

UC Berkeley

CEGA Working Papers

Title

Big Data Privacy in Emerging Market Fintech and Financial Services: A Research Agenda

Permalink

<https://escholarship.org/uc/item/9nw1w6nd>

Authors

Blumenstock, Joshua E.
Kohli, Nitin

Publication Date

2023-08-22

DOI

10.26085/C3WK53

Series Name: WPS
WPS Paper No. 231
Issue Date: 8/21/23

Big Data Privacy in Emerging Market Fintech and Financial Services: A Research Agenda

Joshua E. Blumenstock and Nitin Kohli



CEGA

Center for Effective Global Action

Working Paper Series

Center for Effective Global Action
University of California



This paper is posted at the eScholarship Repository, University of California. http://escholarship.org/uc/cega_wps Copyright © 2023 by the author(s).

The CEGA Working Paper Series showcases ongoing and completed research by faculty affiliates of the Center. CEGA Working Papers employ rigorous evaluation techniques to measure the impact of large-scale social and economic development programs, and are intended to encourage discussion and feedback from the global development community.

Recommended Citation: Blumenstock, Joshua E.; Kohli, Nitin (2023): Big Data Privacy in Emerging Market Fintech and Financial Services: A Research Agenda. CEGA Working Paper Series No. WPS-231. Center for Effective Global Action. University of California, Berkeley. Text. <https://doi.org/10.26085/C3WK53>

Big Data Privacy in Emerging Market Fintech and Financial Services: A Research Agenda*

Joshua E. Blumenstock
UC Berkeley
jblumenstock@berkeley.edu

Nitin Kohli
UC Berkeley
nitin.kohli@berkeley.edu

August 2023

Abstract

The data revolution in low- and middle-income countries is quickly transforming how companies approach emerging markets. As mobile phones and mobile money proliferate, they generate new streams of data that enable innovation in consumer finance, credit, and insurance. Already, this new generation of products are being used by hundreds of millions of consumers, often to use financial services for the first time. However, the collection, analysis, and use of these data, particularly from economically disadvantaged populations, raises serious privacy concerns. This white paper describes a research agenda to advance our understanding of the problem and solution space of data privacy in emerging market fintech and financial services. We highlight five priority areas for research: conducting comprehensive landscape analyses; understanding local definitions of “data privacy”; documenting key sources of risk, and potential technical solutions (such as differential privacy and homomorphic encryption); improving non-technical approaches to data privacy (such as policies and practices); and understanding the tradeoffs involved in deploying privacy-enhancing solutions. Taken together, we hope this research agenda will focus attention on the multi-faceted nature of privacy in emerging markets, and catalyze efforts to develop responsible and consumer-oriented approaches to data-intensive applications.

*We thank Seth Garz, Zoe Kahn, Dan Cassara, and Bilal Zia for constructive feedback on this paper. We gratefully acknowledge financial support from Wells Fargo, the Institute for Business and Social Impact, the Bill and Melinda Gates Foundation, the Center for Effective Global Action, and the National Science Foundation under CAREER Grant IIS-1942702.

Introduction

Today, 95% of the global population has mobile-phone coverage, and the number of people who own a phone is rising fast. These devices generate troves of personal data on billions of people – including on people who live on a few dollars a day, and who historically have been less visible in national statistics and other administrative datasets. This “data revolution” is transforming the way that private companies approach emerging markets, and is revolutionizing different sectors including consumer finance (Heaton, Polson and Witte, 2017; Hambly, Xu and Yang, 2021), credit (Francis, Blumenstock and Robinson, 2017), and insurance (Nagendra, Narayanamurthy and Moser, 2020).

In one particularly striking example, new sources of data – and new computational tools for processing those data – have transformed the consumer lending landscape in many developing countries. A suite of immensely popular new “digital credit” products use data from mobile phones to construct alternative credit scores, which are now used to originate loans to hundreds of millions of historically unbanked individuals (Francis, Blumenstock and Robinson, 2017). The key insight behind this technology is that some people, such as those who frequently make international calls or who have more Facebook friends than others in the same area, are more likely to repay their debts. Machine learning algorithms can detect these patterns (Björkegren and Grissen, 2020) and construct credit scores for hundreds of millions of mobile phone owners who would otherwise been excluded from formal financial services due to a lack of collateral, traditional financial history, or access to a bank.

The data generated from mobile phones is also changing how governments and humanitarian organizations approach public health management (Oliver et al., 2020), social protection (Blumenstock, 2020), and humanitarian aid (Aiken et al., 2022). With some modifications, the same algorithms that Google, Facebook, Amazon, and other companies use to target advertisements to people online can also be used to target resources to people living in poverty. These algorithms leverage personal data from mobile-phone networks and satellite imagery to identify “digital signatures” of poverty (Blumenstock, Cadamuro and On, 2015; Blumenstock, 2018b). For example, in many African countries wealthier individuals tend to make more international phone calls – as observed in mobile-phone networks – while poorer individuals are more likely to live in houses with thatched roofs – as observed in satellite images (Blumenstock, 2018a).

In addition to enabling new consumer-facing applications, the responsible use of private “big”

data can also improve the functioning of entire sectors and industries. For instance, data sharing between financial institutions can facilitate more effective approaches to fighting money laundering (Sangers et al., 2019; van Egmond and Rooijakkers, 2021) and fraud (Hipgrave, 2013; Henecka and Roughan, 2014).¹ This, in turn, can reduce the costs that are borne by consumers, institutions, and society at large.

At the same time, the widespread use of personal data – particularly from economically disadvantaged populations – raises serious concerns about the implications for the personal privacy of consumers (De Montjoye et al., 2018). As mobile phones and other digital technologies become ubiquitous in everyday life, the data captured by such devices has become particularly sensitive. After all, these data describe more than just the digital transactions of an individual; they can also describe their way of life by revealing their associations, their behaviors, their habits, and their beliefs (Kohli, 2021).²

Thus, while these new sources of data offer immense opportunities for new commercial and humanitarian applications, it is imperative to find consistent and secure ways to process and analyze the sensitive data. Unfortunately, contemporary approaches to data privacy in many low- and middle-income countries are inconsistent and often nonexistent. A wide range of technical and non-technical “solutions” exist, but most are not developed with emerging market consumers or companies in mind. Regulatory frameworks are patchwork and rarely enforced (Ademuyiwa and Adeniran, 2020). Furthermore, scholars have questioned the ethics of using such data (Taylor, 2016), and point out that data sharing arrangements are frequently extractive, rather than collaborative (Abebe et al., 2021).

The failure to protect personal privacy – and the high-profile examples of the harm caused by such failures (cf. Crawford and Schultz, 2014a) – threatens to undermine financial and humanitarian applications that could otherwise positively impact society. However, this need not be the case. Appropriately leveraged, privacy-enhancing technologies, policies, and practices have the ability to simultaneously enable innovative uses of big data while providing robust and meaningful privacy protections for personal data.

¹See also *Data Sharing for the Prevention of Fraud: Code of practice for public authorities disclosing information to a specified anti-fraud organisation under sections 68 to 72 of the Serious Crime Act 2007* presented to the UK Parliament, March 2015.

²For more examples, see Justice Sotomayor’s concurrence in *US v. Jones*, 565 U.S. 400, 132 S. Ct. 945, 181 L. Ed. 2d 911 (2012).

This white paper describes a research agenda to advance our understanding of the problem and solution space of big data privacy in emerging market fintech and financial services, as well as related applications of big data in low- and middle-income countries (LMICs). Central to our proposal is that robust approaches to privacy will necessarily pair technical approaches alongside privacy-enhancing policies and practices.

We highlight five priority areas for research. Section 1 describes the need for more comprehensive landscape analysis. Section 2 highlights the lack of a basic understanding of what “privacy” means in non-Western contexts. Section 3 is broken into two parts, with Section 3.1 highlighting what we perceive to be five key sources of risk related to data privacy, and Section 3.2 outlining promising avenues for technical research to attend to these risks. Section 4 discusses non-technical approaches to data privacy, and how they can be used to complement the technical solutions described in the prior section. Section 5 provides a deep dive into one of the most pressing areas for focused research: the tradeoff between privacy and other common objectives, such as accuracy, consumer welfare, and profitability. Section 6 concludes with discussion of this research agenda’s relevance for consumers, financial institutions, and society.

1 Priority 1: Landscape analysis

It is only in the last few years that emerging markets have seen widespread use of large-scale personal data. And while the adoption of digital credit and other products has been rapidly on the rise – with hundreds of millions of active users across LMICs – we are unaware of any systematic and rigorous analysis describing the landscape of big data privacy in fintech and related emerging market applications. Developing this systematic landscape is important to understand (1) how personal data is currently being used in fintech and related applications; and (2) if and how the privacy of consumers and firms is currently being protected. We therefore propose, as a foundational step, a landscape analysis to identify major areas of opportunity and concern, and to highlight important gaps in technology, investment, policy, and practice.

While such an analysis is complicated by the historical opacity of the financial services industry, current approaches to big data privacy in the tech sector are relatively better publicized. For instance, it is increasingly common for large US-based tech firms to voluntarily disclose information

about their consumer privacy practices. Google (Erlingsson, Pihur and Korolova, 2014), Uber (Near, 2018; Johnson, Near and Song, 2018), Amazon,³ Meta,⁴ Microsoft,⁵ Apple (Tang et al., 2017), and LinkedIn (Kenthapadi, Chudhary and Ambler, 2017; Kenthapadi and Tran, 2018) all have adopted and publicized their use of differential privacy for some of their big data applications. Additionally, Google publicly states that it also protects data privacy using k -anonymity⁶ and federated learning,⁷ while Meta publicly discloses that it also uses secure multiparty computation⁸ and federated learning.⁹

Outside of the big tech firms that voluntarily disclose the types of technologies they use to protect individual data, not much is known about industry practices more broadly. At times, when firms poorly attend to individual privacy, a regulatory body (such as the Federal Trade Commission) steps in and reveals information surrounding the specific actions that led to privacy harms (Hoofnagle, 2016). But regulatory enforcement is far less common in emerging markets (Ademuyiwa and Adeniran, 2020).

Landscape Analysis Research Questions In order to understand the current landscape of big data privacy in financial services and related emerging market applications, three major research questions need to be addressed.

1. *Use cases:* What are the key existing and upcoming products, services, and applications that require personal data in emerging markets (in financial services and related settings)?
2. *Data access and control:* What data are being used, where do the data come from, and who controls and regulates those data? Are data commonly shared between institutions and organizations? If so, for what purposes and under what conditions?
3. *Existing protections:* What technical and non-technical approaches are currently being used to safeguard data privacy? How are those approaches typically adopted and/or adapted? And where do these approaches fall short?

³See <https://www.amazon.science/tag/differential-privacy>

⁴See <https://privacytech.fb.com/differential-privacy/>

⁵See <https://blogs.microsoft.com/ai-for-business/differential-privacy/>

⁶See <https://policies.google.com/technologies/anonymization?hl=en>

⁷See <https://ai.googleblog.com/2017/04/federated-learning-collaborative.html>

⁸See <https://privacytech.fb.com/multi-party-computation/>

⁹See <https://privacytech.fb.com/on-device-learning/>

2 Priority 2: Defining “privacy” in emerging markets

The concept of privacy has been described in a multitude of ways by legal scholars (Solove, 2005), philosophers (Nissenbaum, 2004), social scientists (Altman, 1977), and technical researchers (Dwork et al., 2006a). Conceptions of privacy differ not only between these fields but also within each of these fields. For example, legal discourse has described privacy as an individual’s “right to be let alone” (Brandeis and Warren, 1890), an individual’s ability to have “control over their information” (Westin, 1968), and as a “zone of refuge” for an individual (Sklansky, 2014).

The fact that privacy can be described in multiple ways does not mean that some notions of privacy are “wrong.” Nor does it mean that one notion of privacy is “more right” than another. Rather, privacy is best understood as a *pluralistic* (Koops et al., 2016), *social* (Post, 1989), *contextual* (Nissenbaum, 2004), and *contested* (Mulligan, Koopman and Doty, 2016) concept. It is plural in the sense that privacy can mean multiple things. It is social in that privacy is about individuals and groups of individuals. It is contextual in that actions that violate privacy in one setting may be appropriate in another.¹⁰ And it is contested in sense that disagreements about “what privacy is” are essential to provide robust privacy protections across different contexts. Taken together, these imply that there is not a single definition of privacy that can be applied in every setting.

This has two implications for *data privacy*. First, this means that there is no “silver bullet”: no single technology, law, or regulation can be used to address every single privacy issue stemming from the use of data (Mulligan, Koopman and Doty, 2016). Second, this means that the design of data privacy protections should be aimed at addressing privacy *for a specific setting*. Data collected from different contexts have different sensitivities, and thus require different notions of privacy; data privacy protections should be calibrated based on the particular setting in which they are being deployed.

Much of our current understanding of privacy is rooted in a Western perspective. For example, the above notion of privacy as a “right to be let alone” first appeared in a Harvard Law Review article in the 1890’s as a reaction to the rise of photography and printed media to capture and distribute images of individuals’ private lives (Brandeis and Warren, 1890). Similarly, the notion of of privacy as “control over information” arose in the US in the 1970’s. Fueled by concerns surrounding the

¹⁰For example, within the US, the warrantless wiretapping of an individual’s house by the police has been viewed as a privacy violation, but the wiretapping of an individual’s house for whom the police have secured a warrant for is not viewed as such.

adoption and expansion of government record-keeping systems (Ware, 1973), as well as inappropriate government surveillance activities, including the Watergate scandal and the wiretapping of Martin Luther King Jr (Bajpai and Weber, 2017), the notion of privacy as “control over information” sought to mitigate the harms posed by secretive surveillance systems. In addition to being US-centric, these notions of privacy were articulated to deal with the specific social problems of their time.

However, the most salient privacy concerns in emerging markets are not necessarily the same as those that are most important in wealthy nations. For instance, Abebe et al. (2021) have documented that some African communities view the sharing of aggregate statistics about their community with individuals outside their community as a privacy violation, regardless of the technical protections used to secure the information. This conception differs from the US, where the public disclosure of statistics is common and at times required by law (Abowd, 2018).

Differences in privacy concerns are not exclusively confined to data sharing practices. Rather, these differences can permeate throughout everyday interactions. For example, within the US it is common to present ID when purchasing over-the-counter medication at a pharmacy. However, in Saudi Arabia this practice is not common except in special circumstances, so the requirement to disclose your identity to purchase over-the-counter medication can be viewed as an intrusive solicitation for too many details about the consumer (Abokhodair and Vieweg, 2016), resulting in a privacy violation. More generally, on-the-ground concerns about privacy in Lagos and São Paulo may fundamentally diverge from those found on the ground in Los Angeles and San Francisco.

Privacy Definition Research Questions Given that much of our understanding of data privacy – and privacy more generally – is rooted in the US context, there is a need to document how privacy is conceptualized *locally* by individuals whose privacy is at risk in these emerging markets. A nuanced and comprehensive understanding of data privacy concerns is a necessary precursor to developing technical and non-technical solutions to address those concerns. Failing to do this work risks developing solutions that do not actually address the concerns of people in their everyday lives. Examples of outstanding questions include:

1. *Definitions:* How do specific communities and constituents conceptualize data privacy? How do consumers conceptualize the privacy harms that could arise from the use of personal data

sources for financial decisions? How much do these definitions and concerns vary from one region to another? What commonalities exist?

2. *Methodology, Conceptualization, and Measurement*: How should researchers engage local communities in conversations surrounding privacy in a meaningful way? Once articulated, how can privacy concerns best be captured and characterized? To what extent do Silicon Valley approaches to quantifying privacy (or the loss thereof) adequately address the local concerns, norms, and expectations of privacy in emerging market contexts?
3. *Tradeoffs*: When centered in local perspectives, what is the relative value of data privacy, vis-à-vis other common tradeoffs (e.g., product innovation, availability, and cost) in different local contexts?

3 Priority 3: Technical Solutions to Data Privacy

3.1 Background and Motivation: Privacy Threat Analysis

Data privacy is typically described in relation to privacy harms that arise from the (i) collection, (ii) storage, (iii) processing, or (iv) dissemination of information (Solove, 2005). While this 4-part categorization points to *when* failures of data privacy can occur when data are used, it provides an incomplete view of privacy harms. For example, it does not tell us *who* infringed upon privacy, *what* this infringement looks like from a technical perspective and for the individual, *where* the vulnerability and threat came from, and *why* and *how* it occurred. In other words, this 4-part categorization under-specifies the “who-what-when-where-why-hows” of actions that violated privacy.

Since, as discussed in Section 2, privacy is inherently a social concept, the analysis of privacy should center around the individuals and groups whose privacy is at stake. This motivates a clear understanding of *how* individuals can be harmed by their data. Broadly, we view data privacy harms (that occur after data are collected¹¹) as a two-step process: data *reveals something* about an individual, and from these revelations *something occurs* that negatively impacts an individual. Examples of the latter include the disclosure of information (such as doxing) or the threat of disclosure

¹¹Our focus on data privacy harms that occur post-collection does not imply that the collection of data can never constitute a privacy violation. Rather, we focus on post-collection data privacy harms in this section since we will be documenting five different ways in which harmful information can be inferred from personal data that has already been collected.

of information (such as blackmail), financial harm, and physical harm (Solove, 2005), as well as reputational harms and changes in the opinion of the individual whose information was revealed.

How can potentially harmful information be revealed about individuals from their data? We organize our discussion according to five different types of revelations: (1) revelations from one’s own data; (2) revelations from another individual’s data; (3) revelations from the social meaning of the dataset; (4) re-identification of individuals from “anonymized” data; and (5) reconstruction of data from statistics. We explain each of these 5 vulnerabilities in turn below.

3.1.1 Privacy threats from an individual’s own data

An individual’s data can reveal much more about them than first meets the eye. As a motivating example, consider a dataset that contains the locations an individual visited over the past month. Such datasets can be useful for a wide range of tasks, ranging from city planning purposes (Bassolas et al., 2019) to humanitarian response (Kohli, Aiken and Blumenstock, 2023).

While these locations may “only” represent geographical coordinates on a map, the social meaning attributed to the physical locations can be used to draw inferences about an individual’s life, including their sexual preferences (Ovide, 2021; Boorstein, Iati and Shin, 2021) and political preferences (Thompson and Warzel, 2019). For example, a visit to a place of worship evokes different social perceptions about the individual than a visit to a gentleman’s club, which evokes different perceptions than a visit to a drug or alcohol rehabilitation center.

The process of generating these inferences by leveraging the social meaning of data are called *social meaning attacks*, and they can be notoriously hard to defend against. When multiple datapoints are collected about an individual, social meaning attacks can paint an even more detailed picture about an individual’s life (Ohm, 2009). For instance, a visit to an OB/GYN followed a week later by a visit to an expecting mother’s clothing store describes more about an individual’s life than a visit to each of these locations in isolation.¹²

Takeaway: While the literal representation of data may seem innocuous, the information that the data reveals can be far more noxious for individual privacy, particularly when combined with the social

¹²This is inspired by an example provided in *US v. Maynard*, 615 F.3d 544 (D.C. Cir. 2010), where the DC Circuit Court remarked that the “sequence of a person’s movements may reveal more than the individual movements of which it is composed.”

meaning of data.

3.1.2 Privacy threats from other people’s data

Inferences based on similarly situated individuals can also lead to revelations that are noxious for individual privacy. This type of *predictive privacy harm* (Crawford and Schultz, 2014b) arises in situations where statistical associations reveal private information that an individual has not disclosed.

Examples of this abound. In 2009, researchers showed that it was possible to infer a Facebook user’s sexual orientation based on their social network (Jernigan and Mistree, 2009). For individuals who do not choose to disclose their sexual orientation on Facebook, the use of predictive technologies to infer this sensitive fact creates unique privacy challenges. And in 2012, Target discovered that a teenage girl was pregnant based on the similarity between her shopping patterns and those of other pregnant women, despite her never revealing to Target that she was pregnant (Hill, 2012).

Takeaway: Big data technologies can leverage statistical associations to infer non-disclosed attributes, and reveal sensitive information that individuals may have intentionally withheld.

3.1.3 Privacy threats from the social meaning of a dataset

In Section 3.1.1, we described how social meaning attacks can reveal information about an individual beyond what is explicitly stated in their data. One specific type of social meaning attack occurs when the “type of database” an individual is associated with reveals information about them.

One such example is the Ashley Madison data breach. For context, Ashley Madison is an online dating site for married individuals who are looking to have an affair. In 2015, hackers published troves of the company’s data online, including their members’ information (Mansfield-Devine, 2015). This included each member’s name, email address, relationship status, gender, birthdate, billing address, and debit / credit card information.¹³ This online repository of user details would eventually come to be known as the infamous Ashley Madison dataset.

While the release of this hacked information constitutes a privacy breach in its own right, the social meaning of the Ashley Madison dataset automatically reveals even more salacious information

¹³See Paragraph 37 in the Federal Trade Commission’s complaint against Ashley Madison, available [here](#).

about each individual in the dataset. Since Ashley Madison was a website for married individuals seeking an extramarital affair, the existence of a user profile in the Ashley Madison dataset implies that the individual was seeking (or had) an affair. That is, the social meaning of the dataset reveals a characteristic about the individuals whose data is housed within the dataset.

This is not an isolated incident. [Homer et al. \(2008\)](#) were able to determine the involvement of an individual in a Genome Wide Association Study based on aggregate statistics.¹⁴ Since all the individuals in this study shared a common disease, the discovery of an individual in this dataset reveals that the individual has the disease as well.

Takeaway: The specific values of an individual’s data record are not required to reveal a sensitive fact about them. For datasets with social stigma or particular sensitivities about them, simply being identified as part of such a dataset can reveal sensitive information about an individual.

3.1.4 Privacy threats from “anonymized” datasets

When data are used for computational purposes, they are typically stripped of overtly identifying information, such as an individual’s name, social security number, phone number, etc. Due to the lack of an explicit identifier, such data are often assumed to be “anonymized.”

However, the absence of an identifier within a dataset does not preclude the remaining information in the dataset from de facto identifying an individual ([Narayanan and Felten, 2014](#); [Kroll, Kohli and Laskowski, 2019](#)). For example, in the late 1990’s, Latanya Sweeney was able to re-identify the then governor of Massachusetts, William Weld, in a supposedly “anonymized” medical dataset by linking his birthdate, zip code, and sex to a voter registration database that contained his name ([Sweeney, 2002](#)). Similar methods have been successfully re-identified individuals in “anonymized” datasets in different domains, including Netflix data ([Narayanan and Shmatikov, 2008](#)), HIPAA-compliant hospital data ([Yoo et al., 2018](#)), and clinical drug study data ([Branson et al., 2020](#)). For location data in particular, [De Montjoye et al. \(2018\)](#) famously showed that four spatio-temporal datapoints were sufficient to uniquely identify 95% of individuals in an “anonymized” mobile phone dataset.

The method of linking fields in an “anonymized” record to another data source is known as a

¹⁴The technical process by which this involvement was determined is example of a *membership inference attack* ([Shokri et al., 2017](#)), also known as a *tracing attack* ([Dwork et al., 2017](#)). We will discuss these attacks in more detail in Section 3.1.5.

linkage attack, and is an example of a *re-identification attack* more broadly (Sweeney, 2002). The goal of a re-identification attack is to attach the identity of an individual to a data record without one. That is, re-identification attacks reveal the identity of an individual despite the absence of an “obvious” identifier within the data.

Takeaway: Datasets that have been stripped of explicitly identifying information are not immune to privacy risks. There may still be sufficient information in an “anonymized” dataset to reveal the identity of individuals.

3.1.5 Privacy threats from aggregate statistics and big data computations

The previous four harms stem from the possibility of inferring information about an individual beyond what their data explicitly represented. Just as individual data can reveal more than first meets the eye, aggregate statistics – i.e., statistics computed on multiple individual’s data – can reveal quite a great deal about each individual as well.

Statistics are numerical summaries of a dataset. Statistics take datasets – which can contain hundreds, thousands, and even millions of numbers – and compress them into digestible parcels of information. However, since statistics summarize data, they necessarily contain information about the data. As such, statistics may inadvertently reveal the underlying elements of the dataset.

The seminal work of Dinur and Nissim (2003) demonstrated that collections of statistics could be analyzed to reverse engineer the input dataset that generated them. For settings where statistics are released because the sharing of individual data is forbidden, such *reconstruction attacks* contradict the purpose of sharing statistics in the first place. The work of Dinur and Nissim, and the subsequent papers that their work inspired, led to what is known as the *Fundamental Law of Information Recovery*: “too many overly accurate” statistics necessarily violate the privacy of individuals whose data contributed to the statistics (Dwork, Kohli and Mulligan, 2019). As such, “sufficiently noisy” statistics are required to thwart reconstruction attacks (Dwork et al., 2017).¹⁵ Reconstruction attacks aren’t just a theoretical possibility. For example, internal testing by the US Census uncovered that the 2010 census statistics were susceptible to reconstruction attacks, violating the privacy of 46% of the US population.¹⁶

¹⁵We discuss how to use differential privacy to construct “sufficiently noisy” statistics in Section 3.2.

¹⁶See Amicus Brief of Data Privacy Experts in *Alabama v. US Department of Commerce et al.*, Document 99-1 Filed

Classic summary statistics are not the only information sources vulnerable to reconstruction attacks. Since the parameters of machine learning models are statistics computed on training data, reconstruction attacks can also be levied against machine learning models to reconstruct the training data (Balle, Cherubin and Hayes, 2022). Machine learning models are also vulnerable to *membership inference attacks*,¹⁷ which aim to determine whether an individual datapoint is in the model’s training dataset (Shokri et al., 2017). The key insight that enables membership inference attacks is that machine learning models – particularly those that are suffering from overfitting – tend to be more confident in their predictions on datapoints in their training data (Jarin and Eshete, 2022). Consequently, an adversary can query the model with input data, observe the model’s confidence on the input, and infer whether the datapoint was likely used to train the model.¹⁸

Once individual data records have been reconstructed or inferred to be part of a training dataset, all four of the prior revelatory threats are possible: re-identification attacks can be levied to attach an identity to the data record; inferences can be drawn about the individual from their own record; inferences can be drawn about the individual based on characteristics of similar individuals; and, depending on the dataset in which these records belong to, the social meaning of the dataset may reveal further information about the individual. As we will see in Section 3.2, differential privacy can be used to provably thwart reconstruction attacks and membership inference attacks (Dwork et al., 2017).

Takeaway: “Too many overly accurate” statistics (and other big data computations, such as machine learning models) are vulnerable to reconstruction attacks and membership inference attacks. These attacks allow an adversary to determine what data was used to create a statistic or machine learning model, which can then be used as a starting point to reveal even more information about individuals.

3.2 Overview of Technical Solutions

Technical solutions to data privacy vary based on the specific threat to privacy they seek to mitigate. In general, current technical approaches to computational data privacy seek to protect the underlying

04/23/21, Case 3:21-cv-00211-RAH-ECM-KCN.

¹⁷Classic statistics are also vulnerable to these sorts of attacks. Within the privacy literature, membership inference attacks are also called *tracing attacks* (Dwork et al., 2017).

¹⁸Membership inference attacks, while noxious for individual privacy, can sometimes be used as a tool to audit machine learning algorithms. For instance, in a recent high profile example, Chang et al. (2023) performed a membership inference attack on ChatGPT and GPT-4 to determine which books these large language models were (likely) trained on.

data from being “exposed” by a computation.¹⁹ There are many ways this exposure occurs. This section explores these different exposures and the current technical solutions that are commonly used to address them.

3.2.1 Differential Privacy: For Statistics (and Other Big Data Applications)

In Section 3.1, we discussed how overly accurate statistics and machine learning models are susceptible to reconstruction attacks and membership inference attacks. Fortunately, a mathematical approach to privacy known as *differential privacy* can be used to thwart both of these threats (Dwork et al., 2017).

Dwork et al. (2006a) introduced the concept of ϵ -indistinguishability, which would subsequently become known as ϵ -differential privacy, to protect statistics from exposing the underlying data they were computed on. Differential privacy provides data subjects with a mathematical promise, via a formal proof, that guarantees “very little” will be revealed about them when their data is analyzed (Dwork, Kohli and Mulligan, 2019). In doing so, differential privacy enables statistical learning from datasets while protecting the privacy of individuals whose data was used to generate the statistics.

The rationale for constructing a computation that reveals “very little” about any individual is as follows: if a statistic reveals “very little” about any individual, then any adversary who examines the statistic can only learn “very little” about the individuals whose data created the statistic. Hence, data subjects incur “very little” privacy loss when their data are analyzed. This notion of “very little” is made mathematically precise with the introduction of a privacy parameter ϵ (read as “epsilon”). As such, smaller values of ϵ correspond to smaller privacy losses, and larger values of ϵ correspond to larger privacy losses.

In order to achieve this mathematical promise, differentially private algorithms introduce noise when computing statistics.²⁰ That is, differentially private algorithms release “noisy statistics” instead

¹⁹Here, we focus on threats to data privacy that stem from computing statistics on data. In particular, we use the term *statistics* to refer to the output of any computation run on multiple individual’s data. Such computations include the calculation of elementary summary statistics (such as the mean, median, and mode) as well as the execution of complicated machine learning training algorithms (that compute the parameters of machine learning models). Additionally, we will refer to data privacy concerns that arise from the computation of statistics as *computational data privacy*. We do so to delineate it from other data privacy concerns, such as the unauthorized sharing of data to third parties, that are important yet distinct from computational concerns.

²⁰The presence of any noise in a statistical calculation does not automatically make the statistic differentially private. The noise that differential privacy requires is very particular – it must be carefully calibrated and come from appropriate noise distributions (Dwork et al., 2006a).

of true statistics in order limit the statistic’s ability to expose individual data. The amount of noise introduced is related to the privacy parameter ϵ . While smaller values of ϵ yield more privacy, they introduce more noise into the statistic, yielding less accurate results. Alternatively, larger values of ϵ require the introduction of less of noise, which provides more accurate results at the expense of weaker privacy protections. For this reason, differential privacy induces a *privacy-accuracy* tradeoff, which is controlled by the parameter ϵ . The value ϵ is a tunable by the algorithm developer, which allows for different tolerances for the privacy-accuracy tradeoff to be achieved in practice. Since different datasets may have different sensitivities about them, and since different statistics may require different levels of accuracy to be useful, there is no universal consensus on how to set ϵ (Dwork, Kohli and Mulligan, 2019). However, several different methods exist to identify appropriate values of ϵ based on context-specific goals (Hsu et al., 2014; Kohli and Laskowski, 2018; Dwork, Kohli and Mulligan, 2019; Kohli, Aiken and Blumenstock, 2023).

There are many different algorithms that one can use to achieve differential privacy, such as the Laplace mechanism (Dwork et al., 2006a), Gaussian mechanism (Dwork et al., 2006b), multiplicative weights mechanism (Hardt and Rothblum, 2010), and general mechanisms for arbitrary statistical analyses (Mohan et al., 2012; Kohli and Laskowski, 2023). Since the parameters of a machine learning model are statistics computed from training data, differentially private algorithms can be leveraged to produce differentially private machine learning models (Ji, Lipton and Elkan, 2014; Abadi et al., 2016; Papernot et al., 2018; Ha et al., 2019). We note, however, that differential privacy is often incorrectly described as a process by which noise is added to the data to protect privacy. While some differentially private algorithms may perturb individual data,²¹ other algorithms do not. In fact, many algorithms that satisfy differential privacy perturb the output of a statistic while leaving the original data intact.

Differential privacy is accompanied with a whole host of formal guarantees that make it useful for privacy engineering. First, differential privacy makes it possible to keep track of the cumulative privacy loss accrued over multiple analyses. For instance, suppose an analyst runs an ϵ_1 -differentially private algorithm on a dataset and then runs an ϵ_2 -differentially private algorithm on the same data. In the first analysis, data subjects have incurred ϵ_1 privacy loss, whereas in the second analysis, data subjects have incurred ϵ_2 privacy loss. So in total, data subjects have incurred is $\epsilon_1 + \epsilon_2$ privacy loss. This *composition result* informs us that, when multiple differentially private analyses are run

²¹A variant of differential privacy, known as *local differential privacy*, can be used to perturb user information. We will elaborate more on this in Section 3.2.5.

on a dataset, the privacy loss is cumulative. For this reason, differential privacy enables rigorous accounting of privacy loss.

The composition result tells us what happens when multiple private analyses are performed on the same dataset. Another property of differential privacy, known as the *post-processing result*, tells us what happens when someone analyzes a differentially private statistic: without access to the original data, data subjects do not incur additional privacy loss when the output of differentially private computation is analysed. That is, differential privacy fully inoculates a statistic in such a manner to ensure that no matter what anyone else uses this statistic for, data subjects will incur no further privacy loss. Fascinatingly, this result holds *regardless* of (1) the auxiliary information an adversary may have, (2) the specific attack strategy they employ, and (3) the computing power and technology they have – now or in the future. For this reason, differential privacy is often said to be *future proof and adversary agnostic* (Dwork, Kohli and Mulligan, 2019).

3.2.2 Differential Privacy: For Synthetic Data

While differential privacy was introduced as a means to compute privacy-preserving statistics, techniques from differential privacy have been used to generate privacy-preserving synthetic datasets. Conceptually, a synthetic dataset is a “fake dataset” that can be used in place of a real dataset for computations. To ensure that synthetic datasets are representative of the real dataset they seek to replace, synthetic datasets are constructed to replicate numerous statistical properties of the real dataset.

By themselves, synthetic datasets have no formal privacy guarantee. Since statistics can expose underlying elements of a dataset, synthetic data built from these statistics can similarly expose elements of the real dataset (Giomi et al., 2022).²² However, when a synthetic dataset is constructed using techniques from differential privacy, the dataset inherits all of the protections that differential privacy has to offer. Most relevant for differentially private synthetic data is the post-processing guarantee: any use of a synthetic dataset will not incur additional privacy loss for data subjects. This is intuitively appealing for many data analysis tasks, as differentially private synthetic data would enable anyone to analyze the synthetic data without compromising the privacy of the individuals in the real dataset.

²²For a gentle and conceptual introduction to synthetic data, see Ullman (2022).

In their seminal paper, [Blum, Ligett and Roth \(2008\)](#) showed that it is possible to construct a differentially private synthetic dataset that preserves a large number of statistical characteristics of the original data. However, their method is computationally expensive, making it impractical to employ in practice (except on small datasets). Currently, the creation of differentially private synthetic data is an active and promising area of research ([Zhang et al., 2021](#); [Aydore et al., 2021](#); [Tantipongpipat et al., 2021](#); [Bun et al., 2023](#)).

3.2.3 Homomorphic Encryption

Differential privacy aims to protect a statistic from exposing elements of the underlying dataset. In this sense, differential privacy seeks to protect the data *post-computation* by inoculating the statistic with carefully calibrated noise. In order to protect data from being exposed *during the computation process*, other privacy paradigms are required. One such paradigm is *homomorphic encryption*.²³

The concept of homomorphic encryption was introduced by [Rivest, Adleman and Dertouzos \(1978\)](#). The goal of homomorphic encryption is to design an encryption protocol in such a manner that permits computations over encrypted data without the need to decrypt it. Homomorphic encryption derives its name from these two desiderata: *homo* (meaning same) and *morphic* (meaning structure). By design, homomorphic encryption ensures that the mathematical structure of computations over encrypted data is the same as in the unencrypted state.

The first fully homomorphic scheme – a scheme that allows for arbitrary computations – was introduced by [Gentry \(2009\)](#). This scheme, however, was computationally inefficient to implement. Since then, other homomorphic encryption schemes have been introduced that are more efficient, but restrict the types of computations that can be performed ([Evans, Kolesnikov and Rosulek, 2017](#)). Despite these restrictions, schemes exist that allow for addition, subtraction, and multiplication, which can be applied in many practical settings involving statistics and machine learning. One prominent example was the 2016 advent of *CryptoNets*, which allows for the training of neural networks on homomorphically encrypted data ([Gilad-Bachrach et al., 2016](#)).

While both homomorphic encryption and differential privacy operate in the statistical setting, they address different privacy concerns. Differential privacy is concerned with attacks on statistics

²³For a more technical introduction to homomorphic encryption, see [Halevi \(2017\)](#).

that can be used to reveal information about individuals in the underlying dataset post-computation. Homomorphic encryption, on the other hand, is concerned with attacks that can occur when data is in an unencrypted form during computation. Both are concerned about protecting the underlying data – however, they provide protection against different threats that arise at different points in the computation process.

3.2.4 Secure Multiparty Computation

So far we have considered computational settings where data resides in a central location. In some situations, however, data is not able to be housed in a single location. This can occur for a myriad of reasons, such as a lack of trust between multiple data holders, as well as legal restrictions that outright forbid the sharing of data.

In situations such as these, *decentralized computing* – also known as *distributed computing* – can be used to calculate statistics from data without the need to store the data in a central location. While the lack of a single central repository of data can alleviate data privacy concerns surrounding the storage of information, decentralization of data alone does not yield any formal privacy guarantees. However, when combined with other privacy paradigms, formal privacy guarantees can be established.

One such paradigm is *secure multiparty computation*. Secure multiparty computation, often referred to as multiparty computation (MPC), was introduced by Yao (1982) to enable the computation of statistics that rely on multiple parties' data, such that each party learns nothing more about another party's data other than what is revealed by the statistic. MPC is established through protocols that specify the way in which parties act so as to properly compute information in such a manner that, when aggregated together, the result is correct and does not expose any party's input data.

The first protocol was Yao's garbled circuit protocol (GPC), which can be used to compute an arbitrary discrete function (provided that this function can be represented as a fixed-size circuit) (Evans, Kolesnikov and Rosulek, 2017). To date, Yao's GPC remains a foundational building-block for many MPC applications. MPC applications may also use other protocols, such as those involving secret sharing and homomorphic encryption.

Recently, researchers partnered with two Dutch banks to perform privacy-preserving money laun-

dering detection using MPC (Sangers et al., 2019; van Egmond and Rooijackers, 2021). Since data sharing between financial institutions raises privacy concerns, MPC enabled the banks to collaboratively compute over each other’s data to detect money laundering without the need to ever share consumer data.

3.2.5 Local Differential Privacy

In similar spirit to decentralization, a variant of differential privacy, known as *local differential privacy* (LDP), was introduced to compute privacy-preserving statistics in settings where personal data needs to remain private even from the data collector. In our prior discussion of differential privacy, a central authority was tasked with computing noisy statistics from centrally located data. In this case, noise is introduced into the statistic *after* the central authority receives the data. As such, this classic setting is often called the *central* differential privacy setting.

In the LDP setting, a central authority is still tasked with computing noisy statistics. However, in this setting noise is introduced into information *before* the central authority receives it. This is appealing at an intuitive level, as the central authority never receives real personal information – they only receive noisy personal information.

To elucidate discussion, consider a financial institution that wants to understand the reliability of their mobile banking app. The financial institution will assess this by examining the total number of app crashes across their entire user base. Since the company is interested in the total number of crashes, the company could request a noisy count of crashes from each subscriber’s phone. Summing up the noisy individual estimates would provide a noisy estimate of the total number of app crashes across the company’s user base without ever requiring the transfer of crash logs that could contain sensitive mobile phone metadata. As was the case with (central) differential privacy, a wide variety of mechanisms exist to achieve local differential privacy (Erlingsson, Pihur and Korolova, 2014; Kairouz, Oh and Viswanath, 2016; Bebensee, 2019)

3.2.6 Federated Learning

Similar in spirit to MPC is *federated learning* (FL), which enables the decentralized development of machine learning models (McMahan et al., 2017). Following on the theme of decentralization, the

idea of FL is to develop a single machine learning model based on n individuals data without the need to consolidate individual data prior to the model’s development.

At a high level, FL enables the local training of models in such a manner that the local model updates can be combined to a global model update. Federated learning presents an opportunity to train machine learning models from individual data without the need for personal data transfer. Instead, local model updates are transferred for global model development.

However, as noted above, decentralization alone does not provide a rigorous privacy guarantee. Absent any additional requirements, local model updates can leak information about the local data. When coupled with mathematically rigorous techniques from MPC and LDP, FL can now be done in a manner with strong privacy guarantees. For example, [Truex et al. \(2019\)](#) showed that a wide array of machine learning models (such as decision trees, neural networks, and support vector machines) could be trained in a federated manner in conjunction with techniques from MPC and LDP.

3.2.7 Technical Privacy Solutions for LMICs: What Research is Needed?

Techniques from differential privacy (in both the central and local settings), homomorphic encryption, secure multiparty computation, and federated learning can be used to address different privacy concerns related to the exposure of sensitive information. Recently, promising research has emerged that combines multiple techniques from these privacy paradigms to afford individuals stronger levels of privacy ([Wagh et al., 2021](#)). This remains an active and exciting area of research within computer science and computer security.

‘Technical Solution’ Research Questions However, the vast majority of these innovations and advances have been developed for use cases and markets in developed economies – typically by companies operating in those markets or by researchers receiving support from companies or funding agencies from those markets. By contrast, very little – if any – research has been done on privacy-preserving big data applications in emerging markets. For billions of users, this leaves many important questions unanswered:

1. Which, if any, of the existing privacy-enhancing technologies could be used “off the shelf” in emerging market applications? Are there particular settings where minimal tinkering is required?

2. How can these approaches be used in conjunction with one another to mitigate a wider family of privacy harms? Would this more adequately address the key privacy concerns in emerging markets?
3. Are there particular settings where existing technical solutions fail to address the underlying privacy concerns that are most important in that setting?
4. What new technical innovations can best protect the data privacy concerns that are unique to emerging markets in general, or specific local contexts?

4 Priority 4: Non-technical Approaches to Data Privacy

In addition to understanding the potential and limitations of technical solutions to data privacy in emerging markets, *non-technical* approaches – such as policies and practices – can also be very effective at proactively addressing privacy concerns. For instance, in the US context, government organizations are frequently bound by written laws and policies that dictate appropriate uses of data (Bushkin and Schaen, 1976; Officer and Keefer, 2009). In Europe, privacy restrictions tend to be more stringent. For instance, the General Data Protection Regulation (GDPR) places legal restrictions on the collection, storage, analysis, and sharing of personal data. And even when the use of personal data is admissible, the GDPR requires compliance with seven data protection and accountability principles. Failure to abide by the GDPR can result in a maximum fine of 4% of global revenue or €20 million, whichever is larger.²⁴ These stringent fines aim to incentivize compliance with the GDPR, and hence incentivize proactive data privacy practices.

In the private sector, corporations also frequently adopt data use guidelines that enable employees and products to better attend to privacy (Bamberger and Mulligan, 2015). While these policies vary across organizations, sector-specific privacy and data protection laws provide baseline requirements on how and when information may be collected, stored, protected, accessed, used, and deleted. Common practices include data minimization, use limitation, sharing restrictions and requirements, and de-identification (Sedberg and Mulligan, 2015). While these policies lay out the rules of the road, and while these practices help organizations stay in the correct lane, additional accountability mechanisms, such as access controls and audit logging, can also be used to detect deviations from these policies

²⁴See <https://gdpr.eu/what-is-gdpr/>.

and practices (Bowman et al., 2015).

However, to ensure that policies and practices are well-calibrated to the setting at hand, these policies and practices must build off of local conceptions of privacy. As noted earlier in Section 2, several prominent researchers have argued that the “Silicon Valley paradigm” of data privacy is inappropriate in emerging market settings (Abebe et al., 2021), meaning that many data sharing practices that are common in the US may be viewed as privacy violations in the Global South.

‘Non-Technical Solution’ Research Questions Since privacy is a social concept, actions that violate norms of expected uses of data are often viewed as privacy violations by data subjects (Nissenbaum, 2004). Consequently, there is a need for qualitative research to understand *what* particular uses of data are viewed as appropriate by a community within a given setting, and what uses of data are viewed as inappropriate as well. In order to understand potential mismatches in current data privacy policies and practices, several questions require deeper investigation.

1. How does current or proposed regulation address potential privacy harms and consumers’ expectations of privacy? How does it fall short? To what extent is regulation enforced, and to what extent does it more broadly affect product innovation?
2. How can financial service providers best engage proactively with customers – by providing information, requesting consent, etc. – to better align customer expectations with current or future practice?
3. To what extent do existing practices align with consumer’s expectations of privacy? Where do these practices miss the mark? Is it with respect to information collection, access, use, sharing, or deletion? And what new practices can be implemented to build and foster consumers’ trust that their local notions of privacy are being respected?

5 Priority 5: Understanding the Real-World Tradeoffs Between Privacy and Other Key Priorities

Strong privacy guarantees do not come for free. For example, differentially private machine learning algorithms frequently induce a tradeoff between privacy and accuracy: the more private the system,

the less accurate the algorithm. When those algorithms are used to drive key decisions – such as credit scoring and risk management – the loss in accuracy can have a direct consequence on the bottom line (see also the discussion in Section 3.2).

In some settings, quantitative analysis can mathematically demonstrate the exact nature of the privacy-accuracy tradeoff for a given context (Kohli, Aiken and Blumenstock, 2023). For instance, it may be possible to determine that an $X\%$ increase in privacy would translate to a $Y\%$ decrease in accuracy (and thus a $Z\%$ loss in profits). In the context of digital credit, for example, it may be possible to construct a credit score based on a private version of the borrower’s personal data – but that credit score would be less accurate than a version constructed from the raw data, and as such, might not be as profitable for the financial institution. In other cases, lower and upper bounds can be derived to describe the range of the privacy-accuracy tradeoff. But this type of rigorous analysis is rarely documented even in places like the US; to our knowledge, no such case studies exist from emerging markets.²⁵

The accuracy of statistics and machine learning models can be tied to many different objectives depending on what they are used for. When statistics are used for business decisions, accurate statistics promote profitability. Other times, statistics may be used to inform public health interventions and humanitarian response. In these settings, accurate statistics promote social welfare. Alternatively, US Census statistics are used by the government for a whole host of tasks, such as providing election information in minority languages and allocation educational funds (Pujol et al., 2020). In this setting, accurate statistics promote dignity, representation, and fairness in governance.

Improvements in privacy often involve a wider range of tradeoffs beyond the privacy-accuracy tradeoff that is salient with differential privacy (Kohli, 2021). For instance, recent machine learning research has documented a frequent tradeoff between privacy and *fairness*: in settings where fairness criteria are not explicitly built into the learning process, differentially private algorithms can exacerbate inequities (Bagdasaryan, Poursaeed and Shmatikov, 2019). In other settings, there may exist similarly pointed tradeoffs between privacy and other social or corporate objectives such as transparency (more transparent decision rules may more visibility into user data), social welfare (maximizing social welfare may require accessing private information), and competition (monopolistic

²⁵One exception is our recent work in Kohli, Aiken and Blumenstock (2023), available [here](#), that characterizes the privacy-accuracy tradeoff of using differentially private mobility matrices in humanitarian response.

behavior could be inhibited by data sharing requirements), to name just a few examples.

To complicate matters further, tradeoffs are often multidimensional (Rolf et al., 2020). For example, a lending institution may want to preserve privacy and ensure that their lending algorithm does not discriminate on protected attributes, but they may also want to ensure that the algorithm still promotes profitability. When multiple objectives are considered, striking a balance between effective privacy protection and other objectives is a non-trivial and highly contextual task.

“Tradeoffs” Research Questions If and how these privacy tradeoffs relate to emerging markets is very poorly understood, leaving many basic questions unanswered:

1. What are the theoretical and empirical tradeoffs between privacy and accuracy in real-world use cases (such as with alternative credit scores)? How do the reductions in accuracy translate to other key corporate objectives (like risk and profits)?
2. In addition to accuracy, does data privacy induce tradeoffs with other key performance indicators? How can those tradeoffs best be documented and quantified?
3. To what extent do different notions of privacy – and in particular those that are most relevant in emerging markets – make the tradeoff between privacy and accuracy more or less stark?
4. Are financial service providers willing to make small sacrifices in order to protect consumer privacy? If so, under what circumstances?

6 Conclusion

To enable continued use and innovation of personal big data in financial services and related emerging market applications, we believe it is critical to launch a coordinated research agenda to develop a more complete understanding of the problem and solution space of big data privacy in low- and middle-income countries. Too little is known about how to most effectively protect consumers in these markets, or even of what the biggest threats to consumer privacy might be. This lack of knowledge creates systemic risk that threatens to undermine service innovation and adoption. We believe that proactively identifying and addressing potential threats and opportunities *now* — rather than in

reaction to a data privacy scandal — is the best way to ensure the health and welfare of consumers as well as the fintech industry.

To seed this research agenda, this paper provides an overview of modern approaches to data privacy, and identifies five major themes for understanding how best to adapt and improve those methods for use in emerging markets. These themes include:

1. A landscape analysis of key use cases and current privacy protections in financial services and related emerging market applications, with a focus on identifying key vulnerabilities, current best practices, and gaps in knowledge and technology.
2. Defining “data privacy” in emerging markets, with a focus on understanding if and how the concerns coming from low-income countries differ from those found in wealthy nations.
3. Expanding and adapting technical solutions to data privacy for use in emerging markets. Key technologies include differential privacy, homomorphic encryption, secure multiparty computation, and federated learning. To date, however, these technologies have been almost exclusively developed and deployed in high-income countries.
4. Developing and improving non-technical approaches to data privacy (such as policies and practices) that complement technical solutions.
5. Understanding the tradeoffs incurred by financial institutions, fintechs, and individuals when privacy-enhancing solutions are deployed in practice.

The increasing ubiquity of digital technology in emerging markets has the ability to revolutionize fintech and financial service provision in low- and middle-income countries. However, inadequate consideration for data privacy has the potential to undermine this innovation. This need not be the case. Appropriately leveraged, privacy-enhancing technologies, policies, and practices have the ability to simultaneously enable innovative uses of big data while also providing robust and meaningful privacy protections for personal data. We believe this research agenda can help catalyze a more responsible and consumer-oriented approach to big data innovation in financial services.

References

- Abadi, Martin, Andy Chu, Ian Goodfellow, H Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang.** 2016. “Deep learning with differential privacy.” 308–318.
- Abebe, Rediet, Kehinde Aruleba, Abeba Birhane, Sara Kingsley, George Obaido, Sekou L Remy, and Swathi Sadagopan.** 2021. “Narratives and counternarratives on data sharing in Africa.” 329–341.
- Abokhodair, Norah, and Sarah Vieweg.** 2016. “Privacy & social media in the context of the Arab Gulf.” 672–683.
- Abowd, John M.** 2018. “The US Census Bureau adopts differential privacy.” 2867–2867.
- Ademuyiwa, Idris, and Adedeji Adeniran.** 2020. “Assessing Digitalization and Data Governance Issues in Africa.” Centre for International Governance Innovation.
- Aiken, Emily, Suzanne Bellue, Dean Karlan, Chris Udry, and Joshua E Blumenstock.** 2022. “Machine learning and phone data can improve targeting of humanitarian aid.” *Nature*, 603(7903): 864–870.
- Altman, Irwin.** 1977. “Privacy regulation: Culturally universal or culturally specific?” *Journal of social issues*, 33(3): 66–84.
- Aydore, Sergul, William Brown, Michael Kearns, Krishnaram Kenthapadi, Luca Melis, Aaron Roth, and Ankit Siva.** 2021. “Differentially private query release through adaptive projection.” *arXiv preprint arXiv:2103.06641*.
- Bagdasaryan, Eugene, Omid Poursaeed, and Vitaly Shmatikov.** 2019. “Differential privacy has disparate impact on model accuracy.” *Advances in Neural Information Processing Systems*, 32: 15479–15488.
- Bajpai, Kartikeya, and Klaus Weber.** 2017. “Privacy in public: Translating the category of privacy to the digital age.” In *From Categories to Categorization: Studies in Sociology, Organizations and Strategy at the Crossroads*. Emerald Publishing Limited.
- Balle, Borja, Giovanni Cherubin, and Jamie Hayes.** 2022. “Reconstructing Training Data with Informed Adversaries.” *arXiv preprint arXiv:2201.04845*.

- Bamberger, Kenneth A, and Deirdre K Mulligan.** 2015. *Privacy on the ground: driving corporate behavior in the United States and Europe.* MIT Press.
- Bassolas, Aleix, Hugo Barbosa-Filho, Brian Dickinson, Xerxes Dotiwalla, Paul Eastham, Riccardo Gallotti, Gourab Ghoshal, Bryant Gipson, Surendra A Hazarie, Henry Kautz, et al.** 2019. “Hierarchical organization of urban mobility and its connection with city livability.” *Nature communications*, 10(1): 1–10.
- Bebensee, Björn.** 2019. “Local differential privacy: a tutorial.” *arXiv preprint arXiv:1907.11908*.
- Björkegren, Daniel, and Darrell Grissen.** 2020. “Behavior revealed in mobile phone usage predicts credit repayment.” *The World Bank Economic Review*, 34(3): 618–634.
- Blum, Avrim, Katrina Ligett, and Aaron Roth.** 2008. “A learning theory approach to non-interactive database privacy.” 609–618.
- Blumenstock, Joshua.** 2018a. “Don’t forget people in the use of big data for development.”
- Blumenstock, Joshua.** 2020. “Machine learning can help get COVID-19 aid to those who need it most.” *Nature*.
- Blumenstock, Joshua E.** 2018b. “Estimating Economic Characteristics with Phone Data.” *American Economic Review: Papers and Proceedings*, 108: 72–76.
- Blumenstock, Joshua, Gabriel Cadamuro, and Robert On.** 2015. “Predicting poverty and wealth from mobile phone metadata.” *Science*, 350(6264): 1073–1076.
- Boorstein, Michelle, Marisa Iati, and Annys Shin.** 2021. “Top U.S. Catholic Church official resigns after cellphone data used to track him on Grindr and to gay bars.” *The Washington Post*.
- Bowman, Courtney, Ari Gesher, John K Grant, Daniel Slate, and Elissa Lerner.** 2015. *The architecture of privacy: On engineering technologies that can deliver trustworthy safeguards.* O’Reilly Media, Inc.”
- Brandeis, Louis, and Samuel Warren.** 1890. “The right to privacy.” *Harvard law review*, 4(5): 193–220.

- Branson, Janice, Nathan Good, Jung-Wei Chen, Will Monge, Christian Probst, and Khaled El Emam.** 2020. “Evaluating the re-identification risk of a clinical study report anonymized under EMA Policy 0070 and Health Canada Regulations.” *Trials*, 21(1): 1–9.
- Bun, Mark, Marco Gaboardi, Marcel Neunhoeffler, and Wanrong Zhang.** 2023. “Continual Release of Differentially Private Synthetic Data.” *arXiv preprint arXiv:2306.07884*.
- Bushkin, Arthur A, and Samuel I Schaen.** 1976. *The Privacy act of 1974: a reference manual for compliance*. System Development Corporation McLean, Va.
- Chang, Kent K, Mackenzie Cramer, Sandeep Soni, and David Bamman.** 2023. “Speak, memory: An archaeology of books known to chatgpt/gpt-4.” *arXiv preprint arXiv:2305.00118*.
- Crawford, Kate, and Jason Schultz.** 2014a. “Big Data and Due Process: Toward a Framework to Redress Predictive Privacy Harms.” *Boston College Law Review*, 55(1): 93–128.
- Crawford, Kate, and Jason Schultz.** 2014b. “Big data and due process: Toward a framework to redress predictive privacy harms.” *BCL Rev.*, 55: 93.
- De Montjoye, Yves-Alexandre, Sébastien Gambs, Vincent Blondel, Geoffrey Canright, Nicolas De Cordes, Sébastien Deletaille, Kenth Engø-Monsen, Manuel Garcia-Herranz, Jake Kendall, Cameron Kerry, et al.** 2018. “On the privacy-conscientious use of mobile phone data.” *Scientific data*, 5(1): 1–6.
- Dinur, Irit, and Kobbi Nissim.** 2003. “Revealing information while preserving privacy.” 202–210.
- Dwork, Cynthia, Adam Smith, Thomas Steinke, and Jonathan Ullman.** 2017. “Exposed! a survey of attacks on private data.” *Annual Review of Statistics and Its Application*, 4: 61–84.
- Dwork, Cynthia, Frank McSherry, Kobbi Nissim, and Adam Smith.** 2006a. “Calibrating noise to sensitivity in private data analysis.” 265–284, Springer.
- Dwork, Cynthia, Krishnaram Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor.** 2006b. “Our data, ourselves: Privacy via distributed noise generation.” 486–503, Springer.
- Dwork, Cynthia, Nitin Kohli, and Deirdre Mulligan.** 2019. “Differential Privacy in Practice: Expose your Epsilons!” *Journal of Privacy and Confidentiality*, 9(2).

- Erlingsson, Úlfar, Vasyl Pihur, and Aleksandra Korolova.** 2014. “Rappor: Randomized aggregatable privacy-preserving ordinal response.” 1054–1067.
- Evans, David, Vladimir Kolesnikov, and Mike Rosulek.** 2017. “A pragmatic introduction to secure multi-party computation.” *Foundations and Trends® in Privacy and Security*, 2(2-3).
- Francis, Eilin, Joshua Blumenstock, and Jonathan Robinson.** 2017. “Digital credit: A snapshot of the current landscape and open research questions.” *CEGA White Paper*, 1739–1776.
- Gentry, Craig.** 2009. “Fully homomorphic encryption using ideal lattices.” 169–178.
- Gilad-Bachrach, Ran, Nathan Dowlin, Kim Laine, Kristin Lauter, Michael Naehrig, and John Wernsing.** 2016. “Cryptonets: Applying neural networks to encrypted data with high throughput and accuracy.” 201–210, PMLR.
- Giomi, Matteo, Franziska Boenisch, Christoph Wehmeyer, and Borbála Tasnádi.** 2022. “A Unified Framework for Quantifying Privacy Risk in Synthetic Data.” *arXiv preprint arXiv:2211.10459*.
- Halevi, Shai.** 2017. “Homomorphic encryption.” In *Tutorials on the Foundations of Cryptography: Dedicated to Oded Goldreich*. 219–276. Springer.
- Hambly, Ben, Renyuan Xu, and Huining Yang.** 2021. “Recent Advances in Reinforcement Learning in Finance.” *arXiv preprint arXiv:2112.04553*.
- Hardt, Moritz, and Guy N Rothblum.** 2010. “A multiplicative weights mechanism for privacy-preserving data analysis.” 61–70, IEEE.
- Ha, Trung, Tran Khanh Dang, Tran Tri Dang, Tuan Anh Truong, and Manh Tuan Nguyen.** 2019. “Differential privacy in deep learning: an overview.” 97–102, IEEE.
- Heaton, James B, Nick G Polson, and Jan Hendrik Witte.** 2017. “Deep learning for finance: deep portfolios.” *Applied Stochastic Models in Business and Industry*, 33(1): 3–12.
- Henecka, Wilko, and Matthew Roughan.** 2014. “Privacy-preserving fraud detection across multiple phone record databases.” *IEEE Transactions on Dependable and Secure Computing*, 12(6): 640–651.

- Hill, Kashmir.** 2012. “How target figured out a teen girl was pregnant before her father did.” *Forbes, Inc.*
- Hipgrave, Shaun.** 2013. “Smarter fraud investigations with big data analytics.” *Network Security*, 2013(12): 7–9.
- Homer, Nils, Szabolcs Szelinger, Margot Redman, David Duggan, Waibhav Tembe, Jill Muehling, John V Pearson, Dietrich A Stephan, Stanley F Nelson, and David W Craig.** 2008. “Resolving individuals contributing trace amounts of DNA to highly complex mixtures using high-density SNP genotyping microarrays.” *PLoS genetics*, 4(8): e1000167.
- Hoofnagle, Chris Jay.** 2016. *Federal Trade Commission privacy law and policy*. Cambridge University Press.
- Hsu, Justin, Marco Gaboardi, Andreas Haeberlen, Sanjeev Khanna, Arjun Narayan, Benjamin C Pierce, and Aaron Roth.** 2014. “Differential privacy: An economic method for choosing epsilon.” 398–410, IEEE.
- Jarin, Ismat, and Birhanu Eshete.** 2022. “MIAShield: Defending Membership Inference Attacks via Preemptive Exclusion of Members.” *arXiv preprint arXiv:2203.00915*.
- Jernigan, Carter, and Behram FT Mistree.** 2009. “Gaydar: Facebook friendships expose sexual orientation.” *First Monday*.
- Ji, Zhanglong, Zachary C Lipton, and Charles Elkan.** 2014. “Differential privacy and machine learning: a survey and review.” *arXiv preprint arXiv:1412.7584*.
- Johnson, Noah, Joseph P Near, and Dawn Song.** 2018. “Towards practical differential privacy for SQL queries.” *Proceedings of the VLDB Endowment*, 11(5): 526–539.
- Kairouz, Peter, Sewoong Oh, and Pramod Viswanath.** 2016. “Extremal mechanisms for local differential privacy.” *The Journal of Machine Learning Research*, 17(1): 492–542.
- Kenthapadi, Krishnaram, Ahsan Chudhary, and Stuart Ambler.** 2017. “LinkedIn Salary: A system for secure collection and presentation of structured compensation insights to job seekers.” 13–24, IEEE.
- Kenthapadi, Krishnaram, and Thanh TL Tran.** 2018. “Pripearl: A framework for privacy-preserving analytics and reporting at LinkedIn.” 2183–2191.

- Kohli, Nitin.** 2021. “Leveraging Differential Privacy While Attending to Social and Political Commitments.” PhD diss. University of California, Berkeley.
- Kohli, Nitin, and Paul Laskowski.** 2018. “Epsilon Voting: Mechanism Design for Parameter Selection in Differential Privacy.” 19–30, IEEE.
- Kohli, Nitin, and Paul Laskowski.** 2023. “Differential Privacy for Black-Box Statistical Analyses.” *Proceedings on Privacy Enhancing Technologies*, 3: 418–431.
- Kohli, Nitin, Emily Aiken, and Joshua Blumenstock.** 2023. “Privacy Guarantees for Personal Mobility Data in Humanitarian Response.” *arXiv preprint arXiv:2306.09471*.
- Koops, Bert-Jaap, Bryce Clayton Newell, Tjerk Timan, Ivan Skorvanek, Tomislav Chokrevski, and Masa Galic.** 2016. “A typology of privacy.” *U. Pa. J. Int’l L.*, 38: 483.
- Kroll, Joshua A, Nitin Kohli, and Paul Laskowski.** 2019. “Privacy and policy in polystores: a data management research agenda.” In *Heterogeneous Data Management, Polystores, and Analytics for Healthcare*. 68–81. Springer.
- Mansfield-Devine, Steve.** 2015. “The ashley madison affair.” *Network Security*, 2015(9): 8–16.
- McMahan, Brendan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Agueray Arcas.** 2017. “Communication-efficient learning of deep networks from decentralized data.” 1273–1282, PMLR.
- Mohan, Prashanth, Abhradeep Thakurta, Elaine Shi, Dawn Song, and David Culler.** 2012. “GUPT: privacy preserving data analysis made easy.” 349–360.
- Mulligan, Deirdre K, Colin Koopman, and Nick Doty.** 2016. “Privacy is an essentially contested concept: a multi-dimensional analytic for mapping privacy.” *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 374(2083): 20160118.
- Nagendra, Narayan Prasad, Gopalakrishnan Narayanamurthy, and Roger Moser.** 2020. “Satellite big data analytics for ethical decision making in farmer’s insurance claim settlement: minimization of type-I and type-II errors.” *Annals of Operations Research*, 1–22.
- Narayanan, Arvind, and Edward W Felten.** 2014. “No silver bullet: De-identification still doesn’t work.” *White Paper*, 1–8.

- Narayanan, Arvind, and Vitaly Shmatikov.** 2008. “Robust de-anonymization of large sparse datasets.” 111–125, IEEE.
- Near, Joe.** 2018. “Differential privacy at scale: Uber and Berkeley collaboration.”
- Nissenbaum, Helen.** 2004. “Privacy as contextual integrity.” *Wash. L. Rev.*, 79: 119.
- Officer, Chief Privacy, and Timothy J Keefer.** 2009. “Privacy and Civil Liberties Policy Guidance Memorandum.”
- Ohm, Paul.** 2009. “Broken promises of privacy: Responding to the surprising failure of anonymization.” *UCLA L. Rev.*, 57: 1701.
- Oliver, Nuria, Bruno Lepri, Harald Sterly, Renaud Lambiotte, Sébastien Deletaille, Marco De Nadai, Emmanuel Letouzé, Albert Ali Salah, Richard Benjamins, Ciro Cattuto, et al.** 2020. “Mobile phone data for informing public health actions across the COVID-19 pandemic life cycle.”
- Ovide, Shira.** 2021. “The Nightmare of Our Snooping Phones.” *The New York Times*.
- Papernot, Nicolas, Shuang Song, Ilya Mironov, Ananth Raghunathan, Kunal Talwar, and Úlfar Erlingsson.** 2018. “Scalable private learning with pate.” *arXiv preprint arXiv:1802.08908*.
- Post, Robert C.** 1989. “The social foundations of privacy: Community and self in the common law tort.” *Calif. L. Rev.*, 77: 957.
- Pujol, David, Ryan McKenna, Satya Kuppam, Michael Hay, Ashwin Machanavajjhala, and Gerome Miklau.** 2020. “Fair decision making using privacy-protected data.” 189–199.
- Rivest, Ronald L, Len Adleman, and Michael L Dertouzos.** 1978. “On data banks and privacy homomorphisms.” *Foundations of Secure Computation*, 4(11): 169–180.
- Rolf, Esther, Max Simchowitz, Sarah Dean, Lydia T. Liu, Daniel Björkegren, Moritz Hardt, and Joshua Blumstock.** 2020. “Balancing Competing Objectives with Noisy Data: Score-Based Classifiers for Welfare-Aware Machine Learning.” *International Conference on Machine Learning (ICML 2020)*. arXiv: 2003.06740.

- Sangers, Alex, Maran van Heesch, Thomas Attema, Thijs Veugen, Mark Wiggerman, Jan Veldsink, Oscar Bloemen, and Daniël Worm.** 2019. “Secure multiparty PageRank algorithm for collaborative fraud detection.” 605–623, Springer.
- Sedenberg, Elaine M, and Deirdre K Mulligan.** 2015. “Public Health as a model for cybersecurity information sharing.” *Berkeley Tech. LJ*, 30: 1687.
- Shokri, Reza, Marco Stronati, Congzheng Song, and Vitaly Shmatikov.** 2017. “Membership inference attacks against machine learning models.” 3–18, IEEE.
- Sklansky, David Alan.** 2014. “Too much information: How not to think about privacy and the Fourth Amendment.” *Calif. L. Rev.*, 102: 1069.
- Solove, Daniel J.** 2005. “A taxonomy of privacy.” *U. Pa. L. Rev.*, 154: 477.
- Sweeney, Latanya.** 2002. “k-anonymity: A model for protecting privacy.” *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 10(05): 557–570.
- Tang, Jun, Aleksandra Korolova, Xiaolong Bai, Xueqiang Wang, and Xiaofeng Wang.** 2017. “Privacy loss in apple’s implementation of differential privacy on macos 10.12.” *arXiv preprint arXiv:1709.02753*.
- Tantipongpipat, Uthaipon Tao, Chris Waites, Digvijay Boob, Amaresh Ankit Siva, and Rachel Cummings.** 2021. “Differentially private synthetic mixed-type data generation for unsupervised learning.” *Intelligent Decision Technologies*, 15(4): 779–807.
- Taylor, Linnet.** 2016. “No place to hide? The ethics and analytics of tracking mobility using mobile phone data.” *Environment and Planning D: Society and Space*, 34(2): 319–336.
- Thompson, Stuart A, and Charlie Warzel.** 2019. “Twelve million phones, one dataset, zero privacy.” In *Ethics of Data and Analytics*. 161–169. Auerbach Publications.
- Truex, Stacey, Nathalie Baracaldo, Ali Anwar, Thomas Steinke, Heiko Ludwig, Rui Zhang, and Yi Zhou.** 2019. “A hybrid approach to privacy-preserving federated learning.” 1–11.
- Ullman, Jonathan.** 2022. “What is Synthetic Data?” *DifferentialPrivacy.org*, <https://differentialprivacy.org/synth-data-0/>.

- van Egmond, Marie Beth, and Thomas Rooijackers.** 2021. “Privacy-Preserving Collaborative Money Laundering detection.” *ERCIM NEWS*, 27.
- Wagh, Sameer, Xi He, Ashwin Machanavajhala, and Prateek Mittal.** 2021. “DP-cryptography: marrying differential privacy and cryptography in emerging applications.” *Communications of the ACM*, 64(2): 84–93.
- Ware, Willis H.** 1973. “Records, computers and the rights of citizens.” RAND CORP SANTA MONICA CALIF.
- Westin, Alan F.** 1968. “Privacy and freedom.” *Washington and Lee Law Review*, 25(1): 166.
- Yao, Andrew C.** 1982. “Protocols for secure computations.” 160–164, IEEE.
- Yoo, Ji Su, Alexandra Thaler, Latanya Sweeney, and Jinyan Zang.** 2018. “Risks to patient privacy: a re-identification of patients in Maine and Vermont Statewide Hospital Data.” *J Technol Sci*, 2018: 2018100901.
- Zhang, Zhikun, Tianhao Wang, Ninghui Li, Jean Honorio, Michael Backes, Shibo He, Jiming Chen, and Yang Zhang.** 2021. “Privsyn: Differentially private data synthesis.”