DOI:10.1145/3597173

Alessandro Acquisti and Ryan Steed

# Privacy
# Learning to Live with Privacy-Preserving Analytics

*Seeking to close the gap between research and real-world applications of PPAs.*

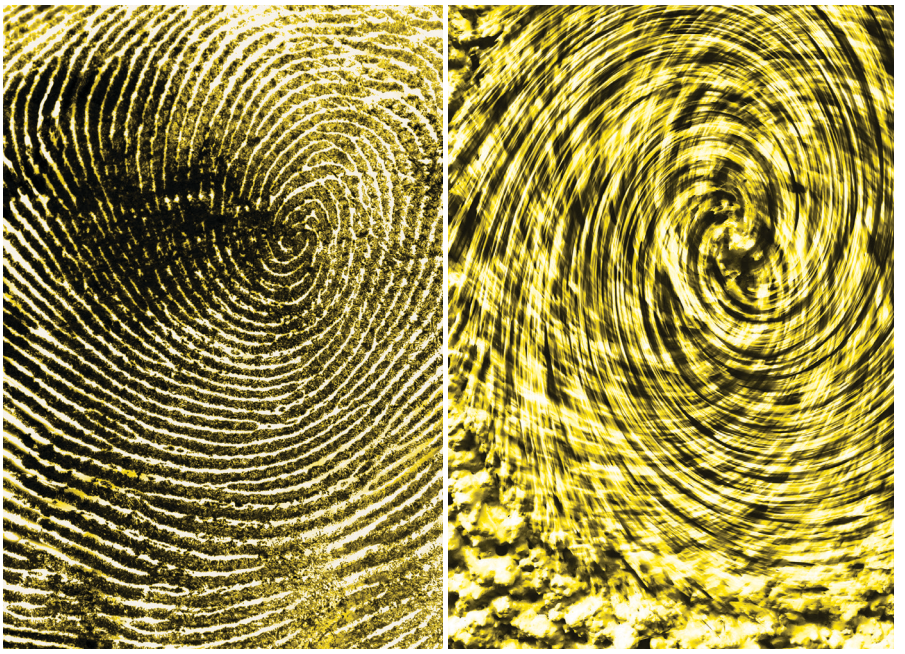

Over the past few decades, computer scientists and statisticians have developed tools to achieve the dual goals of protecting individuals' private data while permitting beneficial analysis of their data. Examples include techniques and standards such as blind signatures, *k*-anonymity, differential privacy, and federated learning. We refer to such approaches as *privacy-preserving analytics*, or PPAs. The privacy research community has grown increasingly interested in these tools. Their deployment, however, has also been met with controversy. The U.S. Census Bureau, for instance, has faced a lawsuit over its differentially private disclosure avoidance system; opposition to the new privacy plans garnered support from both politicians and civil rights groups.[3,11]

In theory, PPAs can offer a compromise between user privacy and statistical utility by helping researchers and organizations maneuver through the trade-offs between disclosure risks and data utility. In practice, the effects of these techniques are complex, obfuscated, and largely untested. While the interest in PPAs has been growing particularly among computer scientists and statisticians, there is a need for complementary social science research on the downstream impacts of these tools[13]—that is, the concrete ripples those technologies may cast for individuals and societies. In this Viewpoint, we advocate for an interdisciplinary, empirically grounded research agenda on PPAs that connects social and computer scientists.

A complete survey of this area of research would vastly exceed the space limitations for this Viewpoint.[a] Instead, after briefly summarizing what PPAs are, we focus on discussing how their rising deployment in real-world applications has been a cause of controversy, but also why PPAs are promising tools that both deserve and necessitate interdisciplinary research attention.

a For an annotated bibliography of additional references not included in this Viewpoint, see https://bit.ly/3I2lGHN

## What Are Privacy-Preserving Analytics?

The term "privacy-enhancing technologies" (PETs) has been used to refer, broadly, to methods, tools, or technologies devised to protect individual privacy. By the term "privacy-preserving analytics," we refer to a particular subset of PETs that attempt to protect individual data while permitting some degree of data analytics.

Some PPAs are defined by specific methods for ensuring private analysis. Differential privacy, for example, can be applied to certain kinds of federated learning—a class of machine learning methods in which a model is trained across multiple devices without send-

ing the original training data to a centralized location—as well as to stochastic gradient descent (a common optimization technique in deep learning). Homomorphic encryption—another PPA on which applications such as blind signatures are based—permits computations (including predictive analytics) performed entirely on encrypted data without access to the decrypted version of the data. Similarly, the field of multiparty computation (MPC) is devoted to protocols that allow multiple parties to participate in aggregate computations without revealing their private data.

PPAs can also be described by the privacy definitions they satisfy. Standards such as *k*-anonymity—which requires an individual's attributes match no less than *k* other records in a dataset—or *l*-diversity are commonly used to evaluate the efficacy of particular PPA approaches, though these standards have since been shown to be vulnerable to certain re-identification attacks. For future-proof, formal guarantees, PPAs may be required to conform to differential privacy, which bounds the influence of any individual's data on output statistics. These standards are often achieved using the methods in the previous paragraph—for example, Google's abandoned Federated Learning of Cohorts (FLoC) proposed to provide *k*-anonymity through federated clustering.[8] Differential privacy can also be federated to provide local differential privacy, where each device applies a differentially private mechanism before aggregation.

## Growing Applications

While several government organizations have been using, for a long time, an array of traditional statistical disclosure limitation methodologies, in recent years industry and government deployments of novel PPAs have become more common. Perhaps the most high-profile example is the U.S. Census Bureau's new Disclosure Avoidance System (DAS), which provides differential privacy guarantees for the 2020 Decennial Census. A team of bureau scientists realized the need for stronger protections after internal research suggested that published census data could be linked with commercial data to re-identify more than 52 million individuals.[3] The bureau had previously leveraged differential privacy to release additional, previously unavailable data through the 2008 On-The-Map tool.

**PPAs are promising tools that both deserve and necessitate interdisciplinary research attention.**

Another prominent example is Google's proposal to replace third-party cookies with interest-based alternatives. Using a federated approach, Google initially proposed to automatically group users into *k*-anonymous cohorts based on their personal data. More recently, Google proposed Topics API, a taxonomized approach to grouping users into much larger interest categories.[8]

PPAs are often proposed as a means to facilitate data sharing with researchers. In partnership with Social Science One, Facebook released in 2020 a large, differentially private set of URLs shared on the platform to aid studies on social media and democracy.[9] Google, LinkedIn, and Microsoft have also released public datasets with differential privacy.[6] Researchers have also extensively discussed and debated the use of *k*-anonymity to protect individuals' health data and satisfy the Health Insurance Portability and Accountability Act (HIPAA) Privacy Rule.[10]

Less publicized applications also exist. Differential privacy has been applied in deployed systems including the Facebook advertising stack: various internal SQL systems at Google, Uber, Oracle, SAP, and other tech companies; and local telemetry in Apple devices.[6] Federated learning is even more widespread, deployed in notable products such as Google's GBoard keyboard.

## Growing Dissent

Our Viewpoint is motivated by the observation that, as PPAs have grown in

**Many of the implementations of PPAs have raised controversy *precisely* because they engender a trade-off between statistical utility and data privacy.**

popularity in research circles and as their real-world applications proliferate, their deployment has also been met by controversy. For instance, the introduction of differential privacy into Census Bureau operations has been particularly controversial, with a number of scholars, politicians, and civil rights groups questioning its use. Google delayed, and then shelved its initial third-party cookie replacement plans after privacy advocates protested, asserting it would make fingerprinting and discriminatory advertising easier.[5] Social Science One publicly disagreed with Facebook's interpretation that differential privacy was required for the Facebook URLs dataset, arguing that merely de-identified or aggregated academic data sharing should be enough to satisfy the General Data Protection Regulation and Facebook's FTC consent decree.[9] And researchers have questioned the use of privacy-preserving techniques in high-stakes settings (such as healthcare), fearing the utility trade-offs may especially affect vulnerable minority groups for whom PPAs may not perform as well.

We believe many of the implementations of PPAs have raised controversy *precisely* because they engender a trade-off between statistical utility and data privacy. In the census differential privacy example, demographers and politicians protested the deleterious effects that injected noise might have on critical research and political processes. In the FLoC example, privacy advocates protested the opposite—that Google sacrificed user privacy in order to preserve advertisers' ability to precisely target consumers.[5]

Granted, many studies in computer science and statistics already analyze these trade-offs. However, and key to our point, most approach the issue from formal perspectives, analyzing the trade-offs between abstract privacy parameters (for example, epsilon, for differential privacy) and abstracted loss functions (for example, root mean squared error).[1] In practice, the utility trade-offs reference concrete, not abstract, downstream benefits and harms; politicians are concerned that their electoral maps will be overhauled to the advantage of their opponents, and advertisers are concerned that the new Topics API will stifle sales.[12] Privacy trade-offs, on the other hand, often live in the realm of risk. Data stewards (as in the case of the census) must assess the chance of a potentially catastrophic reconstruction attack, while translating complex, probabilistic privacy parameters into concrete confidentiality guarantees.[11] Stakeholders may have practical preferences and concerns over these outcomes. However, the mostly theoretical literature makes it difficult to assess the ability of PPAs to satisfy all parties. We argue there is a need for research emphasizing empirical methods and participatory approaches involving a diverse set of stakeholders focusing on downstream outcomes and a more holistic set of consequences.

## Downstream Implications, and Looking at the Bigger Picture

By downstream outcomes, we refer to empirical studies of the organizational and managerial considerations behind PPA development, along with the impact of their usage on those organizations in consumer products, in research, and in policymaking.

For example, in recent work,[15] we considered the controversy over the introduction of differential privacy to the 2020 Census and looked at potential but practical (and empirically measurable) implications of a deployment of differentially private algorithms in policymaking. Every year, Census estimates are used to guide the allocation of over $1.4 trillion in federal funding, including more than $16 billion dollars

in education funding divided among school districts in 2019. The allocation of that funding relies on a formula based on the Census Bureau's estimate of the number of children in poverty across the country. We simulated the effects of two different kinds of statistical uncertainty on the allocation of these Title I grants: the effect of noise injected to achieve differential privacy, and the effect of *existing* data error in the Census Bureau's estimates. What we found was that, while enforcing differential privacy did cause "misallocation" of funding relative to the official allocations, the misallocations caused by simulated data error were much larger. In other words, our results indicated the decision to augment privacy only adds to much larger costs of statistical uncertainty in census-guided grant programs. Weakening privacy protections would do little to mitigate these deeper disparities—in fact, if respondents are not confident in the privacy of their responses, the usefulness of census data will be further reduced, especially for harder-to-count groups.

In our study, the addition of differential privacy exposed deeper flaws in the way census-guided funding programs distribute the impacts of statistical uncertainty. But we also found that simple policy changes can reduce these impacts, ameliorating the costs of existing uncertainty and making the addition of stronger privacy protections less daunting. Our study shows the trade-offs involving PPAs may not be as stark as they seem: When PPAs conflict with existing data infrastructures, scientists and policymakers have an opportunity to revisit and improve statistical practices. Indeed, Oberski and Kreuter[13] argue that the constraints imposed by differential privacy (on repeated and granular analysis) could act as a barrier to *p*-hacking and other dodgy techniques.

We can already see glimpses of more robust, privacy-preserving research in studies of prominent, differentially private datasets. Evans and King[7] adapt existing methods for correcting naturally occurring data error to help construct valid linear regression estimates and descriptive statistics and evaluate their corrections on the differentially private Facebook URLs Dataset released in 2020.[9] Buntain et al.[4] develop a robust measure of ideological position on the same dataset and conduct an extensive empirical study of news platforms and audience engagement. In many cases, these more robust estimators significantly reduce the costs of privacy. For example, Agarwal and Singh[2] propose methods for noise-adjusted causal inference, recovering the results of a seminal study on import competition with differentially private 2020 Decennial Census data. And PPAs may require other, less technical changes with beneficial side effects—for example, interviews by Sarathy et al.[14] with differential privacy practitioners suggest a need for better data documentation and more context-specific education for data analysts.

## What Is Next for PPAs?

In theory, PPAs can offer a compromise between user privacy and statistical utility. In practice, the effects of these techniques are still to a large extent untested. PPAs still pose challenges for scientists and policymakers. Work on correcting methods to account for noise and other privacy protections is still ongoing, and not all costs can be mitigated by more robust statistical practices. For example, studies by Sarathy et al.[14] and others describe practitioners' struggles with understanding and successfully applying PPA tools. And not all applications may be suitable for PPAs: some uses may be more suited to more traditional legal protections, and some uses may entrench harmful data practices. For example, the EFF protested Google's third-party cookie replacement not only because it could aid browser fingerprinting, but also because any form of targeted advertising can be used for exploitation and discrimination.[5] On the other hand, some uses of PPAs could help prevent harms like these: for example, PPAs could help facilitate sensitive data sharing with auditors, helping hold online platforms accountable for discrimination and other harms.

The controversies surrounding the deployment of PPAs highlight a critical gap in current research: While the interest in PPAs has been growing particularly among computer scientists and statisticians, there is still a paucity of applied, empirical research on the downstream implications of the development and deployment of PPAs across a variety of usage cases. Thus, there is a need for complementary social-science research on the ripples those technologies may cast—that is, the impacts of these tools. In the past couple of years, more applied research has started appearing. Our hope is that computer and social scientists will increasingly collaborate on research in this area to build generalizable knowledge and develop a grounded theory of the trade-offs involved, thus highlighting both the benefits as well as the burdens of PPA adoption for various stakeholders in real-world settings. C


**References**

1. Abowd, M. and Schmutte, I.M. An economic analysis of privacy protection and statistical accuracy as social choices. *American Economic Review 109*, 1 (2019), 171–202.
2. Agarwal, A. and Singh, R. Causal Inference with Corrupted Data: Measurement Error, Missing Values, Discretization, and Differential Privacy. (2022); https://bit.ly/3I2wi9j.
3. Bouk, D. and boyd, d. Democracy's Data Infrastructure. Knight First Amendment Institute. https://bit.ly/3pwPANM.
4. Buntain, C. Measuring the Ideology of Audiences for Web Links and Domains Using Differentially Private Engagement Data. (2021); https://bit.ly/41oCyzj
5. Cyphers, B. Google's FLoC Is a Terrible Idea. Electronic Frontier Foundation. (Oct. 2021); https://bit.ly/42LiWX9.
6. Desfontaines, D. A list of real-world uses of differential privacy. Ted is writing things. (2021); https://bit.ly/3BkS8kW.
7. Evans, G. and King, G. Statistically valid inferences from differentially private data releases, with application to the Facebook URLs dataset. *Political Analysis 31*, 1 (Jan. 2023); https://bit.ly/44PEtQj.
8. Goel, V. Get to know the new Topics API for Privacy Sandbox. Google. (May 16, 2022); https://bit.ly/3Bjk34t.
9. King, G. and Persily, N. Unprecedented Facebook URLs dataset now available for academic research through Social Science One. Social Science One. (202)); https://bit.ly/3MlOCwZ.
10. Malin, B. et al. Never too old for anonymity: A statistical standard for demographic data sharing via the HIPAA Privacy Rule. *Journal of the American Medical Informatics Association 18*, 1 (Jan. 2011); https://bit.ly/42tQyZU.
11. Nanayakkara, P. and Hullman, J. What's driving conflicts around differential privacy for the U.S. Census. *IEEE Secur. Privacy* (2022); https://bit.ly/42RKU3D.
12. Nguyen, G. Google's Topics API: Advertisers share concerns about topic diversity and other potential challenges. Search Engine Land. (2022); https://bit.ly/3VXKupP.
13. Oberski. D.L. and Kreuter, F. Differential privacy and social science: An urgent puzzle. *Harvard Data Science Review 2*, 1 (Jan. 2020); https://bit.ly/44PQvcm.
14. Sarathy, J. et al. Don't Look at the Data! How Differential Privacy Reconfigures the Practices of Data Science. (2023); https://bit.ly/3BkOmb0.
15. Steed, R. et al. Policy impacts of statistical uncertainty and privacy. *Science 377*, 6609 (Aug. 2022), 928–931; https://bit.ly/3BgQXTB.



**Alessandro Acquisti** (acquisti@cmu.edu) is the Trustees Professor of Information Technology and Public Policy at the Heinz College, Carnegie Mellon University, Pittsburgh, PA, USA.

**Ryan Steed** (ryansteed@cmu.edu) is a Ph.D. student at the Heinz College, Carnegie Mellon University, Pittsburgh, PA, USA.



The authors thank Priyanka Nanayakkara and Roy Rinberg for feedback on an earlier version of this Viewpoint. Alessandro Acquisti gratefully acknowledges support from the Alfred P. Sloan Foundation.