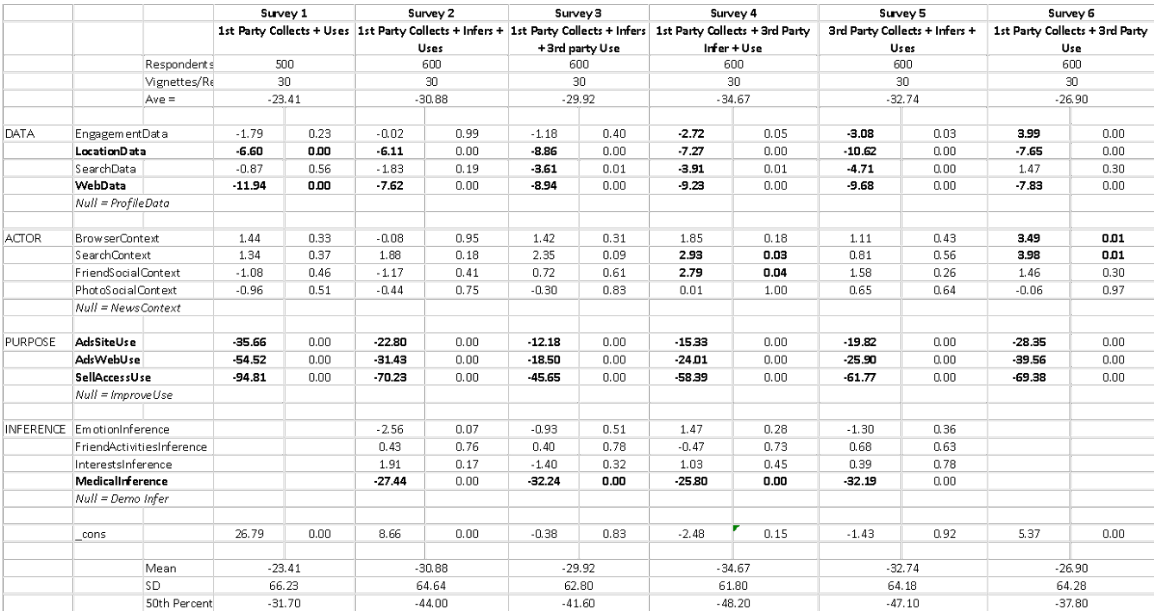

			Survey 1		**Survey 2**		**Survey 3**		**Survey 4**		**Survey 5**		**Survey 6**	
			1st Party Collects + Uses		**1st Party Collects + Infers + Uses**		**1st Party Collects + Infers + 3rd party Use**		**1st Party Collects + 3rd Party Infer + Use**		**3rd Party Collects + Infers + Uses**		**1st Party Collects + 3rd Party Use**	
		Respondents	500		600		600		600		600		600	
		Vignettes/Re	30		30		30		30		30		30	
		Ave =	-23.41		-30.88		-29.92		-34.67		-32.74		-26.90	
DATA	EngagementData		-1.79	0.23	-0.02	0.99	-1.18	0.40	**-2.72**	0.05	**-3.08**	0.03	**3.99**	0.00
	LocationData		**-6.60**	**0.00**	**-6.11**	0.00	**-8.86**	0.00	**-7.27**	0.00	**-10.62**	0.00	**-7.65**	0.00
	SearchData		-0.87	0.56	-1.83	0.19	**-3.61**	0.01	**-3.91**	0.01	**-4.71**	0.00	1.47	0.30
	WebData		**-11.94**	**0.00**	**-7.62**	0.00	**-8.94**	0.00	**-9.23**	0.00	**-9.68**	0.00	**-7.83**	0.00
	Null = ProfileData													
ACTOR	BrowserContext		1.44	0.33	-0.08	0.95	1.42	0.31	1.85	0.18	1.11	0.43	**3.49**	**0.01**
	SearchContext		1.34	0.37	1.88	0.18	2.35	0.09	**2.93**	**0.03**	0.81	0.56	**3.98**	**0.01**
	FriendSocialContext		-1.08	0.46	-1.17	0.41	0.72	0.61	**2.79**	**0.04**	1.58	0.26	1.46	0.30
	PhotoSocialContext		-0.96	0.51	-0.44	0.75	-0.30	0.83	0.01	1.00	0.65	0.64	-0.06	0.97
	Null = NewsContext													
PURPOSE	**AdsSiteUse**		**-35.66**	0.00	**-22.80**	0.00	**-12.18**	0.00	**-15.33**	0.00	**-19.82**	0.00	**-28.35**	0.00
	AdsWebUse		**-54.52**	0.00	**-31.43**	0.00	**-18.50**	0.00	**-24.01**	0.00	**-25.90**	0.00	**-39.56**	0.00
	SellAccessUse		**-94.81**	0.00	**-70.23**	0.00	**-45.65**	0.00	**-58.39**	0.00	**-61.77**	0.00	**-69.38**	0.00
	Null = ImproveUse													
INFERENCE	EmotionInference				-2.56	0.07	-0.93	0.51	1.47	0.28	-1.30	0.36		
	FriendActivitiesInference				0.43	0.76	0.40	0.78	-0.47	0.73	0.68	0.63		
	InterestsInference				1.91	0.17	-1.40	0.32	1.03	0.45	0.39	0.78		
	MedicalInference				**-27.44**	0.00	**-32.24**	**0.00**	**-25.80**	**0.00**	**-32.19**	0.00		
	Null = Demo Infer													
	_cons		26.79	0.00	8.66	0.00	-0.38	0.83	-2.48	0.15	-1.43	0.92	5.37	0.00
		Mean	-23.41		-30.88		-29.92		-34.67		-32.74		-26.90	
		SD	66.23		64.64		62.80		61.80		64.18		64.28	
		50th Percent	-31.70		-44.00		-41.60		-48.20		-47.10		-37.80	

A. Use of Inferences Versus Raw Data

We first tested whether the purpose or use of inferences was judged more positively compared to using raw data. Survey 1 includes scenarios where an actor (e.g., search engine) collects data (e.g., location) and then uses that data (e.g., place ads). Survey 2 is the same as Survey 1, but the data recipient first derives inferences (e.g., emotions) about the user and then uses those inferences for the same purposes (e.g., to place ads). Figure 3 shows the comparison of Survey 1 versus Survey 2.

FIGURE 3: Designs for Surveys 1 and 2 to Test Use of Inferences

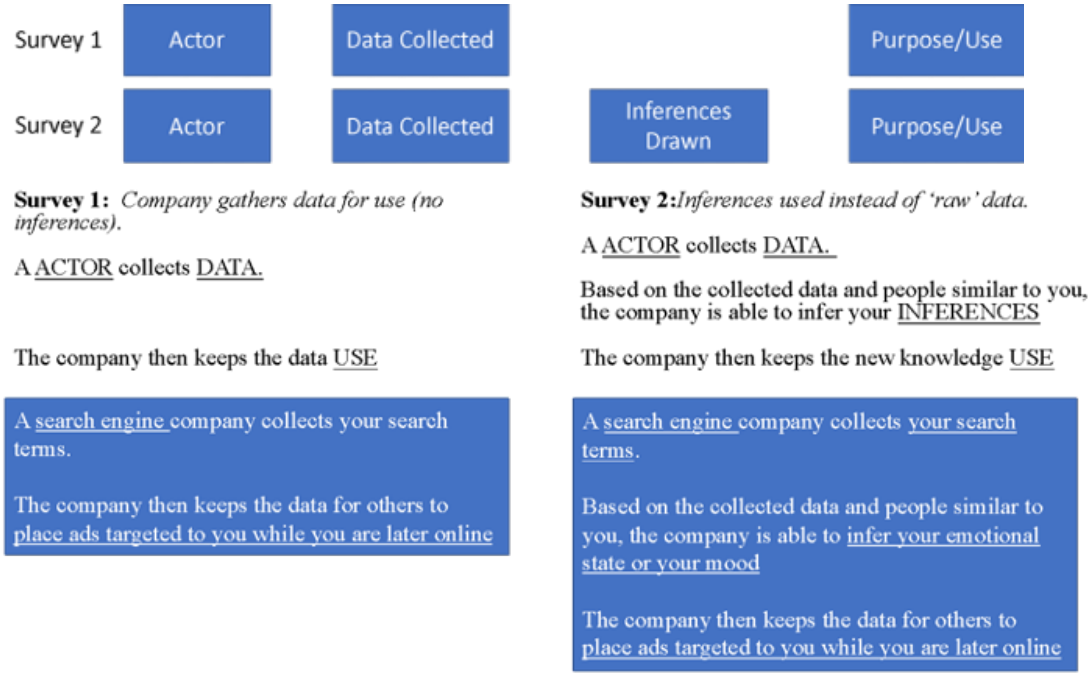

Survey 1: *Company gathers data for use (no inferences).*

A ACTOR collects DATA.

The company then keeps the data USE

A search engine company collects your search terms.

The company then keeps the data for others to place ads targeted to you while you are later online

Survey 2:*Inferences used instead of 'raw' data.*

A ACTOR collects DATA.

Based on the collected data and people similar to you, the company is able to infer your INFERENCES

The company then keeps the new knowledge USE

A search engine company collects your search terms.

Based on the collected data and people similar to you, the company is able to infer your emotional state or your mood

The company then keeps the data for others to place ads targeted to you while you are later online

The results in Table 2 show that, on average, the use of inferences is judged as less appropriate than the use of raw data. The average rating task for Survey 2 (-30.88) is significantly lower than the average rating for Survey 1 (-23.41, t = 10.22, p < 0.00). The results are contrary to the argument that using inferences rather than raw data better meets users' privacy expectations.

We also compared the use of inferences versus raw data for specific data types in Figures A1-5 in the Appendix. For example, Figure 4 below shows the average rating (degree scenario is OK) for different uses of Web history (found in Survey 1) compared to inferences *drawn from Web* history data found in Survey 2. In other words, the averages plotted in Figure 4 "to improve services" are calculated from vignettes where Web history data is included and used to improve services (Survey 1) as well as where inferences based on Web history data are used to improve services (Survey 2). A box around Web history data and an inference designates statistically identical averages (based on t-tests and p<0.05). The results show that users do not judge inferences drawn from Web history data as more appropriate than Web history data itself, for any use. The same result holds for each data type in Figures A1-5.

FIGURE 4: Use of Web History Data Compared to use of Inferences Based on Web History Data

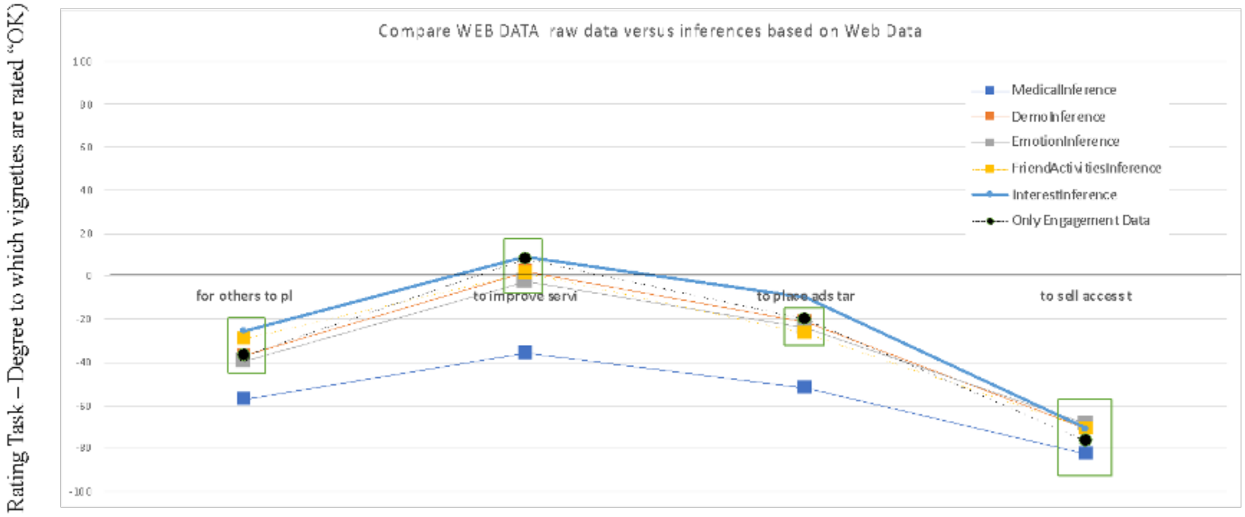

The use of inferences is not judged as more appropriate than the use of raw data as shown in Figures A1-5 in the Appendix, and in the majority of cases, both the use of data and inferences are judged to be privacy violations. In fact, the use of location data is rated the same as the use of emotion, demographic, and friend inferences for placing ads (Figure A1), and all are judged to be privacy violations. The use of raw data is rated *more* appropriate than the use of inferences when used for improving the site for engagement and search data (Figures A2 and A3). And the use of raw data is rated *better* than the use of medical inferences for all purposes aside from selling access to data/inferences (e.g., to a potential employer or lender), which has a negative rating (-60 to -80) for any type of data or knowledge.

In fact, respondents' judgements were sensitive to the purposes for which the information was used. Specifically, when information flows served contextual ends and values (e.g., to improve a given service) the data flows were found to be appropriate consistently across contextual actors (e.g. news site, search engine, and social network). In the regression results in Table 2, ratings for using information to place ads on a website, offering the data for targeted advertising later online, and selling the data to a data broker are all statistically different from ratings for using the data for the purpose of improving services across all six surveys. Additional analysis in the Appendix in Tables A1-5 shows that the same holds for each type of data/inference: for each type of data collected and type of inference created, respondents differentiated between the four different uses of data. In addition, the use of the same information to promote non-contextual purposes, such as to place ads when subsequently online or to sell to others, did not meet privacy expectations and was rated more negatively. See Tables A1-5 in the Appendix. The Appendix includes the comparison of the respondents' privacy judgements about raw data versus inferences for search engine, browser, and social network. In each case, the use of raw data meets greater approval than inferences when the purpose is to improve services. Using data for the purpose of subsequently placing ads is rated poorly across all three actors, no matter whether the data is raw or inferred. See Tables A6-8.

These findings are consistent with previous empirical work on uses of data for contextual versus non contextual purposes, where respondents evaluated the use for non-contextual purposes to be inappropriate.[98] (Martin and Nissenbaum). This finding undermines companies' claims that the use of inferences rather than raw data is a privacy solution. Our results also undermine companies' claims that whether or not they share data is more important than the use of that data for a particular purpose.

We then examined five specific cases common in practice:

- The collection of search data by search engines
- The collection of location data by search engines
- The collection of engagement data by social networks
- The collection of Web history data by browsers
- The collection of location data by news sites

Figures A9-13 include the average rating tasks for the specific cases. For example, the results in Figure 5 (Figure A9a in the Appendix) show that the use of (raw) search terms for the purpose of improving the search engine is statistically preferable to using inferences based on searches, such as medical visits, interests, emotional state, and friends. Similarly, the use of (raw) search data to place ads on the search result site is *better* than using inferences about demographics or medical visits. Figures A10-A13 show similar results with the use of raw data being preferred or statistically equivalent to using inferences for each case, respectively.

[98] Martin & Nissenbaum, *supra* note 20, at 177, 207, 214.

FIGURE 5a: Rating Average for Use of Search Data versus Inferences by Search Engine

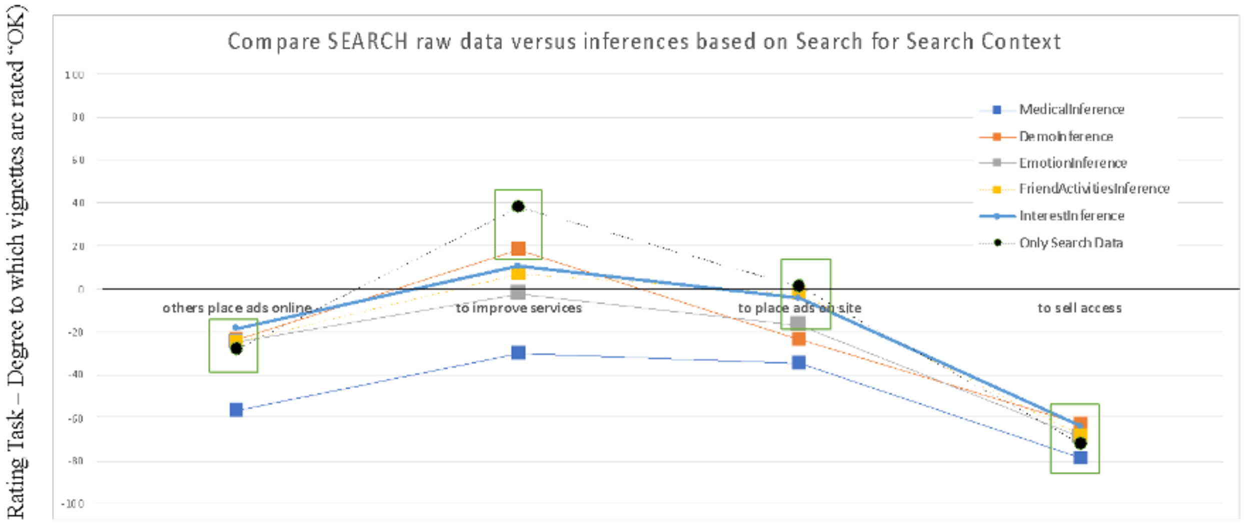

FIGURE 5b: Rating Average for Use of Engagement Data versus Inferences by Social Network

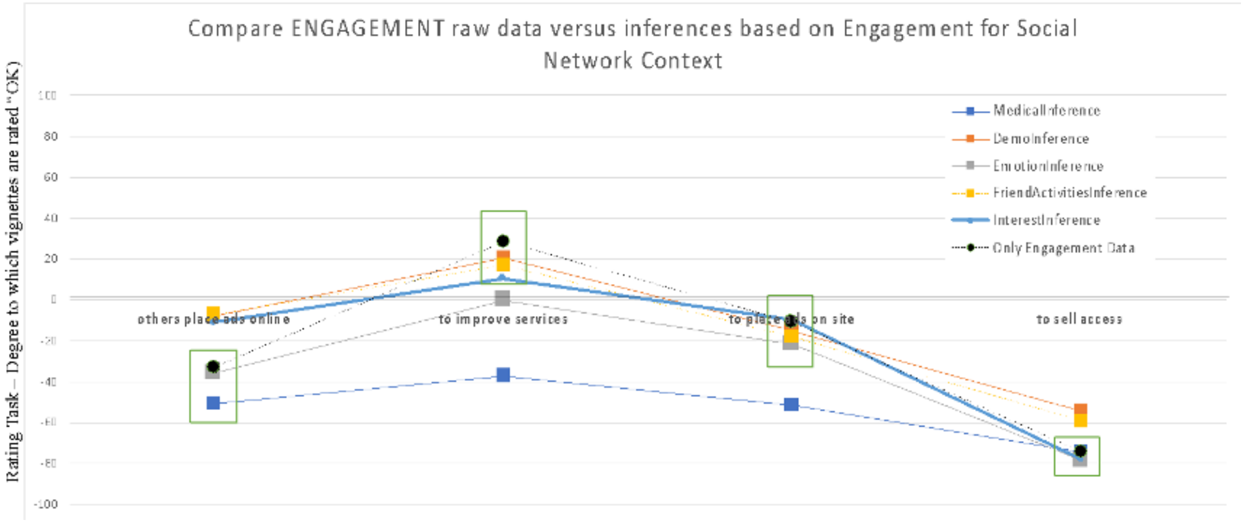

B. *Using Inferences versus Raw Data - for Data Brokers*

To ascertain whether the collection and use of inferences is judged as more appropriate than raw data when the recipient is a third-party data broker, we test two scenarios: (1) whether a data broker *buying* inferences rather than raw data is judged as more appropriate and (2) whether a data broker being able to *create* and use inferences rather than use raw data is judged as more appropriate.

1. Data Broker Buying Inferences v. Raw Data

In order to test if a data broker buying inferences about users is judged as more appropriate than a data broker buying raw data about users, we compared the results of Surveys 3 and 6. Survey 3 describes scenarios where an actor (e.g., search engine) collects data (e.g., location),creates inferences based on the data collected (e.g., medical visits), and then sells those inferences to a data broker to use (e.g., place ads). Survey 6 includes the same actors collecting data but with the actor selling the *raw data* to a data broker to use. The only difference between the two scenarios is whether the data broker purchases raw data (Survey 6), or purchases inferences based on that raw data (Survey 3).

FIGURE 6: Design of Surveys 3 and 6

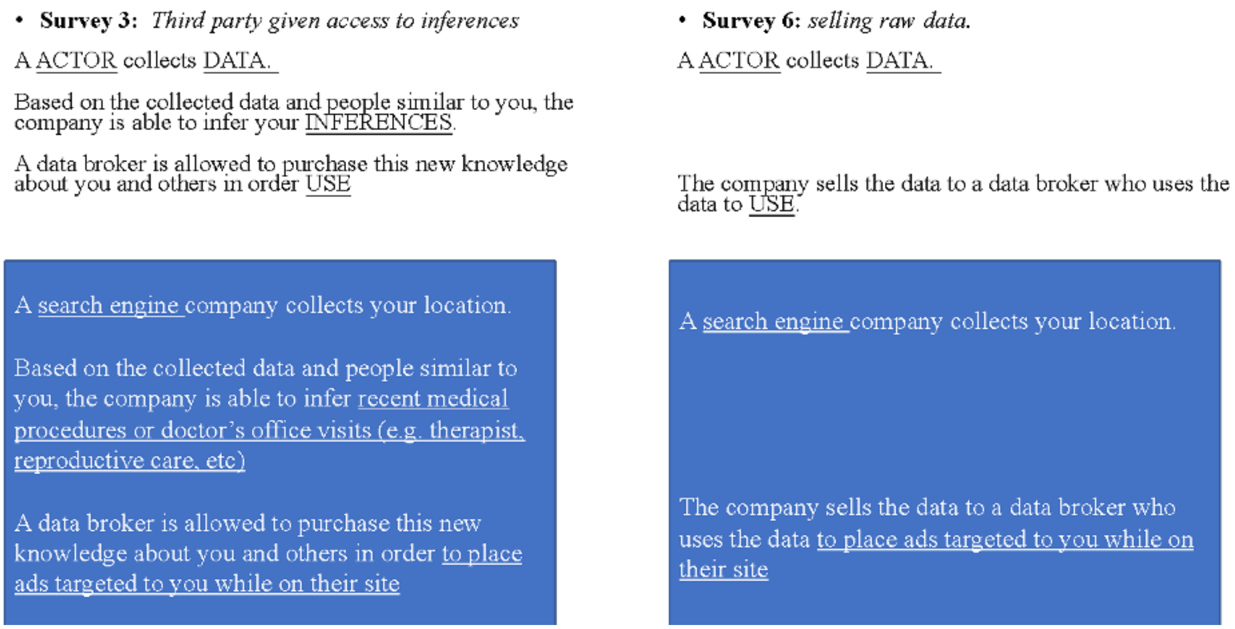

- **Survey 3:** *Third party given access to inferences*

A ACTOR collects DATA.

Based on the collected data and people similar to you, the company is able to infer your INFERENCES.

A data broker is allowed to purchase this new knowledge about you and others in order USE

A search engine company collects your location.

Based on the collected data and people similar to you, the company is able to infer recent medical procedures or doctor's office visits (e.g. therapist, reproductive care, etc)

A data broker is allowed to purchase this new knowledge about you and others in order to place ads targeted to you while on their site

- **Survey 6:** *selling raw data.*

A ACTOR collects DATA.

The company sells the data to a data broker who uses the data to USE.

A search engine company collects your location.

The company sells the data to a data broker who uses the data to place ads targeted to you while on their site

We find that the purchase and use of inferences by data brokers in Survey 3 (average rating = -29.92) does not meet the privacy expectations of consumers and is judged to be less appropriate than buying and using raw data in Survey 6 (-26.90; t = -4.4683, p<0.00). The purchase and use of both raw data and inferences do not meet the privacy expectations of users as both Surveys have average ratings that are negative.

We then examined whether raw data (Survey 6) or inferences (Survey 3) are judged as more appropriate to purchase by a data broker for each data type. Figures A14-18 in the Appendix compare the privacy judgements of respondents for the purchase of specific raw data versus inferences based on that same raw data. Similar to the analysis above, data brokers buying raw data to improve the site better meets privacy expectations compared to the purchase of inferences for each type of data (search, location, Web history, engagement, or profile data). In addition, respondents do not differentiate between the use of raw data and inferences when data brokers target users with ads later online (and both are judged a privacy violation). For example, Figure A14 and Table A14 show that the average privacy judgments for the collection of location data by data brokers are privacy violations, regardless of whether raw data or inferences are used to place ads on a site, target users later online, or sell the data to other interested parties.

2. Data Broker Creating Inferences Versus Buying Data

We then examined whether allowing a data broker to *create* inferences without buying the raw data is better than that same data broker buying and using raw data. We compared the results of Surveys 4 and 6. Survey 4 describes scenarios where an actor (e.g., search engine) collects data (e.g., location) and allows a data broker to create inferences based on the data collected (e.g., medical visits), which the data broker then uses (e.g., by placing ads). Survey 6 includes the same actors collecting data but then the actor sells the raw data to a data broker to use. The only difference between the two scenarios is whether the data broker purchases raw data (Survey 6) or purchases the ability to create inferences based on that raw data (Survey 4).

FIGURE 7: Design of Surveys 4 and 6

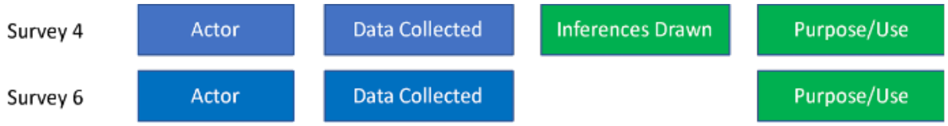

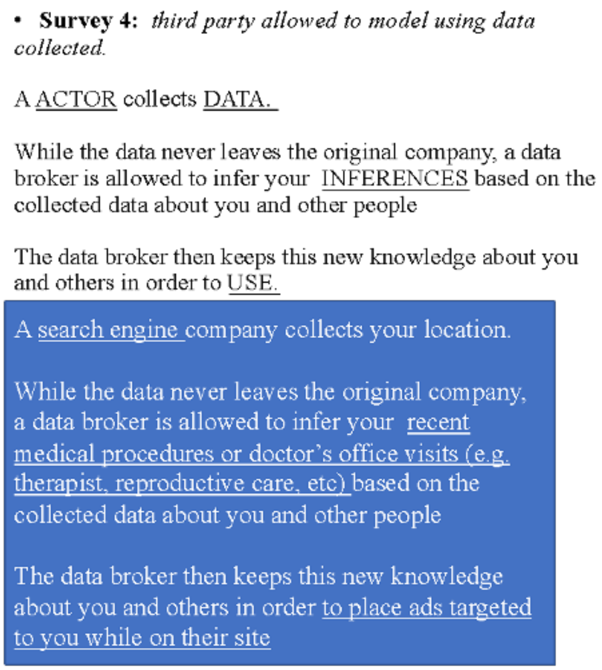

- **Survey 4:** *third party allowed to model using data collected.*

A <u>ACTOR</u> collects <u>DATA.</u>

While the data never leaves the original company, a data broker is allowed to infer your <u>INFERENCES</u> based on the collected data about you and other people

The data broker then keeps this new knowledge about you and others in order to <u>USE.</u>

A <u>search engine</u> company collects your location.

While the data never leaves the original company, a data broker is allowed to infer your <u>recent medical procedures or doctor's office visits (e.g. therapist, reproductive care, etc)</u> based on the collected data about you and other people

The data broker then keeps this new knowledge about you and others in order <u>to place ads targeted to you while on their site</u>

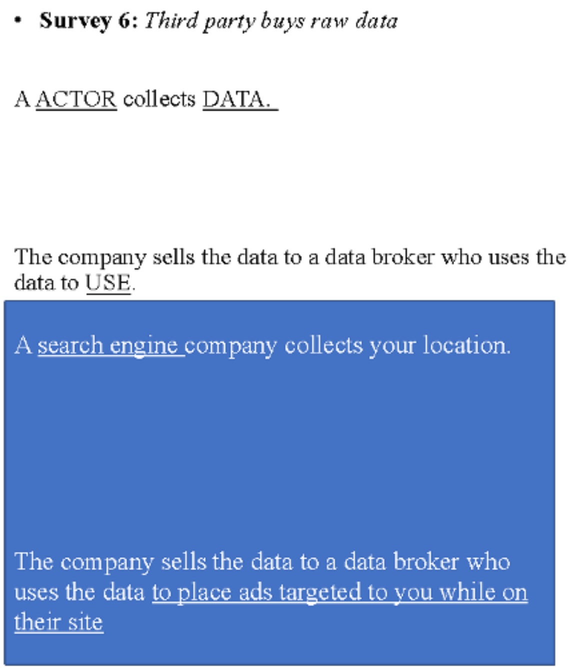

- **Survey 6:** *Third party buys raw data*

A <u>ACTOR</u> collects <u>DATA.</u>

The company sells the data to a data broker who uses the data to <u>USE</u>.

A <u>search engine</u> company collects your location.

The company sells the data to a data broker who uses the data <u>to place ads targeted to you while on their site</u>

We find that the creation and use of inferences by data brokers (-34.67 for Survey 4) is judged to be a privacy violation and is judged less appropriate than buying and using raw data in Survey 6 (-26.90; t = -11.4956, p<0.00). It should be noted that allowing a data broker to create inferences by accessing the data collected (while never leaving the original company, as shown in Survey 4) is considered *less appropriate* than the first party that collected the data selling the inferences to the same data broker (Survey 3).

We also tested whether allowing a data broker to create inferences without taking the actual data (Survey 4) was better at meeting privacy expectations than being sold the raw data and creating the same inferences (Survey 5). We find that allowing a data broker to draw inferences without collecting the data was considered *worse* in meeting privacy expectations of respondents (Survey 4 average = -34.67) as compared to allowing the data broker to just collect the raw data through trackers and then draw inferences (Survey 5 average = -32.74; t=-2.8461, p = 0.002).

In sum, the use of inferences rather than raw data collected by a primary site is not a privacy solution for users. In most instances, respondents judged the use of raw data such as browsing history, location, search terms, and engagement data to be statistically the same as using inferences based on that same data. Further, for improving services across contexts, consumers judged the use of raw data as more appropriate compared to using inferences based on that same raw data. The purpose of using the data or inferences was statistically significant for respondents – respondents differentiated between using data to improve services, place ads on a website, provide to companies to place ads later when online, and sell to a data broker. In fact, where users did not differentiate between data and inferences, the respondents did differentiate between different purposes of using the knowledge in making privacy judgements. These findings undermine the presumption that the users' privacy expectations are addressed when a company uses inferences about individuals rather than raw data.

IV. RESULTS 2: 1ST PARTY VERSUS 3RD PARTY AS A PRIVACY SOLUTION

The second section of analysis focuses on comparing a first- party versus a third-party accessing user information as a privacy solution. To test the relative importance of first parties versus third parties collecting data, creating inferences, or using information, we created a series of surveys that systematically changed the actor (first or third party) that collects, creates inferences, or uses the information. Figure 2 above includes a diagram where the factors and values remain the same (actor, data, inferences, purpose) but the vignette template changes *who* takes each action.

For example, Survey 2 includes scenarios where the first party (browser, search engine, etc.) collects the data, creates inferences, and uses that new knowledge. Survey 5 includes scenarios where third parties (trackers and data brokers) collect the data, create inferences, and use that new knowledge. The results are in Table 2 above with the summary statistics for each survey run. The regression results are in Table 3 where the dependent variable was regressed on the vignette factors for each condition. In the analysis below, we first measure whether having a first party versus a third-party *using* inferences is judged as more appropriate. We then move to comparing first and third parties for the *creation of inferences* as well as the *collection* of data.

A. Comparing First Versus Third Party in the Use of Information

One privacy solution offered by companies is to have the first party – a social network, browser, or search engine – store and use data but preclude any third parties from accessing the same knowledge. For example, a search company may collect consumer data, draw inferences based on the data, and, while using the inferences themselves, argue that they protect privacy by withholding these inferences from third parties.. A social network may gather engagement data, create inferences about their users, and argue it protects privacy by not allowing third parties to have access to those same inferences.

In order to compare whether allowing first- party versus third- party to *use* inferences matters to privacy judgments, we compared Surveys 2 and 3 as in Figure 9. The only difference between the two surveys is whether the first party (browser, search engine, social network, etc.) versus third-party (data broker) has the inference to place ads, sell the inference, or improve the service.

Figure 8: Analysis Design for Comparing First Versus Third-party Use of Inferences

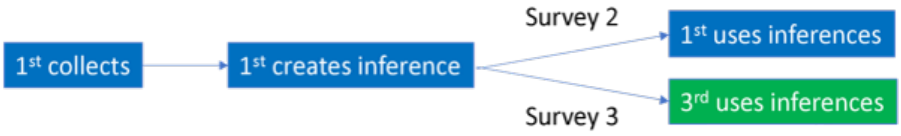

Figure 9: Vignette Template and Example Comparing Surveys 2 and 3

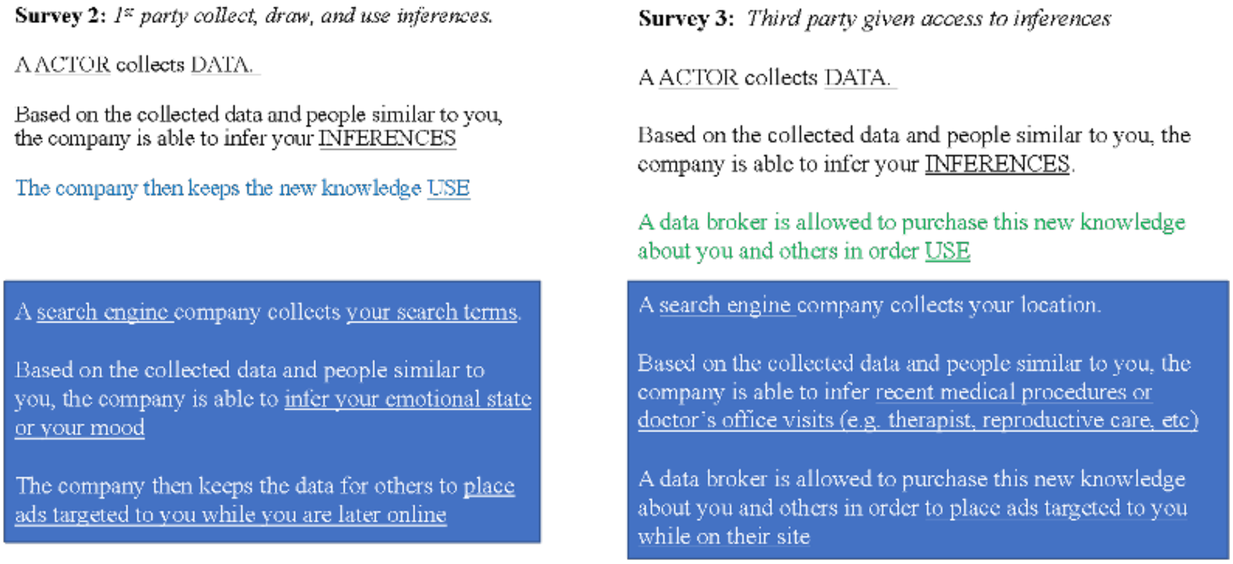
Survey 2: *1st party collect, draw, and use inferences.*

A ACTOR collects DATA.

Based on the collected data and people similar to you, the company is able to infer your INFERENCES

The company then keeps the new knowledge USE

A search engine company collects your search terms.

Based on the collected data and people similar to you, the company is able to infer your emotional state or your mood

The company then keeps the data for others to place ads targeted to you while you are later online

Survey 3: *Third party given access to inferences*

A ACTOR collects DATA.

Based on the collected data and people similar to you, the company is able to infer your INFERENCES.

A data broker is allowed to purchase this new knowledge about you and others in order USE

A search engine company collects your location.

Based on the collected data and people similar to you, the company is able to infer recent medical procedures or doctor's office visits (e.g. therapist, reproductive care, etc)

A data broker is allowed to purchase this new knowledge about you and others in order to place ads targeted to you while on their site

The results show that respondents judged a first party using inferences met their privacy preferences to the same degree as a data broker who purchases and uses the inferences for the same purpose. We see this both overall with the average vignette rating for the first party using inferences in Survey 2 (-30-88) as statistically the same as the third-party using inferences in Survey 3 in Table 2 (-29.92; t = -1.4231, p = 0.0774). Figure 10 shows the same results for specific inferences and uses of knowledge for the search context. The other contextual actors are included in the Appendix in Figures A21-25. In general, having a first party use the inference *to improve services* is judged as more appropriate than a third party's use. However, when a third party later sells those inferences, it is judged as better meeting privacy expectations than the first party doing so (Figure 10). The same trend holds across the contextual actors in the Appendix.

Figure 10: Average Rating for the Search Engine Versus Third Party Use of Inferences

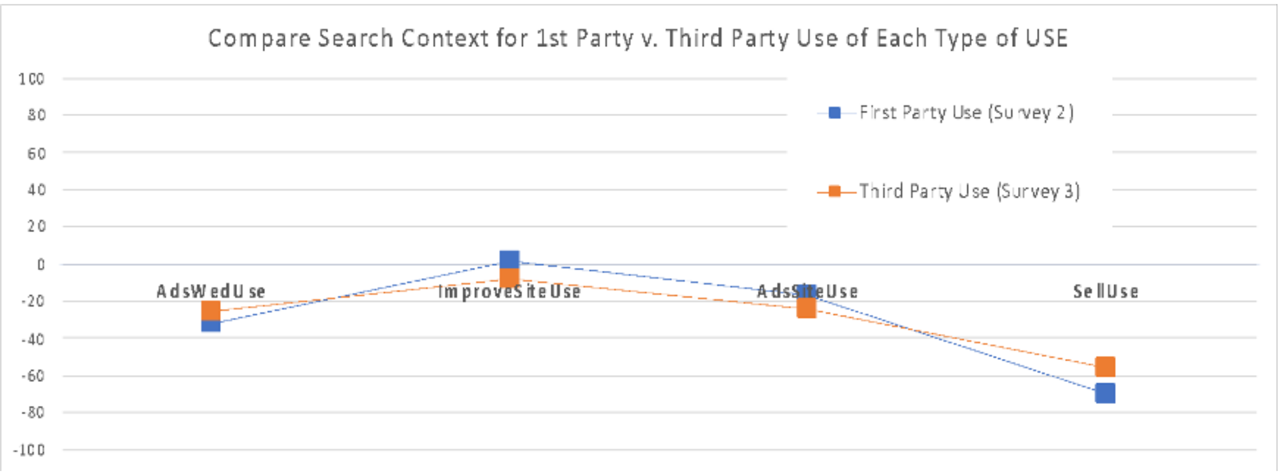

Figure 11: Average Rating for a Browser Versus Third Party Use of Inferences

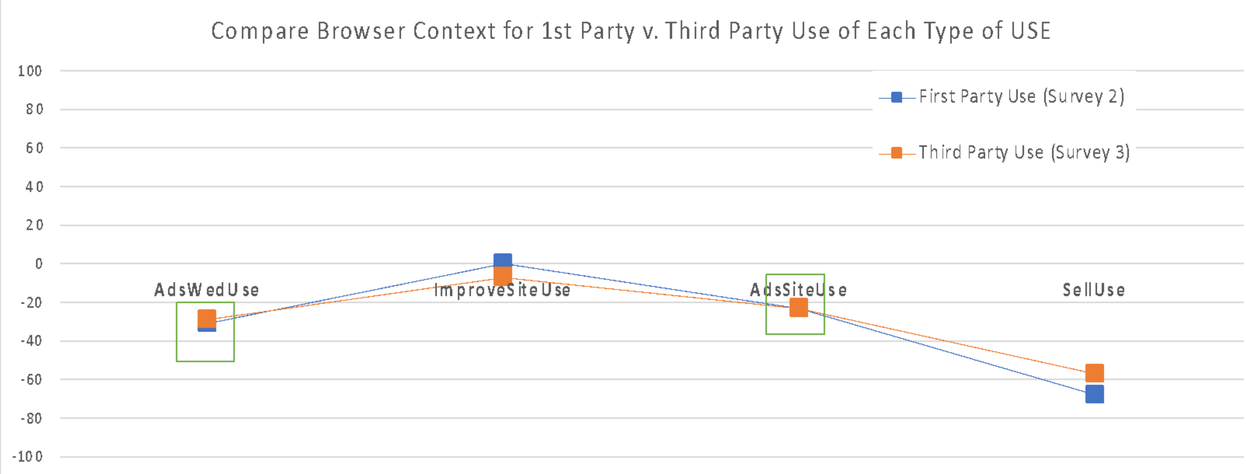

We then examined five specific cases as before. The results are in the Appendix in Figures A26-30. For each, the only use of information that is, on average, positive by either a first or third party is to use inferences to improve services on the site. For example, in considering the capture and use of Web history data within the browser context, the use of inferences by *third parties* (trackers collecting data and a data broker creating inferences for particular purposes in Survey 3) is judged to better meet the privacy expectations of users compared to the first party (the browser) for selling the data to others as shown in Figure 11 ;browsers and data brokers are judged statistically the same for placing ads on a site and for targeting users with ads later online (both are privacy violations).

Figure 12: Comparison of Search Engine Versus Third Party Using Inferences Based on Search Data

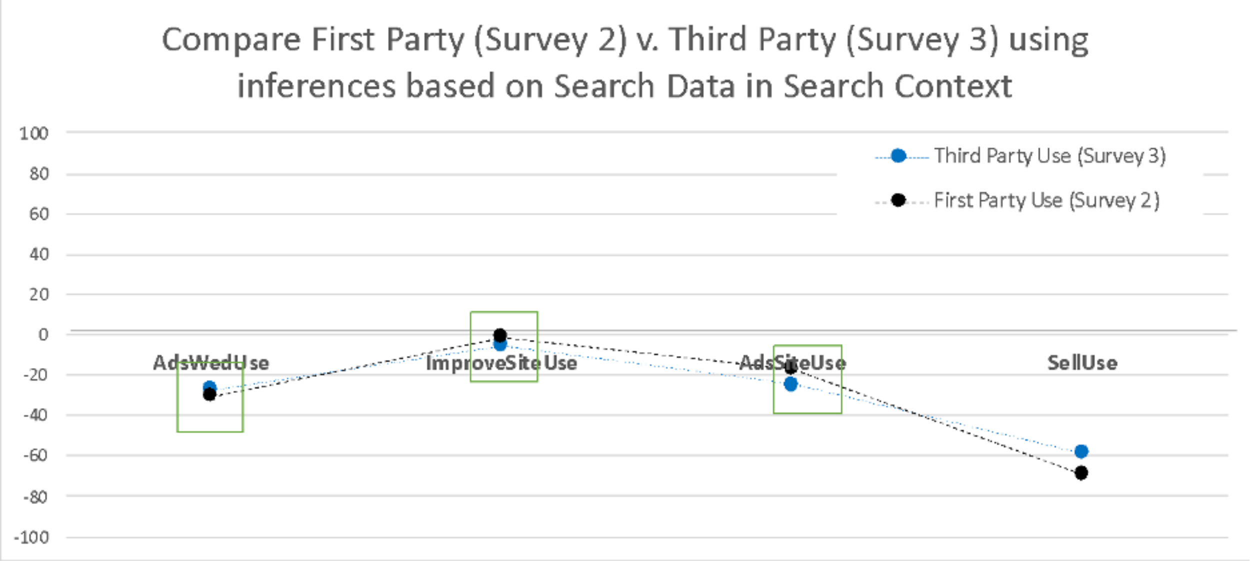

Figure 12 compares the average rating that the scenario is okay (judged as appropriate) for first parties versus third parties using inferences based on search data within the search context. The results show that *selling* the data is judged to be more acceptable for third parties compared to the search engine doing the same activity. The same trend holds for the other special cases in the Appendix.

We extended this analysis to examine if the same trends hold for the collection and use of raw data by first versus third parties through the comparison of Surveys 1 and 6. The results show that the average rating for the appropriateness of the data flow was greater for first parties in Survey 1 (-23.11) compared to data brokers in Survey 6 (-26.90; t = 4.7930, p<0.00). In other words, respondents found both the use of raw data by third parties (data brokers) and first parties (browsers, search engine, social network, etc.) to be a privacy violation since both averages were well below zero. However, respondents found the use of raw data, on average, to be worse when used by the third-party (data broker) to sell to other companies as shown in Figure 13. The Appendix has a detailed analysis in Figures A31-35.

Figure 13: Comparison of Browser Versus Third Party Using Inferences Based on Web History

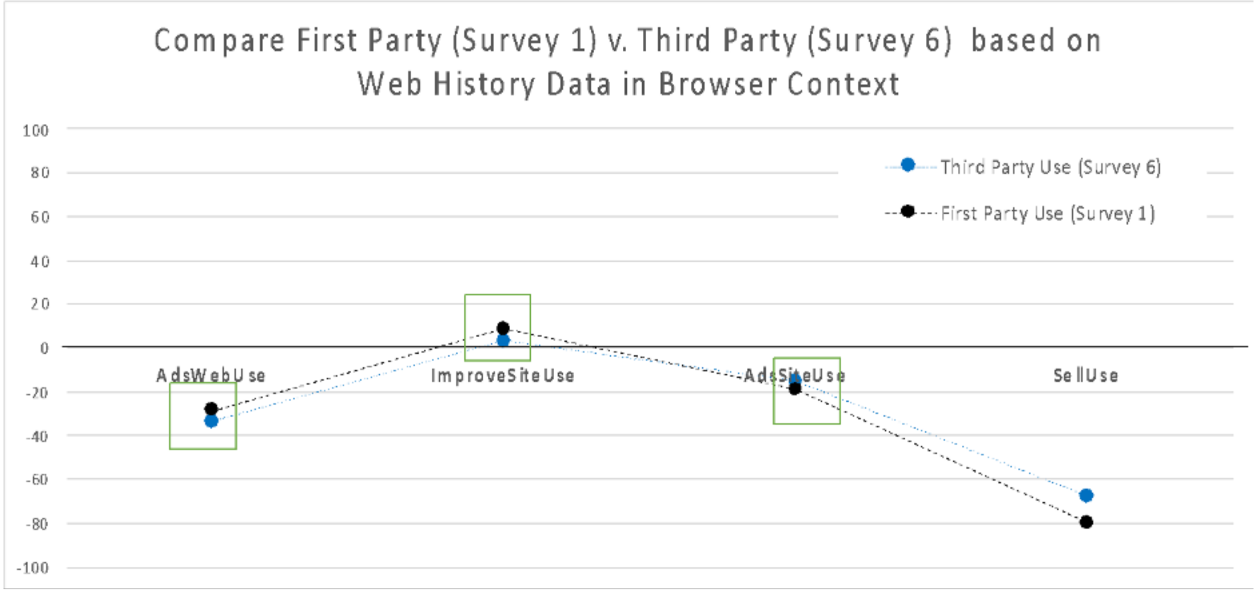

B. *Comparing First and Third Party in Creating Inferences*

Where the previous analysis compared first versus third parties in the use of knowledge, we turn to examine whether limiting a third party from even creating inferences or accessing raw data would be a privacy solution for respondents. For example, a first party (e.g., search engine, social network, browser) may represent that only they can create inferences from the data gathered. We examined whether limiting third parties from creating inferences was a privacy solution by comparing Survey 2 (where first party creates and uses inferences) to Survey 4 (where a data broker is allowed to create and use inferences without accessing the raw data –which is called multi-party computation) as shown in Figure 14.

Figure 14:

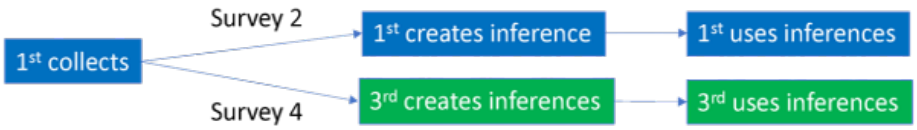

Figure 15: Comparison of Surveys 2-4 Testing Whether Third Party Purchasing or Creating Inferences Impact Meeting Privacy Preferences

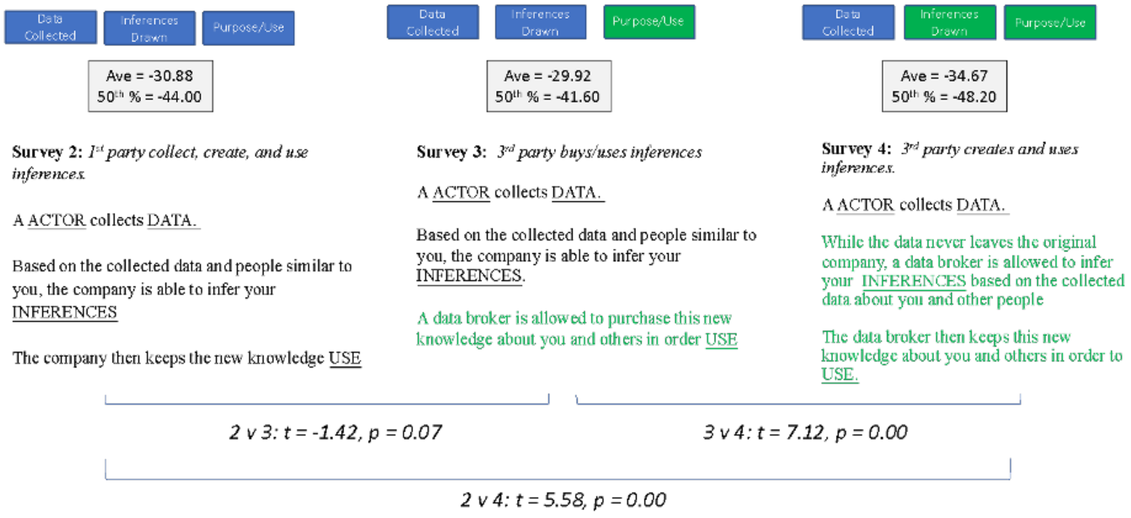

Survey 2: *1st party collect, create, and use inferences.*

A ACTOR collects DATA.

Based on the collected data and people similar to you, the company is able to infer your INFERENCES

The company then keeps the new knowledge USE

Survey 3: *3rd party buys/uses inferences*

A ACTOR collects DATA.

Based on the collected data and people similar to you, the company is able to infer your INFERENCES.

A data broker is allowed to purchase this new knowledge about you and others in order USE

Survey 4: *3rd party creates and uses inferences.*

A ACTOR collects DATA.

While the data never leaves the original company, a data broker is allowed to infer your INFERENCES based on the collected data about you and other people

The data broker then keeps this new knowledge about you and others in order to USE.

Respondents judged the creation and use of inferences by the first party (-30.88) to be more appropriate as compared to a third party (-34.67; t = 5.5823, p<0.00). However, neither scenario met their privacy expectations on average. The Appendix contains a detailed graph for each context (Figures A36-40), where the trend holds across contexts that the respondents judged the creation and use of inferences by either first or third party to be a privacy violation.

C. Comparing First and Third Party in Collecting User Data

Finally, recent privacy solutions have positioned a first party—such as a social network, search engine, or browser—as the only party able to appropriately collect consumer data, create inferences, and use that knowledge (e.g., to place ads). The privacy solution being offered is to preclude any third-party trackers from collecting data as well as forbidding others from gaining access to inferences about their users or being able to target them with advertising. Conceptually, Figure 16 illustrates the two options: a first party collects data, creates inferences, and uses that new knowledge (Survey 2), or trackers collect data that a data broker uses to create inferences for further application (Survey 5).

Figure 16: Comparing Surveys 2 and 5

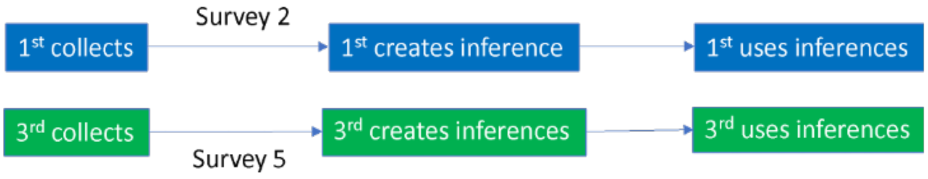

Figure 17: Vignette Design for Surveys 2 and 5

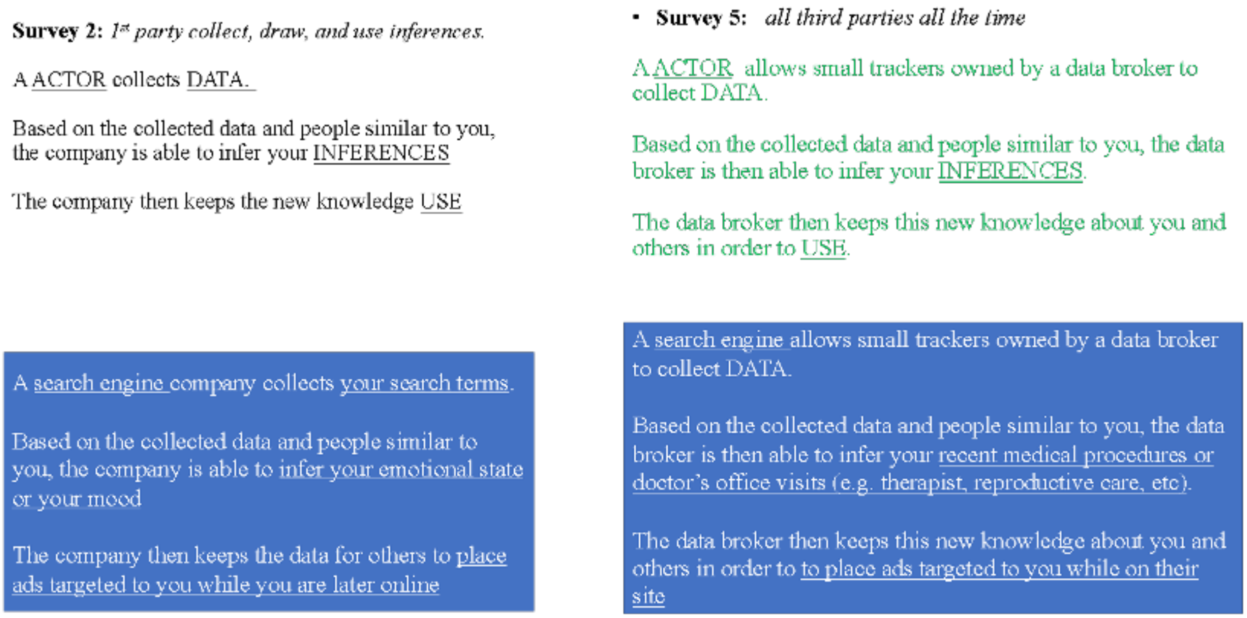

Survey 2: *1st party collect, draw, and use inferences.*

A ACTOR collects DATA.

Based on the collected data and people similar to you, the company is able to infer your INFERENCES

The company then keeps the new knowledge USE

A search engine company collects your search terms.

Based on the collected data and people similar to you, the company is able to infer your emotional state or your mood

The company then keeps the data for others to place ads targeted to you while you are later online

- **Survey 5:** *all third parties all the time*

A ACTOR allows small trackers owned by a data broker to collect DATA.

Based on the collected data and people similar to you, the data broker is then able to infer your INFERENCES.

The data broker then keeps this new knowledge about you and others in order to USE.

A search engine allows small trackers owned by a data broker to collect DATA.

Based on the collected data and people similar to you, the data broker is then able to infer your recent medical procedures or doctor's office visits (e.g. therapist, reproductive care, etc).

The data broker then keeps this new knowledge about you and others in order to to place ads targeted to you while on their site

Overall, respondents judged having either first or third parties collect data, create inferences, and use their data to be a privacy violation. The average rating the vignette is appropriate for a first party (-30.88) is greater than the rating of a third-party (-32.74; t=2.7011, p = 0.004); however, neither is deemed appropriate by the respondents.

The findings are consistent across contextual actors. The results comparing first party versus third-party collection and use of consumer data for browsers is shown in Figure 18, and the results are similar to the other contexts included in the Appendix.

Figure 18: Average Rating for Vignettes with Browser as Contextual Actor for each Type of Use by Search Engine Versus Third Party

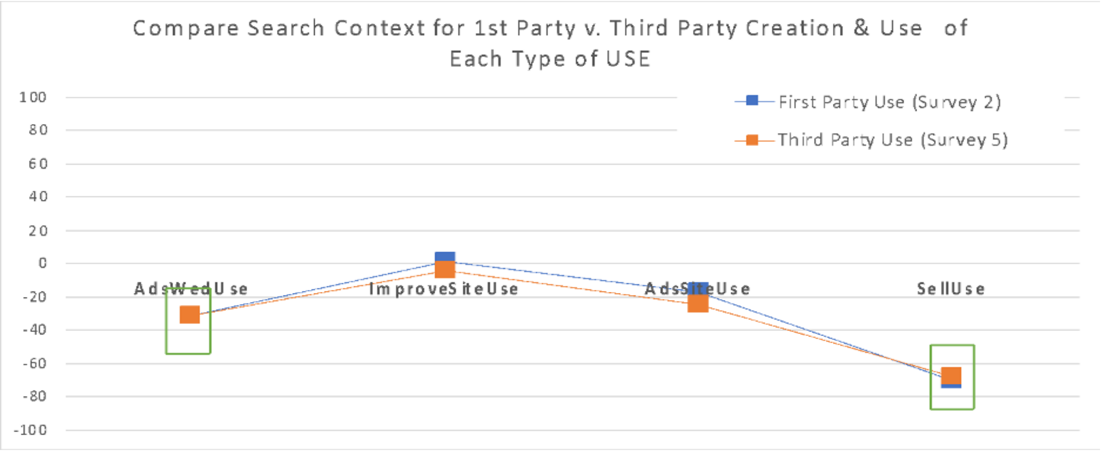

We then compared five special cases with the results in the Appendix. We used the case of tracking Web history data in the browser context, given the recent suggestion that removing third-party trackers from websites and only having the browser collect data, create inferences, and place ads would be a privacy solution. The results in Figure 19 show that the use of inferences by either the browser or a data broker is judged to be a privacy violation. Even the use of inferences to improve services on the site does not meet privacy expectations. In Figure 19, showing the average rating of vignettes focused on a browser as the contextual actor, the respondents judged the use of interest inferences by the browser to be more appropriate compared to a data broker in placing ads later online; however, both ratings are negative.

Figure 19: Comparison of Browser And Third-Party Creation and Use of Inferences

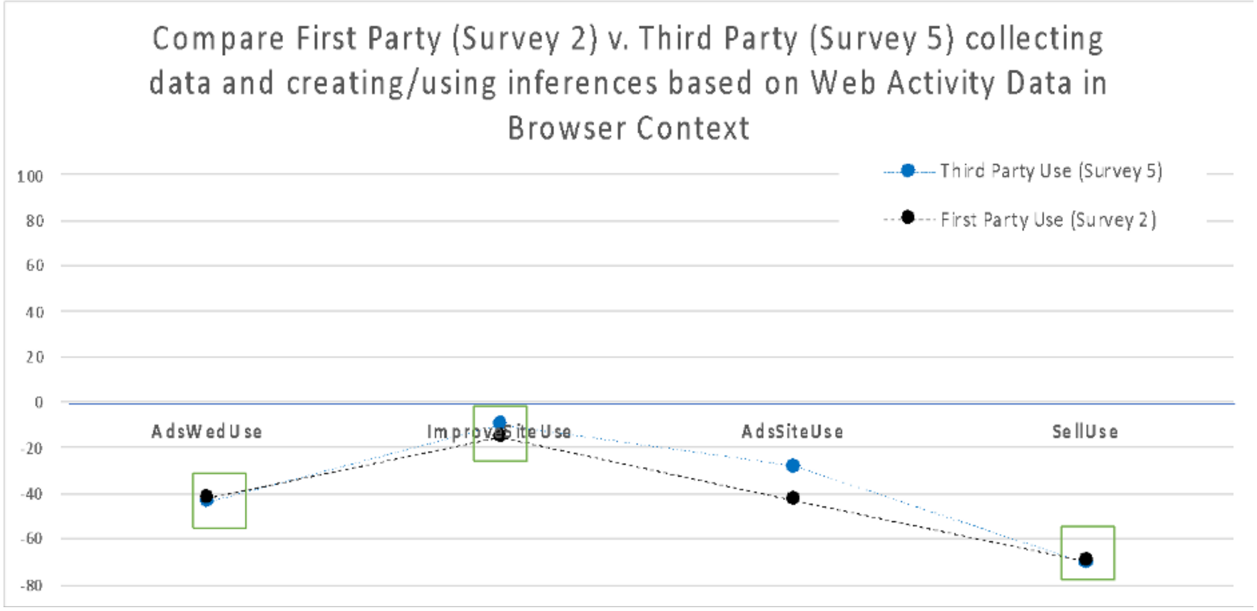

V. RESULTS 3: YOUR DATA STAYS ON YOUR DEVICE AS A PRIVACY SOLUTION

“Your Data Never Leaves Your Device”

Recently proposed privacy techniques center on the solution of having consumer data remain on a given device, and the company creates inferences about the individual for the purpose of advertising. The “data stays on your device” solution combines the two prior solutions: the focus shifts to (1) the creation and use of inferences rather than raw data and (2) the use of those inferences by the browser (first party) rather than third-party ad networks. Our results above reveal that respondents judged the creation and use of inferences to be the same or more of a privacy violation compared to the use of raw data. Respondents did not distinguish between first and third parties in the use of inferences for advertising. We decided to explicitly test the “data stays your device” solution with a factorial vignette survey using the same factors and values in Table 1. But this test used the following vignette template focusing on the browser collecting data and then creating inferences for the purpose of targeted advertising.

Figure 20: Vignette Template and Example Comparing Surveys 2, 3, and 7

.

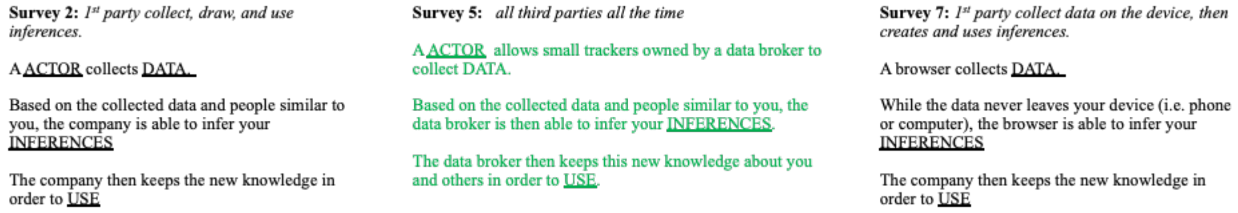

Survey 2: *1st party collect, draw, and use inferences.*

A ACTOR collects DATA.

Based on the collected data and people similar to you, the company is able to infer your INFERENCES

The company then keeps the new knowledge in order to USE

Survey 5: *all third parties all the time*

A ACTOR allows small trackers owned by a data broker to collect DATA.

Based on the collected data and people similar to you, the data broker is then able to infer your INFERENCES.

The data broker then keeps this new knowledge about you and others in order to USE.

Survey 7: *1st party collect data on the device, then creates and uses inferences.*

A browser collects DATA.

While the data never leaves your device (i.e. phone or computer), the browser is able to infer your INFERENCES

The company then keeps the new knowledge in order to USE

The privacy solution proposed is to have a browser collect data, then create inferences about individuals—either (1) from a curated list created by the company or (2) by an advertising company creating ad hoc inferences based on online events. All individuals that are a part of the same event are placed in an interest group category for targeted advertising (e.g., visit a specific website, click on a specific ad or link, etc.). The solution proposed is to have third-party trackers and data brokers blocked from having access to the raw data as in Survey 5 and the company not have access to the raw data (Survey 2) since the data remains locally in the browser on the device.

The privacy solution described in Survey 7 —where a browser collects your data, the data remains on the device locally, and inferences are created and used for advertising —is rated the same (average = -32.01) as Survey 5 (-32.74; t = -1.0062, p = 0.15), where third-party trackers collect the data, and a data broker creates inferences for use in advertising.

The results are the same when we limit the contextual actor to "browser" in Survey 5: while using a browser, third-party trackers collect data, sell the data to a data broker who creates inferences about the individual for use in advertising (Survey 5 (browser) = -32.32; t = -0.2535, p = .399). We see the same results when comparing specific inferences in Figure 21 and when comparing specific purposes of use in Figure 22. Respondents did not judge the privacy solution proposed —to have a browser collect data, keep the data on the local device, create inferences for use in advertising —to be a privacy solution over third-party tracking and the use of data brokers in advertising.

Figure 21 Average Rating for Vignettes with Browser as Contextual Actor for each Type of Inference by First Party Versus Third Party

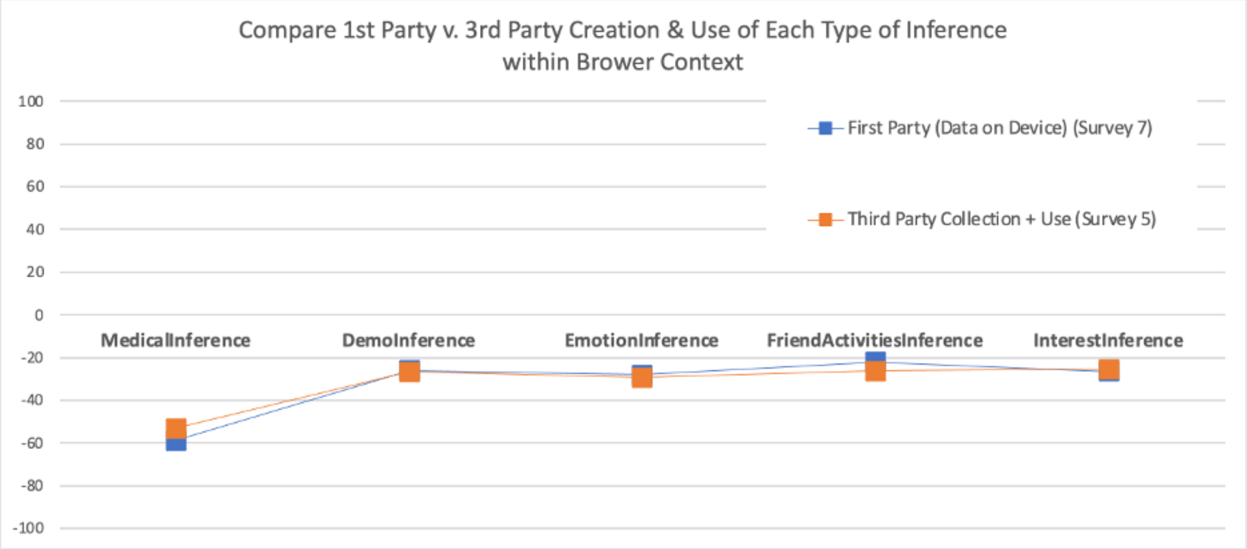

Figure 22: Average Rating for Vignettes with Browser as Contextual Actor for each Type of Use by First Party Versus Third Party

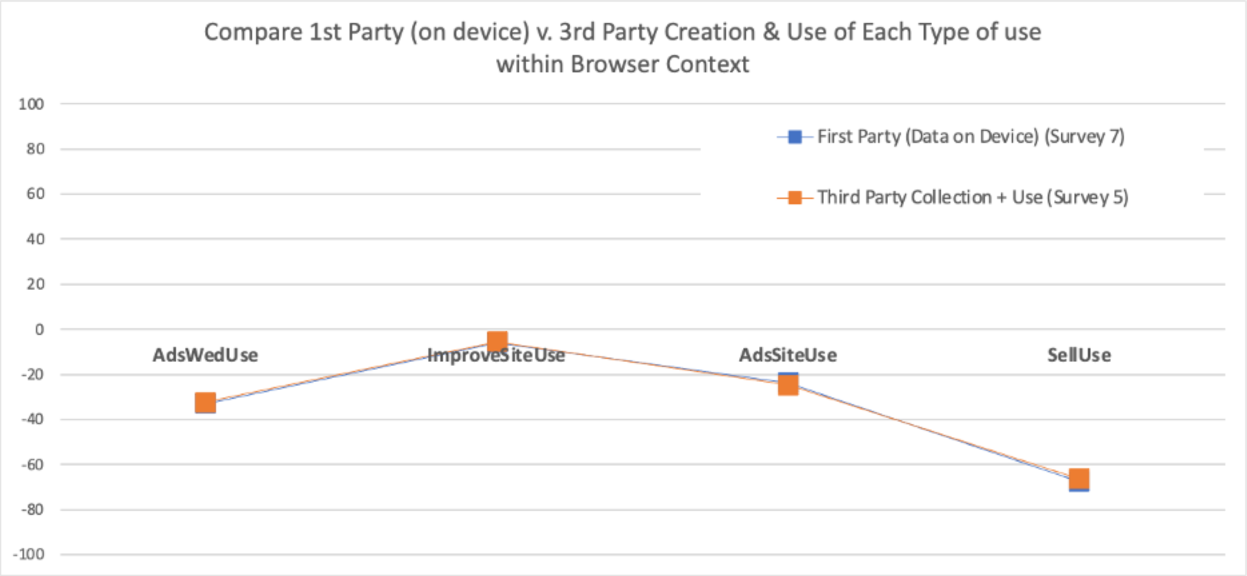
Compare 1st Party (on device) v. 3rd Party Creation & Use of Each Type of use within Browser Context
100
80
60
40
20
0
-20
-40
-60
-80
-100
First Party (Data on Device) (Survey 7)
Third Party Collection + Use (Survey 5)
AdsWedUse
ImproveSiteUse
AdsSiteUse
SellUse

VI. Summary of Key Findings

1. Creating, storing, and using inferences as a substitute for using raw data is not a privacy solution. Respondents judged the use of inferences compared to raw data the same or worse at meeting their privacy expectations. This finding held across types of data (e.g., location versus inferences based on location; search terms versus inferences based on search terms, etc.) and across purposes/uses. In fact, the use of raw data better meets privacy expectations when improving services and, in some circumstances, when placing ads on a website. This undermines the claims of certain techniques, such as Google's Privacy Sandbox, federated learning, or even the focus on using inferences versus collected data for advertising, as a plausible privacy solution. A key privacy technique promoted by the industry is to remove access to or delete raw data. However, advertisers' decisions are based on inferences and this empirical study shows that users care *more* about the collection, storage, and use of inferences for advertising than the comparable use of raw data. This mismatch clearly shows that removing access to raw data but keeping access to inferences is not sufficient to meet users' privacy preferences. These techniques are not judged to better meet the privacy preferences of users compared to the alternative.

2. Purpose matters. Respondents' judgements were sensitive to the purposes for which information is used. Specifically, when information flows served contextual ends and values (e.g., to improve services) they were found to be appropriate consistently across contextual actors (e.g., news site, search engine, and social network). The use of the same information to promote non-contextual purposes, such as to place ads later online or sell to others, did not meet privacy expectations and were judged more negatively as privacy violation. This is consistent with previous empirical work on uses of data for contextual versus noncontextual purposes—where respondents evaluated the use of data by an actor for a non-contextual purpose to be inappropriate.[99] (This finding undermines claims of companies that whether or not they share data is more important than the use of that data for a particular purpose.

3. **Multiparty computation**. When judging whether a third-party data broker should have access to raw data versus inferences based on that raw data, respondents judged selling inferences (#3) to be a privacy violation of greater magnitude than selling raw data (#6). Respondents slightly preferred data brokers to buy the data rather than the inferences based on that same data (however, both

[99] *See* Martin & Nissenbaum, *supra* note 20, at 210.

scenarios were negative). In addition, our studies showed that respondents evaluated creation and use of inferences by a data broker without the raw data leaving the original company (**multiparty computation** in #4) as less aligned with privacy expectations than either (a) selling raw data directly to a data broker (#6) or (b) allowing trackers to collect the same data for a data broker (#5).

4. **First parties and third parties**. Constraints on data flows based on a distinction between first parties and third parties are not necessarily viewed as a privacy solution. Respondents judged the creation and use of inferences by a third-party data broker to be a privacy violation to the same degree as when performed by the first party (e.g., search, news, etc.). For this case, respondents' evaluations were more nuanced than this dichotomy allows. In fact, respondents judged a first party selling their inferences to be *worse* than the data broker selling their inferences. For search in particular, respondents rated a third party broker the same or better than the search firm using inferences to place ads later online.

5. **Sandbox**. "The data stays on your device" is not a privacy enhancing technique. The proposed Sandbox privacy solution does not fully address privacy expectations of users or provide a solution over the alternative of third-party trackers and ad networks placing personalized ads. Specifically, a browser collecting data and creating inferences to place ads online did not meet privacy preferences and was rated the same or lower at meeting privacy preferences compared to third-party trackers and data brokers placing ads.

Map of Results

Inferences v. Raw Data		
In general	For each type of data (e.g., search, location, engagement, etc.), the creation and use of inferences were judged to be the same or worse at meeting privacy preferences of users.	III.A. Figure 4. Appendix A1-5
Specific results	The use of raw data or inferences to place ads later online for search, browser, and social networks was a privacy violation.	Figures 5a and 5b Appendix A6-10
Interesting results	The use of search data to place ads on a search engine is not a privacy violation, but the use of inferences based on history is a privacy violation.	Figure A9a and A9b.
	Creation and use of inferences by data brokers (multiparty computation) do not meet the privacy preferences of consumers and are worse at meeting those preferences compared to a broker buying and using raw data.	III.B.2
1st versus 3rd Party		
In general	The respondents judged a first party using inferences met their privacy preferences to the same degree as a data broker who purchases and uses the inferences for the same purpose.	IV.A. Figure 10, 11 Appendix Figures A21-25
Specific results	Browsers and third-party data brokers are judged statistically the same for placing ads on a site and for targeting users with ads later online (both are privacy violations).	Figure 12
	The use of inferences by either the browser or a data broker does not meet privacy preferences of users. Even the use of inferences to improve services on the site does not meet privacy preferences for either a browser or data broker.	Figure 19
Interesting results	However, having a third party later sell those inferences is judged as better meeting privacy preferences than having the first party sell those inferences.	Figure 10, 11
	Respondents judged that the creation and use of inferences by the first party better meet their privacy preferences compared to by a third-party. However, neither scenario met their privacy preferences on average.	IV.B. Figure 15.
In addition,		
Data stays on device	Respondents did not judge the privacy solution proposed—to have a browser collect data, keep the data on the local device, create inferences for use in advertising—to be a privacy solution over third-party tracking and data brokers using the data for targeted advertising.	V. Figure 21-22
Purpose or u se	Respondents' judgements were sensitive to the purposes for which information is used. Specifically, when information flows served contextual ends and values (e.g., to improve services), the data flows were found to be appropriate (i.e., more OK) consistently across contextual actors (e.g., news site, search engine, and social network).	Appendix Figures A1-5 with Tables A1-5.

VII. IMPLICATIONS

A. Implications for Practice

As we show in this Article, specific inferences made about the user determine whether privacy is protected or violated. Simply removing access to raw data—while exposing inferences—is not sufficient to argue that privacy has been protected. Making this determination requires a detailed analysis of available inferences. In the case of Google's Privacy Sandbox, privacy protection relies on the curation of Topics API (in particular, ensuring that this API does not contain any topics that users consider sensitive). To the extent protected audiences are not curated and the categories are determined by website operators and their marketing partners at will, they may (or may not) violate users' privacy. This is true even if raw data about users' browsing behavior is protected and inferences are not linkable to the specific users, except to show online advertising based on these inferences.

B. Implications for Policy

This Article has implications for policies, regulations, and the courts. First, even where a delineation between first and third parties is clear and consistent, it does not consistently distinguish between flows that are appropriate and defensible on ethical grounds and those that are not. Further, even when restricting access to third parties limits the flow of data to non-contextual recipients, our findings challenge the presumption that first parties are entitled to freely access and use data without restraint. We expose this non-sequitur, due to a fixation on ownership as a bright line between first and third parties with argument and empirical demonstration. Law, policy, and regulation miss the mark by not recognizing that those asserting first party privileges may not be acceptable contextual actors (e.g., Google in its capacity as an ad network), third- parties may have (limited) access privileges in contextually appropriate capacities (e.g., Sonos manufacturer on Amazon or truck drivers delivering products). All actors are subject to restrictions on the purposes for which data is used.

This means that rules or regulations that require notification if sharing with a third party, or require users to opt in to sharing data with a third party, or even limit whether a third party can receive consumer data would not necessarily provide privacy protections since first parties would still be able to collect, combine, and use the data for non-contextual purposes. Further, such rules would limit the appropriate flow of data to contextual actors that are third parties—such as payment processors, delivery drivers, or cybersecurity services—and falsely pit "protecting privacy" (in the form of not sharing data with third parties) against functionality.

Second, rules, regulations, and laws focused on data retention or data minimization may miss the use of inferences based on that consumer data in privacy violations. Regulations that force consumer data to be deleted but do not address the storage and later use and sharing of knowledge derived from it fail to resolve significant sources of privacy violating behaviors. They may remedy this omission by addressing legitimacy of inferences that are drawn from this data, even after its deletion. Our findings were consistent with what CI would have predicted, namely that respondents' privacy judgments are highly variable across the types of inferences drawn, for example, demographic versus medical condition, and across purposes served by these inferences. Our findings would support policies that treat inferences just as seriously as underlying data; if the latter calls for deletion, so must the former.

Third, as already noted, our results unequivocally confirmed that people care about the purpose for which collected data is used. Rules and regulations that focus on what information is collected, stored, or shared but do not take into account the purpose served by these practices are missing one of the key hallmarks of privacy violating behavior, whether attributed to third parties or to first parties, for example, in the landscape of targeted advertising. Our studies show that how information was used, specifically, whether in service of contextual versus non-contextual ends was even more important for our respondents than what data was collected or with whom the data was shared. For example, sing personal data to serve the purpose of targeting consumers in online ads was judged to be a privacy violation whether performed by a first or third party.

Finally, our studies debunked the idea that data never leaving devices was sufficient for privacy. In fact, a browser collecting consumer data on the user's device and drawing inferences for the purpose of targeted advertising was judged to be a privacy violation. If policy is to curtail data practices that address fundamental privacy concerns of consumers, our results demonstrate that a technical focus on where data is held or where processing takes place —on device or elsewhere—makes little difference. It matters most whether the purposes for which the inferences are used are contextually appropriate.