INTERNATIONAL MONETARY FUND

Privacy Technologies & The Digital Economy

A Primer for Supervisors

Parma Bains and Tamas Gaidosch

WP/25/60

2025
MAR

WORKING PAPER

 WP/25/60

IMF Working Paper
Monetary and Capital Markets

Privacy Technologies & The Digital Economy: A Primer for Supervisors

Prepared by Parma Bains and Tamas Gaidosch

Authorized for distribution by Jay Surti
March 2025

ABSTRACT: Ensuring that users and society-at-large derive the maximum benefit from digital technology requires active and open participation in the digital economy. However, such participation is not without risks and users may withhold or withdraw their active participation in response to such factors. One important reason for doing so is users' privacy concerns, which may induce behavior that limits digital footprints in order to shield personal data from third parties and governments. Coupled with regulation, privacy technologies can help build trust in the digital economy. If designed and deployed appropriately, they could form the basis of trust in the digital economy. We offer three considerations for supervisors. First, they should understand the strengths and weaknesses of privacy technologies, and this primer aims to provide a foundational tool to achieve this. Second, domestic collaboration and international cooperation is indispensable to improving knowledge sharing and providing clarity regarding mandates and rules. Third, they need to understand the cybersecurity implications and tradeoffs in using privacy technologies.

RECOMMENDED CITATION: Bains, P. and Gaidosch, T. (2024). *Privacy Technologies & The Digital Economy: A Primer for Supervisors*. IMF Working Paper, Monetary and Capital Markets.

JEL Classification Numbers:	G21, O16, O32, O33
Keywords:	Fintech; cyber; digital; privacy; PET; regulation; supervision; encryption; cryptography
Author's E-Mail Address:	pbains@imf.org; tgaidosch@imf.org

The authors would like to thank Jay Surti, Dirk Jan Grolleman, Nobu Sugimoto, Itai Agur, Tanai Khiaonarong, Yan Carriere-Swallow, Franko Solomon, and Herve Tourpe for their input

WORKING PAPERS

Privacy Technologies & The Digital Economy

A Primer for Supervisors

Prepared by Parma Bains and Tamas Gaidosch

Contents

Glossary **3**

Introduction **4**
- Big Data, Trust, and the Digital Economy 4
- Risks of Big Data 5
- The role of privacy technologies 6

Privacy Technologies **8**
- Input Privacy Technologies 8
  - Homomorphic Encryption 8
  - Secure Multiparty Computation 11
  - Federated Learning 12
- Output Privacy Technologies 13
  - Zero-Knowledge Proofs 14
  - Data Masking 15
  - Differential Privacy 17
  - Synthetic Data 19

Considerations for Supervisors **22**
- Outreach, engagement and understanding trade-offs 22
- Coordination and cooperation 23
- Understanding and managing cyber implications 24

Conclusion **26**

References **28**

Glossary

AE	Advanced Economy
AI	Artificial Intelligence
API	Application Programming Interface
BaaS	Banking-as-a-Service
BCBS	Basel Committee on Banking Supervision
BigTech	Large Technology Conglomerates Operating Across Markets
DLT	Distributed Ledger Technology
EMDE	Emerging Market and Developing Economy
FCA	Financial Conduct Authority (UK)
Fintech	Financial Technology
ICO	Information Commissioner's Office
KYC	Know Your Customer
ZKP	Zero Knowledge Proof

Introduction

Big Data, Trust, and the Digital Economy

The digital data universe has grown from around two zettabytes[1] in 2010, to approximately one hundred zettabytes in 2024.[2] The growth of this large volume of data (commonly known as "Big Data") has served as a prime driver of the application of innovative digital technologies that has transformed financial services.[3] The commercial success of most technologies deployed during the "fintech revolution" is predicated on the use of consumer data, or Big Data. For example, distributed ledger technologies (DLT), at their core, are methods to transfer, store and process data in distributed systems, with the ability to establish and maintain trust between users based on technology. Machine learning, the subset of artificial intelligence (AI) most often deployed in financial services relies on the availability of high-quality data to train algorithms for a desired result.[4] Open Finance is fundamentally about the controlled sharing of data to unlock innovation, often referred to as data portability.[5] The business models of large technology conglomerates (BigTech) are predicated on the generation, use, and regeneration of Big Data.

The growth and increased use of Big Data creates new opportunities to unlock potential efficiencies in financial markets by introducing tailored products and services that increase consumer choice in a more inclusive manner. However, the use of Big Data also creates potential new risks such as pervasive profiling of individuals, targeted misinformation campaigns, and data breaches with more severe consequences. Without appropriate data privacy as well as conduct and prudential regulation and supervision, these risks can increase the likelihood of consumer harm, erode market integrity, adversely impact financial safety and soundness, and may potentially, threaten financial stability. Importantly, a higher likelihood of such risks can erode trust in digital financial services and inhibit the generation and sharing of data by individuals and firms, which is the foundation on which the process of fintech innovation is built, thereby constraining efficiencies and inclusive growth.

Two main ways of creating that trust are through financial regulation and supervision, and the use of privacy technologies. This paper provides a primer for financial services regulators and supervisors to better understand how the use of privacy technologies could manage some of these risks, while preserving the potential for digital financial innovation.[6]

[1] A zettabyte is equivalent to a billion terabytes, or a trillion gigabytes

[2] Data growth worldwide 2010-2025 | Statista

[3] Fintech is technology enabled innovation in financial services underpinned by connectivity, big data, and the ability for end users to directly engage with a product and service.

[4] An algorithm is a method of designing a sequence of actions to solve a problem, and in machine learning these actions optimize automatically through experience with limited human intervention

[5] Open Banking allows regulated firms to share to certain data (often account and transaction data) of individual users with their consent. Open Finance broadens this concept to other types of financial data.

[6] There are other ways to create or enhance trust, for example other methods of privacy by design or minimizing the collection of personally identifiable information but these are not in the scope of this paper

The paper is structured in three main sections and builds on work done by the Fund and other global bodies.[7] It targets financial sector supervisors working in functional areas of fintech, information technology, and cybersecurity. The paper could benefit supervisors who also have a role in catalyzing, designing, developing, or operating a key component of the financial system.

The rest of the Introduction explores the risks of Big Data, the role that financial regulation and supervision can play in managing some of these risks and introduces the concept of privacy technologies. The following section introduces common forms of privacy technologies that could shape financial services in the coming years. The final section offers three key considerations for supervisors to help support the development of privacy technologies in financial markets in a sensible and sustainable way.

Risks of Big Data

Data is generally not owned by data subjects but by large entities like banks and BigTechs that generate revenues by owning and controlling its use, combination, and distribution. In addition to not retaining control and ownership of data that they produce (including personal information), users must face the possibility of data leaks or hacks which could put their sensitive personal information in the hands of unauthorized third parties, including malicious actors. Coupled with a desire to limit digital footprints for privacy and other purposes (like avoiding unwanted marketing and the access of banks, BigTech, and governments to their information), this could create an environment where users' trust in owners of their data gets eroded and their reluctance to use services and share this data grows.[8] Not only does this increase the number of individuals excluded from the digital economy but also reduces innovation and growth due to sub-optimal sharing, and consequently sub-optimal transformation, distribution and reuse of personal information and data. That is, the reduction in trust due to perception of a threat to privacy of personal data can hinder efficient innovation and increase exclusion, both of which are welfare reducing.

For example, excluded individuals could face major challenges in accessing products and services at a reasonable cost. Currently, data granularity might not capture all indicators of a users' financial health, but in the future, they could face a scenario where they are expected to share granular data that reflects poorly on themselves, leading to bad outcomes. Evidence from early Open Banking implementations supports this notion.[9] Excluded populations may also come from specific demographics potentially amplifying existing gaps in access to financial products and services.[10] While some members of these populations may not share data for privacy reasons, data shared by close contacts or similar users could provide entities with information on these excluded individuals (linked data), therefore putting them in a lose-lose situation. This raises concern around individual preferences and collective responsibilities of society, with the possibility of data privacy considered a public good.[11] However, there appears to be a digital privacy paradox where users may say they value privacy but do not act accordingly by ensuring

[7] Staff Discussion Notes Volume 2021 Issue 005: Toward a Global Approach to Data in the Digital Age (2021)

[8] The expectation of privacy differs across cultures and demographics, but the right to privacy and the option of privacy is important. See Privacy (Stanford Encyclopedia of Philosophy)

[9] Bank of England Staff Working Paper No. 1,059

[10] Nothing to hide? Gender and age differences in willingness to share data (bis.org)

[11] "Privacy as a Public Good" by Joshua A. T. Fairfield and Christoph Engel (duke.edu)

their data remains private. Additionally, digital privacy is context, culture, and demographic dependent with older users more like to value privacy, and users more concerned about privacy of transactions and financial data than they are of, say, browsing or app usage data.[12]

The role of privacy technologies

The presence of some of these concerns has led to the development of data protection regulation, but the promulgation of these has been uneven globally. Many emerging and developing economies lack data protection frameworks. Where data protection regulations do exist, some may have significant gaps compared to internationally recognized guidelines (such as the OECD guidelines), while resource constrained authorities may struggle to supervise or enforce them.

Without robust data protection regulations, there will be a lack of trust that may mean users are less likely to play active roles in the digital economy and may withdraw their participation altogether. The OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data provide an important framework to help guide authorities in this area.[13] While the development of strong data protection regulation is a question for governments and data protection agencies, financial regulators consider compliance with these requirements in their supervision where relevant in view of their mandate (e.g., consumer protection, market integrity, risks related to the safety and soundness of financial institutions). It is only when these regulations and proportionate supervision are in place that we can move onto questions around data ownership or data portability (like Open Banking / Finance).[14]

Given that the business models of many tech firms are predicated on learning the most about their users, it is important for data subjects to have some control over their data. In the digital economy, it can be difficult to replicate privacy that we expect in wider society.

The easiest way to retain privacy in the digital economy is to not interact with the digital economy or prohibit data portability, but this is not an ideal outcome and could leave users without access to basic services. It may also act as a drag on the growth of the digital economy. Another option is through data minimization, sharing only the most basic and necessary data needed to get important services. While this allows a degree of access, it doesn't deliver the full benefits for users and markets. A further option could be access through anonymization, where data is shared only in instances where the data subject is not known. While useful, anonymization has considerable drawbacks.[15] A promising option is the development of privacy technologies.

[12] https://www.nber.org/system/files/working_papers/w30943/w30943.pdf

[13] OECD (2002), *OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data*, OECD Publishing, Paris, https://doi.org/10.1787/9789264196391-en.

[14] Key Considerations for Open Finance | CGAP Research & Publications

[15] in the US, gender, date of birth, and zip code are all that are needed to identify the majority of US citizens. By linking these attributes in a supposedly anonymized database it is possible to identify individuals within the anonymized data set Sweeney_Article.pdf (epic.org)

Privacy technologies (sometimes called privacy enhancing or privacy preserving technologies) could enable more data sharing while improving regulatory protections of data subjects. They could form the bedrock of the digital economy, allowing users to participate in the digital economy with greater confidence while maintaining the attractiveness of business cases for firms. Privacy technologies, based on the concept of privacy by design, refers to a set of technologies that aim to protect sensitive data and maintain the confidentiality and integrity of data, while also allowing use of that data.

If designed and deployed in compliance with data protection regulation, privacy technologies could form the basis of trust in the digital economy, allowing users to participate in the digital economy more freely and confidently with different degrees of privacy expectations, just as you would expect in the wider economy. This could allow for the provision of useful products and services like account aggregation as well as tailored (and potentially cheaper) products and services delivered through emerging technologies trained on large, high-quality data. They could also help financial authorities better manage risks, including early and accurate identification of risks to mandates, such as financial stability. The use of privacy technologies by commercial firms and regulatory authorities could feed down into positive outcomes for data subjects at the individual level as privacy embedded infrastructure, products, and services, would mean greater privacy for the individual.

However, privacy technologies are not a panacea. They cannot prevent situations where data subjects choose, or are encouraged, to share more personal data than is good for them or society. They also cannot help where collected data is misused by firms or repurposed beyond the original purpose.

First, it is important to understand that privacy technologies do not replace effective and robust data protection regulation, and if supervisors misunderstand this concept, it could lead to risks to users, markets, and potentially financial stability if the technology is used at scale, for example in underpinning digital transactions, or used as a basis in machine learning algorithms.

Second, the use of privacy technology that supports greater data portability could lead to the entrance of new firms, products or services that lack effective regulatory oversight or that are more challenging to regulate. Data portability could help facilitate the movement of data from firms with robust regulatory oversight, including prudential oversight, to those that are more lightly regulated, and it could lead to the growth of complex business models including new services based on problematic uses of AI.

Third, the misuse or over-reliance of privacy technologies by regulators and supervisors for their own use (for example as SupTech tools) could generate new risks. Relatedly, the use of privacy technologies by (or involving) regulatory authorities with large resources and expertise could result in agencies with limited resources falling further behind in terms of supervisory capabilities and potentially creating opportunities for cross-border regulatory arbitrage.

Fourth, since some of these technologies are complex and not all sufficiently validated, risks, such as those to cybersecurity, could increase with their implementation, because of the greater likelihood of vulnerabilities and less availability of specialist expertise for secure deployment.

Privacy Technologies

At their broadest, privacy technologies can be classified into two categories: ***Input Privacy***, which serve to reduce data access, and ***Output Privacy***, which serve to reduce inference from shared data (ICO,2023).[16] Categorizing privacy technologies is not easy as many overlap several categories, and different institutions prefer different approaches to categorization.

In this section we explore three input technologies (homomorphic encryption, secure multiparty computation, and federated learning), and four output technologies (zero-knowledge proofs, data masking and tokenization, differential privacy, and synthetic data).

Input technologies like homomorphic encryption could be used wherever computations are outsourced, for example where a third-party conducts risk assessments or due diligence on behalf of a financial services firm. Secure Multiparty Computation could be particularly useful where multiple parties collaborate on a common dataset to support fraud detection. Federated Learning is already being used in financial services to allow multiple parties to pool training data to improve products and services.

Output technologies like zero-knowledge proofs are a core component of experiments in digital payments that allow privacy-preserving transactions where payments reveal minimal information. Data masking and tokenization have been used for many years in credit and debit card-based payments and have been a key component of some mobile digital wallets. Differential privacy technologies have been utilized by large technology conglomerates to share insights and trends without revealing the underlying data. Synthetic data is an area that several regulators are exploring to support machine learning innovation in financial services through new tools like digital and data sandboxes.

Input Privacy Technologies

Input privacy technologies aim to reduce data access and encompass a variety of technologies and use cases and entail some important tradeoffs (Table 1). Homomorphic encryption, secure multiparty computation, and federated learning are three of the most popular input privacy technologies.

Homomorphic Encryption

Encryption technologies have evolved over 2000 years from simple substitution cyphers to today's highly sophisticated and virtually unbreakable algorithms, even when considering the threat posed by quantum computing.[17] In a nutshell, encryption means scrambling data using a key in a way that it can only be

[16] More narrowly, privacy technologies can be classified into four categories, data obfuscation tools (increase privacy by altering data), encrypted data processing tools (keeping data encrypted when in use), federated and distributed analytics (not all data is visible to all users), and data accountability tools (allowing subjects control of their own data) (OECD, 2023), although some privacy technologies might fall into multiple categories.

[17] Virtually unbreakable in the sense that correct use of the algorithm results in encrypted data that cannot be decrypted without the key within a meaningful time while expending realistic resources. For highly secure algorithms this can mean hundreds of years or more. There is at least one completely unbreakable encryption algorithm no matter the time and resource expenditure (the one-time pad), but its use is very limited and non-existent in the financial sector.

unscrambled by using the same key or a complementary key. By a key, we mean a string of characters that, when processed through a cryptographic algorithm, can encode or decode data. Only persons or systems in possession of the relevant key have access to data in the sense of being able to learn its meaning and use it. Therefore, encryption can be seen as controlling data access, and so, is an input privacy technology.

A major issue in using classic encryption in the digital supply chain (where various companies are involved in the delivery of a digital product or service) is that data must be decrypted before acting on it, and so third parties in the chain can learn its content. Appropriate legal, governance and technology arrangements can mitigate the risk, but assurance on compliance eventually hinges on trust, and with large amounts of sensitive data, this is not always adequate. The concern is more acute in the context of cloud computing, for example with some user organizations and regulators being wary of processing sensitive data at service providers, especially outside of their home jurisdictions.

Homomorphic encryption is a potential solution to the problem because it allows the processing of encrypted data directly, without first decrypting it by making use of encryption and an algebraic system. Plain text can be converted into unreadable ciphertext, functions can be performed directly on that ciphertext, then the results sent back to the owner who can decrypt it with a private key. We can say a third-party is able to work on encrypted data without the need to decrypt it at all.

Operations on encrypted data lead to the same results as operations on the plain unencrypted data.[18] Homomorphic encryption is, therefore, a win-win for data owners and data processors because it removes the risk of inadvertent or malicious data breaches at processing facilities. It is particularly well suited as a privacy technology in outsourcing arrangements, especially in cloud outsourcing. It means if someone hacks into the dataset at any point, all they get is unusable information and not the information that needs to be kept secure.

Newer homomorphic encryption schemes are *fully homomorphic,* meaning that any operation is supported, as opposed to early schemes in which only one type was.[19] This widens the field of applicability considerably. However, while efficiency is an issue, and notable gains have been achieved in recent times, homomorphic encryption is still several orders of magnitude slower than operations on unencrypted data. Therefore, tradeoffs are necessary about what kind and amount of data firms protect with homomorphic encryption and for now it remains commercially unviable for large-scale computations. There are other forms of homomorphic encryption that may improve efficiency but at the expense of functionality. For example, *somewhat homomorphic encryption* reduces the amount of functions that can be run on the encrypted data, which is fixed in advance, while *partial homomorphic encryption* supports only single functions (addition or multiplication) but provides for better performance.

[18] The result of an operation on some homomorphically encrypted data will remain encrypted. When decrypted, it will be identical to the result of the same operation on the plain unencrypted data.

[19] In homomorphic encryption schemes the operations are, depending on whether partial, somewhat, or full, addition or multiplication, addition and multiplication but on a limited basis, or addition and multiplication a limited amount of times.

Table 1. Input Privacy Technologies compared

	Homomorphic Encryption	Secure Multiparty Computation	Federated Learning
What is it?	Allows computations to be carried out on encrypted data	Enables parties to jointly compute functions over their data while keeping those data private	Trains a machine learning model across multiple decentralized nodes holding local data samples, without exchanging them
Data Privacy Level	**High** Data remains encrypted during processing	**High** Original data are never revealed to other parties	**High** Data remains on the local node
Overheads	**High** Due to the complexity of operations on encrypted data.	**Moderate to high** Depending on the complexity of the computation and the protocol used.	**Low to moderate** Mainly involves training local models and aggregating updates
Scalability	**Low to moderate** Heavily dependent on the encryption scheme and the computing architecture	**Moderate** Difficult to scale well because of both computational and communication overheads between parties	**High** It scales well with the number of participants and the amount of data
Use Cases	Secure data processing in the cloud, secure voting systems, private information retrieval, open banking / finance	Collaborative data analysis, joint financial processing	Decentralized machine learning, collaborative learning, privacy-preserving ML and AI
Less suited for	Real-time and/or large-scale applications	Large-scale applications	Not suited for applications other than ML/AI

Going forward, if made more computationally efficient, homomorphic encryption can play a significant role in protecting training data for machine learning models. The model can use homomorphically encrypted data sets to calibrate its parameters while preserving privacy. It might also be useful in areas like Open Finance where it might allow consumers to request data to be shared in an encrypted format and for third parties to then run functions on this encrypted data to provide new products and services for users (such as financial advice). It could also be used where banks or lending firms use third-party services for credit scoring or risk assessments. Homomorphic encryption could also be useful in terms of regulatory reporting where supervisors are able to request more confidential or sensitive information to support offsite monitoring which can be delivered and manipulated in encrypted format.

Secure Multiparty Computation

Secure multiparty computation (SMPC) enables joint computations by multiple parties who are unwilling to reveal their own data to the others. A participating party only shares a part, or some derivative of its own data, from which the full data cannot be inferred by any other party. By appropriately choosing how and what to share (the process of which is called the secret sharing protocol), computations such as statistics can be done, while preserving the privacy of each participant.

The classic example is computing averages of sensitive data series while each data point in the series is kept private by its owner. Using SMPC, every participant can learn the average of the series, and will not know any data point but their own.[20] The concept can be extended to more sophisticated data structures and processing. The net effect is as if participants entrusted a third party with their data who did the processing on their behalf and shared the results only. SMPC eliminates the need for a trusted third party (e.g., data broker). Moreover, certain protocols can detect dishonest or malicious participants who try to subvert the process. However, it is important to note that SMPC cannot protect against incorrect results or data leaks if the majority colludes to break the protocol.[21]

Homomorphic encryption can be used as an enabler in SMPC. For example, each party encrypts their data and sends it to a central processing facility, which returns the results and never learns the data since it has no keys to it. Of course, it does not need any, because it can work on the encrypted data as it is homomorphically encrypted. Note that the central processing party need not to be trusted beyond executing correct computations.

Just as with homomorphic encryption—and virtually all other privacy technologies—there is a computational overhead with SMPC that increases in-step with the size of datasets. There is a telecommunication overhead as well because more messages need to be exchanges between collaborating parties.

SMPC can be used in financial services, e.g., in privacy-preserving data analytics, including trend analysis, benchmarking, risk assessment and collaborative fraud detection without the need to expose

[20] There are limitations because some information about other parties' input can inferred from the results.

[21] This is conceptually similar to how a blockchain, which also eliminates the need for a trusted third party, is vulnerable to 51 percent attacks.

individual transactions to each other. Other uses include supporting transactions and smart contracts on blockchains, secure cross-border payments, high security digital wallets (splitting the private key among multiple devices), and secure sharing of market data and models. A special use case of SMPC can be in ensuring compliance with data sovereignty requirements as it can support cross-border data transfers that do not reveal sensitive financial or personal information.

Federated Learning

Commonly, machine learning algorithms are trained on data that are kept in a central server and both the data and the server are held by the entity developing the algorithm.[22] While this approach could be beneficial for firms with access to large data sets, it can be costly and challenging for smaller firms. Gathering data from other sources, like other entities or devices held by users, can generate risks to privacy and in some jurisdictions could be limited by data protection or portability frameworks to protect users.

Federated learning operates on the notion that data remains distributed across various local servers thereby enhancing data privacy by avoiding central storage. In its most common form, known as centralized federated learning, a central server initiates a process by sending an initial model (sometime called a global model) to local nodes. The local nodes then train a local model using the data available to them (be that device data, or proprietary data held by an entity) and send summary results and model updates back to the central server – but not the raw data. Part of this local training involves computing "gradients" or updates. These are calculations that help figure out how to adjust the model to make fewer mistakes and improve. The nodes then send these gradient-derived updates and summary results back to the central server. The central server uses these model updates from each local node to train their global model, and this updated model is fed back down to the local nodes and the process is repeated until the model reaches whatever the desired aim is. Raw data is kept at the decentralized local node level and never travels back to the central server.

A less common alternative is the decentralized federated learning model where there is no centralized server and local nodes communicate directly with each other and all update the global model directly. This increases privacy and security but can be computationally more intensive and reduce efficiency. There are several types of federated learning algorithms that describe the way data are aggregated by the central server, with the most popular being federated stochastic gradient descent (FedSGD) and federated averaging (FedAve).[23]

While privacy preserving, federated learning can still suffer from inference attacks where sensitive data on subjects is used in the training of the model. There are three possible types of inference attacks, model inversion (where an attacker tries to invert the model to find the underlying data point),

[22] Servers are also commonly contracted by financial incumbents and other technology driven firms from larger technology conglomerates, the so-called BigTechs.

[23] In the FedSGD the gradients are calculated on a small subset of samples rather than all the data, and these "mini-batches" are then returned to the central server in the form of model updates. This reduces the amount of communication between the central server and local nodes, and potentially improves privacy, at the cost of model accuracy (which may take longer to achieve). It's the most common way of operationalizing the centralized federated learning model. In the FedAve, the model updates received by the central server are aggregated and all the updates are averaged.

membership inference (where an attacker tries to determine if a given data point was present in the underlying data), and property inference (where an attacker tries to determine the presence of some property in the underlying data). These happen when an attacker can learn about the underlying data set, not by gaining access to the data, but through using the model's own responses. This risk increases if the data are bijective, or injective but non-bijective and non-injective data can still generate privacy risks. Inference attacks can be used to make inferences about the data that the model was trained on, for example, by training a second algorithm to reconstruct training data based on model updates and outputs. Another issue arises from memorization where the algorithm seems to memorize raw data where updates might contain patterns or specific features that can be traced back to the original data. Memorization, which is a common problem in neural networks – particularly where the model is too complex relative to the data, could also increase the potential for successful inference attacks,

A growing area of development is combining federated learning with other types of privacy technologies such as secure aggregation which can make use of cryptographic tools like secure multi-party computation and homomorphic encryption, or differential privacy technologies. In the former, the central server in unable to view individual updates that are sent to the server, but rather only the final aggregated output. For example, using homomorphic encryption, the model updates are encrypted by the local server, sent to the central server in the encrypted form which then aggregates all the encrypted local model updates received from different local servers, and only then are the updates decrypted which is then used to update the global model. In the latter, differential privacy (explored below) can be implemented either at the local server or the central server (i.e., noise is added before aggregation at the central server, or by the central server during the aggregation process).[24]

Federated learning has the potential to be useful wherever machine learning algorithms are being developed, and therefore its potential utility in financial services is large. Importantly, it might allow several firms in a market to work closely together on developing new machine learning models while potentially remaining complaint of national data protection frameworks, although such compliance will differ on a case-by-case basis. Some banks are working with third party technology companies that provide the infrastructure for federated learning in order to improve detection of suspicious transactions and other financial crimes without exposing sensitive information. Before submitting the data to the third-party technology company for model aggregation, regulated banks are also adding noise to the data to further improve privacy.[25]

Output Privacy Technologies

Output privacy technologies aim to reduce inference from shared data and there are also many examples of such technologies. Table 2 compares some of the most common output privacy technologies that could be implemented across regulated financial services along with relevant trade-offs and use cases. Here, we explore four of the most popular: zero-knowledge proofs, data masking, differential privacy, and synthetic data.

[24] Adding noise at the local server is likely to generate stronger privacy protection although this is at the expense of model degradation given the central server is unable to control the amount of noise added across servers. A decision may depend on how trusted the central server is to local servers.

[25] Building privacy-preserving federated learning to help fight financial crime - IBM Research

Zero-Knowledge Proofs

Zero-Knowledge Proof (ZKP) protocols are among the most promising privacy technologies and already have applications within financial services markets. Personally Identifiable Information (PII) is often used to prove or verify that a person is in possession of certain information, is authorized to perform or has performed some action, is entitled to something, has certain attributes, etc. (e.g. a document like a driving license can prove age because it has PII.) With ZKP protocols this is possible without revealing the PII or oversharing unnecessary PII (e.g. proving age without sharing other PII like address, gender etc.).

An example of zero knowledge proof of group membership can illustrate this point. If somebody needs to prove membership of a group but wishes to maintain their privacy, they can use a cryptographic key that is shared among group members and the verifier party (and nobody else) to encrypt a message that affirms membership. The verifier then uses the same key to decrypt the message and if they succeed then membership is proven because only members have access to the key and only that specific key will result in a correct decryption. Since the key is decoupled from PII, privacy is preserved.

The concept can be generalized to any kind of sensitive information, proofs, and parties (systems, people, or organizations). It's important to note that ZKPs protect both the prover and the verifier. The prover does not have to disclose sensitive information, and the verifier does not need to collect and process it, so reducing incentives for data theft and cybercrime in general.

There are two main types of ZKP protocols: (i) interactive, which require multiple rounds of interaction between the prover and the verifier so that the verifier can be certain the prover has whatever proof is required; [26] and (ii) non-interactive, with which the prover can generate a proof without any interaction with the verifier. ZKPs should be complete (if a statement is true, an honest prover can convince an honest verifier), sound (is a statement is false, a dishonest provider cannot convince an honest verifier), and zero-knowledge (if a statement is true, a verifier cannot learn anything other than the fact the statement is true).

ZKPs have significant advantages in preserving privacy and enhancing security because sensitive information is not shared. It is precisely for this reason that they have become a popular method of preserving privacy in the digital environment. Different versions of ZKP could be applied broadly across society in terms of identify verification, digital voting, and private data sharing.

However, ZKPs used in real life are complex cryptographic constructs, typically used in complex systems such as blockchains, decentralized authentication systems, or secure multi-party computations. Just as with other cryptographically complex algorithms, these are resource intensive, which can make them unsuitable for extensive use on resource-constrained devices (such as smartphones) and pose a scalability issue with large datasets. In addition, implementation requires deep knowledge of cryptography, which is scarce, increasing the risk of errors and vulnerabilities.

[26] It is often demonstrated using the Alibaba Cave Example Zero-Knowledge Inclusion Proofs

A more fundamental issue that can impact ZKPs' usefulness is that the proofs are not about the correctness or the relevance of information. To continue with the example of group membership, ZKPs are not well suited to detect if the membership was mistakenly granted or obtained in bad faith.[27]

The key use cases of ZKPs in the financial markets include compliance and Know Your Client procedures, identity verification, anonymous blockchain transactions, and credit scoring. They could also possibly support privacy in digital transactions (both blockchain and non-blockchain based). In the blockchain space, cryptos like Zcash and layer 2 protocols like Zero-Knowledge Rollups already make use of ZKPs. In the non-blockchain space, ZKPs to support Open Banking/Finance applications holds considerable promise.

Data Masking

The simplest way to preserve privacy is not to expose sensitive data to risky processes or systems in the first place. This is not a solution if given processes or systems must use them with their original values. However, there are several areas where the exact original value of the sensitive fields in a dataset is not needed such as software development and testing, training, demonstrations, proof of concept systems, or business intelligence algorithms that work with market or consumer population segments (which they typically do). In these cases, one could ask, why not simply omit the sensitive fields, say, from a database. The reason for not doing so is that this can break the business logic. For example, if a program expects a credit card number field, then it would probably run into an error if that field was not present.

A better solution is to replace sensitive fields with made-up data. Replacement data can be realistic, meaning it has the same format and structure as the real one, or arbitrary, depending on the use case.[28] In this way, processing can take place normally and privacy is preserved. The replacement itself is called data masking and it is predominantly used to make PII non-identifiable.

While the concept is simple enough, the application of it can be difficult and complex. On a small scale, a simple macro could "sanitize" home addresses out of spreadsheet. Large-scale financial systems work with terabytes of structured, semi-structured and unstructured data, e.g., databases, pdf forms with tables and text, and images. Here it is difficult to determine how and what to mask automatically at the high speeds required, especially in the case of unstructured data. Therefore, dedicated software is often used for data masking that can handle static and dynamic ("on-the-fly") scenarios, a wide variety of data structures and formats, and work efficiently based on predefined policies.

[27] Note that an adversary masquerading as a legitimate verifier can use the proof to infer PII through compounding: gathering small, seemingly unimportant pieces of information over a longer time period to reveal sensitive information when analyzed in context.

[28] In terms of realistic data used for masking, also see the section on synthetic data. It can be argued that data masking using realistic substitutes is the same as creating synthetic data.

Table 2. Output Privacy Technologies compared

	Zero Knowledge Proofs	Data Masking	Differential Privacy	Synthetic Data
What is it?	Methods to prove an assertion without revealing any underlying information	Obfuscating or replacing sensitive data with non-sensitive while preserving the functionality of applications	Introducing disturbances (noise) in datasets so that individual sensitive information is difficult to infer while not significantly altering analytical results	Generating artificial data that is based on real data and preserves its statistical properties
Use Cases	Authentication systems, blockchain, secure voting, digital transactions	Testing and development, proof of concept and demo systems, business intelligence systems	Statistical analysis, trend analysis, ML/AI	AI/ML model training, testing and development
Data Privacy Level	**High** No underlying information is revealed	**Moderate to high** Depending on the effectiveness and extent of masking	**Moderate to high** It can be set to a desired guaranteed level by altering the parameters of the algorithms involved	**Moderate to high** No real data is involved although inference risks to benchmarked data remain.
Overheads	**Moderate to high** Depending on the protocol	**Low** Once masking policies are set	**Moderate to high** Due to variability of noise addition and desired guarantee	**Moderate to high** Generating high quality synthetic data can be resource intensive
Scalability	**Moderate** Issues with complex protocols and systems with large number of verifications	**High** The masking operations are generally simple and act on small subsets of data	**Moderate** Managing privacy levels (“budgets”) while still ensuring data utility can be difficult at scale	**High** Does not need continues access to real data and can processing can be easily parallelized
Less suited for	Resource-constrained devices, and fast and frequent verifications Any use case that cannot be expressed as proving a question	Application where the exact original data values are needed	Application where the exact original data values are needed	Application where the exact original data values are needed

There are various ways to mask sensitive data but all fall into two categories: (i) irreversible methods that break the link between the data used for masking and the real one; and (ii) reversible methods that only make it difficult but not impossible to learn the real data. For example, masking sensitive fields in pdf documents can be done either by replacing the sensitive data or by applying an actual mask (e.g., a black rectangle) over it (which is often called redacting) and protecting the file against editing. A skilled and well-resourced adversary could remove the mask by hacking the pdf file, and so learn the original data. However, if it is completely replaced, then this cannot happen.[29] The same applies for more sophisticated masking methods as well, such as format-preserving encryption and data tokenization. In theory, the only truly irreversible data masking is to erase the original sensitive data after it was replaced with random data. This is often not practical.

Format-preserving encryption, as the name suggests, outputs encrypted data in the same format as the input data in the clear. For example, numbers remain numbers and words remain words.[30] This is especially important for data masking to work with legacy systems where some specific formats are expected, and it is difficult to make them work with something different.

Data tokenization is a reversible replacement method, but only with access to the system that manages the tokens. Without it, it is not feasible to infer the original data. It can be seen as a form of encryption. In tokenization, instead of sensitive data, non-sensitive data (tokens) are issued, which have no meaning on its own. The tokens refer to the sensitive data, but the reference only works through the token management system. In other words, malicious parties trying to infer the underlying data need to obtain access to (e.g., hack) the token management system first. Because of this, the cybersecurity of the token management system is crucial in maintaining privacy.

Whether and when it is appropriate to use irreversible versus reversible data masking depends on the use case and the risk profile of the processing environment and parties involved. Generally, in third-party processing scenarios it makes more sense to exchange datasets with sensitive fields irreversibly masked. In-company or in-group processing can benefit from reversible masking, and sometimes is the only way that makes economic sense. For example, a large database could use format-preserving encryption to mask sensitive fields in situ and avoid the considerable complexities and costs of creating separate copies or views or restructuring the database.[31] By controlling access to the encryption keys, access to PII can be restricted on a need-to-know basis while having the assurance that even if the database was breached, sensitive information would be protected.

Differential Privacy

Differential privacy is not an algorithm but a mathematical promise that guarantees that when, looking at the output of certain data: (a) you cannot tell whether any single data point is included in the original dataset; or (b) you cannot learn anything more about the data subject than if it was absent from the data

[29] Note that if the field is encrypted as a way of masking, then access to the key makes the masking reversible.

[30] Modern cryptographic systems scramble completely the input format (which is better for security), hence the distinction.

[31] Why not encrypt the entire database, and use granular access controls over tables, views, and fields instead? These methods are used as well but come with additional issues. Wholesale encryption of large databases is costly and granular access controls are difficult to maintain and become ineffective the moment the database is breached.

set. In other words, if some action was performed on two nearly identical data sets that differ in only one data point, that difference should not significantly alter the outcome of the action. Differential privacy provides a mathematical guarantee of privacy, regardless of the resources of a third-party that may want to access data (e.g., unlimited computing power, knowledge of algorithms etc.)

The most common way that differential privacy aims to achieve the promise of privacy is through adding "noise". This is inaccurate information that is injected into the data set in amounts too small to disturb observation of major trends. Noise can be added during aggregation (global differential privacy) or at individual levels before aggregation (local differential privacy). There are two main mechanisms for adding noise into data sets, the Laplace and Gaussian mechanisms, although other mechanisms can be used depending on the need.[32]

The amount of noise added is controlled by a parameter known as epsilon and defines the level of privacy that will be achieved (called the "Privacy Budget")[33]. A low value of epsilon means stronger privacy, but less accurate data given that more inaccurate data will disturb the dataset to a larger extent. Where epsilon is the only parameter and at zero, we refer to this as "pure" differential privacy. However, to ensure data has utility, it might be useful to add a second parameter that allows users to understand the likelihood that privacy could be compromised. This second parameter is known as delta and controls the probability of an extreme privacy breach, and lets the user know the probability of epsilon not holding. The smaller the delta, the lower the chances of privacy being compromised but again, at the expense of data utility. Where delta is a second parameter, we call this "approximate" differential privacy. [34]

While providing privacy in large data sets, it can bring inaccuracy to smaller data sets as the contribution of each data point is larger and so setting a minimum query set size is important. Additionally running an algorithm multiple times over a database increases the epsilon value and so reduces the level of privacy therefore limiting how many times a reviewer can run particular types of queries prevents them from being able to figure out whether a given individual is present in the data set.

Differential privacy already has commercial applications, is well established in the market and may have promising use cases in financial services[35]. It is also already used in the public sector to protect sensitive data.[36] Going forward, credit ratings agencies could use differential privacy by adding noise to data they have on their users to conduct demographic analysis (like average credit scores of a demographic group). Differential privacy for regulatory authorities could be most useful when sharing broad statistics such as general trends in complaints data, consumer spending patterns, consumer financial behavior, or economic surveys where directionality is more important than granularity.

[32] The Laplace mechanism draws from the Laplace distribution which is useful as it can model data where changes occur around a central value, modelling mostly small disturbances and occasional larger ones. The Gaussian mechanism draws from the Normal distribution which is useful as it allows for the addition of noise in a smoother more continuous manner.

[33] More formally, it is the maximum distance between a query on a given database and the same query on a second database.

[34] To determine how much the output of an action can change with a single datapoint's modification L1 sensitivity is most often used in pure differential privacy which measures the difference using the sum of absolute differences, while L2 sensitivity is more commonly used in approximate differential privacy using the square root of the sum of squared differences.

[35] Differential Privacy Overview (apple.com)

[36] Key Parameters Set to Protect Privacy in 2020 Census Results

Synthetic Data

Synthetic data is data which has, usually, been generated algorithmically. It is artificial data that aims to mimic or reflect real-world data, without having some of the issues that come along with using real-world data, like revealing information on the underlying data. Synthetic data can be split across two main types, partially synthetic data (where small parts of real-world data are substituted for synthetic data), and fully synthetic data (where all parts of the data are synthesized).

The utility of synthetic data extends beyond privacy and includes the creation of new data in areas where data are not readily available and the augmentation of data where data are incomplete or scarce. From a privacy perspective, the utility is the use of data that are not merely anonymized or sanitized, but rather lack any use of raw data while retaining and preserving their statistical properties.

However, in order to generate synthetic data, you usually need a real data set against which the synthetic data is benchmarked. This is called fidelity, the statistical similarity of a synthetic dataset to the input real data. The importance you place on fidelity depends on what you want to use the data for. For example, in order to build a model that reflects the creditworthiness of the average British borrower, it makes sense to have synthetic data that closely reflects the properties and behavior of average British borrowers. If you just want data to test system resilience, fidelity is much less important.

There are several ways that synthetic data is generated and three of the most popular are hand-engineered methods,[37] agent-based modelling,[38] and deep learning.[39] Given that most types of synthetic data are based on real world data sets, synthetic data is not automatically privacy preserving. The notional privacy of synthetic data may initially rest on its statistical similarity to real world data. Often (but not always) the most useful data are those that have the highest fidelity.

However, the moment you create synthetic data, it already is not real-world data and may not accurately reflect real-world behavior. The data may be oversimplified, contain biases, or there could be a lack of understanding of the underlying data. In order to manage these risks, you can ensure that synthetic data closely reflects real-world data. However, the more we increase fidelity, the more we decrease privacy.

[37] Hand-engineered methods rely on the manual creation of synthetic data. For example, in rules-based generation, data are subject to pre-defined rules like certain minimum, and maximum values. Any relationships or statistical patterns that you desire are pre-defined. An output is linked to a given input. The resultant synthetic data is only as good as the pre-defined set of rules. By adding multiple rules, you may create overlapping or conflicting rules, as well as issues around scalability. Rules-based data generation is static and reflects the initial input, but real-world data are continuously changing so rules must ideally reflect this drift.

[38] Agent-based modelling relies on models that explain observed behavior among different agents (data points) and then reproduces random data using the same model. Analysis of real-world data starts by understanding the interaction of individual elements within a data set, which means synthetic data more closely reflects the way elements within a dataset might behave in reality. Here you get different outcomes even with the same initial inputs, which is very different from the rules-based approach.

[39] AI generated synthetic data can be the most complex but also the most useful and includes data generated by generative adversarial networks (GANs), and variational autoencoders (VAEs). GANs generate synthetic data using two neural networks in a competitive setting: a generator that creates data and a discriminator that evaluates it. The generator aims to produce data that is indistinguishable from real data, while the discriminator tries to detect if the data is real or not. VAEs encode data into a simplified, compressed form and then expand it back out, aiming to maintain just the general features.

We can increase privacy, but at the expense of fidelity which can lead to model drift over time and make synthetic data less useful for many (but not all) uses.

Additionally, re-identification is possible as, even though the data are synthetic, they likely mirror real datasets. Therefore, the risk exists that the presence of an individual data point could be inferred (membership inference), or an attribute about an individual data point could be learned without knowing whether an individual data point is present in the dataset (attribute inference). Furthermore, some AI based models on synthetic data generation have been shown to have retained a "memory" of the raw data on which synthetic data were created which can provide additional vectors for uncovering potentially sensitive data.

In order for synthetic data to be more private, one approach is to remove outliers or data points with uniquely identifying features. If a real dataset has certain outliers which are mimicked by synthetic data, then it might be possible to make inferences about the underlying data. A second approach is to either make the dataset larger which can make it more difficult to identify individual data points, or to reduce the number of variables (feature reduction) included in the data which can reduce the dimensional space that can be linked back to real-world data by only focusing on necessary or important variables. A third approach is to make the data differentially private by adding noise into the data generation process.

Synthetic data is becoming increasingly popular and understanding some of the pros and cons can help supervisors better manage risks while allowing sensible innovation to occur. When used appropriately, synthetic data has the potential to improve outcomes for markets, consumers, and potentially financial stability. We can't think of synthetic data as appropriate in all scenarios, nor can we dismiss it because of some of the risks it generates, rather we should help guide its use where it is most likely to improve outcomes for markets and users.

In financial services we might see synthetic data being used to improve fraud detection and prevention by creating realistic transaction profiles to train fraud detecting algorithms. Synthetic data might also be used by banks, regulators, or international organizations to simulate stress testing. It could be used to improve customer chatbots, or to test new products or services before they are launched. However, financial sector authorities might face several challenges when dealing with synthetic data. When developing synthetic data themselves, authorities will need access to data, expertise, as well as the ability to absorb the costs associated with synthetic data generation. This is likely to be particularly challenging for resource constrained authorities, or authorities in jurisdictions with smaller financial markets where access to relevant real-world data as well as relevant skills and resources are limited.

Box 1. The Digital Sandbox and Synthetic Data

Since its launch in October 2021, the UK Financial Conduct Authority (FCA), in conjunction with the City of London Corporation, has expanded the Digital Sandbox into a permanent initiative, to accelerate data-driven innovation. By providing access to synthetic data and application programming interfaces (APIs) through an API

marketplace, the platform enables firms to safely test and refine new technologies in a controlled environment, overcoming challenges related to real-world data access and regulatory constraints.

Growth of such innovations rests on the ability to train these models on large training datasets. However, real financial data contains sensitive and confidential information and so, outside of proprietary data that start-ups often lack, it cannot be used to train new algorithms. Data anonymization and/or pseudonymization is one way to resolve this problem where identifying elements of data are removed to prevent identification of the data subject. However, anonymized/pseudonymized data can potentially be reverse engineered or combined with other data sets to identify data subjects within the underlying data. Synthetic data provides an alternative approach which was adopted by the UK FCA, who generated it through a 'DataSprint' that bought together the public and private sectors to explore different methodologies to generate the synthetic data.

In March 2023, the FCA established the Synthetic Data Expert Group (SDEG), consisting of 21 experts from financial services, public sector, data and technology vendors, and consumer groups. The group focuses on identifying key issues, use cases, and challenges surrounding synthetic data in the UK financial markets. As the understanding of data generation techniques and best practices continues to evolve, the journey moves forward, with the SDEG helping to shape the responsible use of synthetic data. Through collaboration across industry, regulators, academia, and civil society, the SDEG promotes the responsible use of synthetic data, driving digital market evolution and supporting the FCA's efforts in achieving positive outcomes and digital transformation.

While synthetic data provides a viable option for regulatory authorities to develop data-based tools, its generation is technically challenging and resource intensive and may not be feasible for many regulatory authorities. However, as the technology evolves and becomes more accessible, it holds the potential to be increasingly utilized by regulators in the future.

_____

Source: Digital Sandbox | FCA

Source: Report: Using Synthetic Data in Financial Services | FCA

Considerations for Supervisors

Privacy technologies have the potential to unlock innovation by inducing greater trust, thereby allowing data to flow better within and across borders, and hence, for more data to be generated. They are a potentially transformative set of technologies that could improve outcomes for markets and consumers.

However, they have not been used at scale and there are outstanding questions about deployment at scale. Some privacy technologies could provide a false sense of security, and many technologies will need a combination of approaches for true privacy to be delivered. Defining true privacy in the digital age can be challenging, there may be frictions between efficiency and regulatory requirements and greater data sharing and flows may also generate cyber risks. Data privacy is not the same as data security and privacy technologies do not replace data security controls or more generally, cybersecurity controls. Therefore, it is useful to regard cybersecurity as a foundation for the proper use of privacy technologies when dealing with personal or sensitive data.

Outreach, engagement and understanding trade-offs

It is important that supervisory authorities have a clear understanding of how privacy technologies are developing within their markets as this can inform the pace of change in policies and assessment of the need for new policies. Where firms are significantly experimenting with privacy technology within markets, authorities may need prioritize their engagement to ensure regulation and supervision are fit for purpose. This is particularly true in instances where data is being shared using privacy technologies or where privacy technologies are a fundamental part of financial services, for example their use in digital payments or crypto. It is also important to understand which privacy technologies are likely to permeate throughout financial markets, and where certain public key encryption-based technologies could be impacted by emerging technologies such as quantum computing.

Outreach and engagement do not require any structural changes, and often, existing supervisory structures should allow supervisors to gather insights and share them with all relevant stakeholders. This means engaging directly with supervised firms where supervisors already have pre-existing relationships. In the case of many markets, a significant amount of digital innovation occurs through authorized firms (banks, e-money providers, insurers etc.) Regulators have also made use of demonstration days which allow firms, both regulated and unregulated, to showcase ideas and innovations to relevant teams within the regulatory authority. This is relatively resource light, allowing regulatory and supervisory authorities to quickly identify trends and levels of development, and is especially important for privacy technologies where development is primarily occurring outside the regulated sector.

Some authorities have developed Innovation Hubs to conduct outreach and engagement. These may supplement existing supervisory structures by providing specialist fintech expertise or by leading on the outreach themselves, which may be particularly effective in conditions where supervisors have the necessary budgets and resources, and fintech developments are occurring rapidly. These units advise on how existing regulation applies to innovative business models and emerging technologies which is

important given that privacy technologies will challenge existing rules around transparency, disclosure, and data protection.

As the level of innovation of privacy technologies in the private sector deepens, making further technical insights necessary, sandboxes could be used to plug those gaps in knowledge.[40] Digital sandboxes may be beneficial for certain privacy technologies, such as federated learning, as they can provide firms with access to data that is necessary to conduct tests with regulatory oversight. Likewise, digital sandboxes can also help facilitate the dissemination of synthetic data, particularly where regulatory authorities have been involved in its development. Synthetic data could enable federated learning to grow more quickly and with less regulatory risk. It could also provide proofs-of-concept for homomorphic encryption, particularly fully homomorphic which carries greater benefits but also possibly greater challenges.

While most financial authorities are technology neutral, fintech developments are challenging existing assumptions. For example, the Basel Committee's proposal to treat crypto assets deployed on permissionless networks differently to those on permissioned networks reflects a position on technology, while regulatory authorities in Japan and the United Arab Emirates have banned the listing of privacy technologies given risks around financial integrity. New business models and emerging technologies will further challenge the notion of technology neutrality and certain privacy technologies might require specific approaches.[41] In this regard, the use of Policy Sandboxes could usefully support consultation exercises. This could also be supported by supervisors issuing advisories to supervised firms with respect to privacy technologies to increase awareness of their benefits and risks, providing opportunities for greater experimentation.

Coordination and cooperation

An integrated approach that considers privacy, competition, innovation, financial stability and data protection frameworks is important.[42] Ensuring compliance with data protection frameworks, including rules on data privacy, is primarily the remit of data protection authorities, but financial supervisors have a stake too. Data breaches and consumer complaints involving improper handling of personal data are operational risk events that can cause reputation risk, and in extreme cases may impact financial safety and soundness and even financial stability. Often such events are indicative of weaknesses in cyber risk management at institutions, a growing concern among financial supervisors.

In the financial sector, consumer protection agencies or conduct authorities could become aware of (potential) data breaches even before a report is filed by the institution, e.g., due to a sudden increase in consumer complaints. Prudential supervisors usually do not receive complaints (data) directly. An

[40] Currently, there is little evidence to suggest that Regulatory Sandboxes (Product Testing Sandboxes) are able to deliver on the objectives set out by authorities, but they have the potential to be useful in certain instances to allow supervisors to better understand privacy technologies in practice. See Institutional Arrangements for Fintech Regulation: Supervisory Monitoring (imf.org)

[41] For example, if ZKPs are shown to provide privacy for digital transactions while retaining compliance for AML/CFT purposes, authorities might mandate their use by payment service providers.

[42] The Economics and Implications of Data: An Integrated Perspective

agreement that makes it possible to share such data from financial conduct and data protection authorities on an aggregated basis can be helpful in detecting a worsening cybersecurity stance at supervised institutions. Of course, unusual patterns in complaints could be caused by various reasons other than data privacy or cybersecurity issues, but then often there is a link, which may be worth investigating.

Many jurisdictions are developing structures to improve domestic, inter-authority coordination, e.g., the South African Intergovernmental Fintech Working Group and the British Digital Regulation Cooperation Forum. Both of these bring together financial and nonfinancial regulatory authorities. Collaboration mechanisms, whether formal or informal, are set to become increasingly indispensable in the digital economy given the reliance digital financial services has on big data and connectivity, and certainly when considering the impact of privacy technologies in financial services.

It is important that supervisors are able to work closely with their peers domestically and globally to understand the emerging impacts of privacy technologies in financial services and their implications for regulation and supervision. Fintech cooperation agreements, when utilized correctly with efforts by all signatories, provide one avenue for formalizing information sharing which could be particularly useful in exchanging information and ideas on how to best use, and supervise, privacy technologies. Their relative informal nature allows for faster information sharing between specialist units of experts, which is important given the fast-paced evolution of fintech. However, without a legal underpinning they may not be appropriate to discuss specific cases or exchanging sensitive or confidential information.

Understanding and managing cyber implications

Data privacy and data security focus on different aspects of how data are handled and protected. The subject of data privacy is the personal data of individuals (data subjects) and refers to the collection, use, storage, and deletion of personal data in compliance with privacy regulation. Data security is a broader concept that has wider scope of application. Its subject is all types of data, and its goals include not only preserving data confidentiality, as is the case for data privacy, but also the integrity and availability of data. It is also easy to see that while preserving the confidentiality of personal data is necessary to ensuring data privacy, it is not sufficient in itself.[43]

Cybersecurity provides a foundational layer of protection for the use of privacy technologies. Even if a dataset is made less sensitive using privacy technologies, it usually still needs to be protected to a degree, and it can be used in systems or flow through networks that could require high levels of protection for other reasons.[44] Authentication and authorization systems, firewalls and anti-malware systems all act in concert to limit the risk of unauthorized access to datasets and systems that process them. Certain privacy technologies can also be embedded in cybersecurity infrastructures, such as authentication and key management systems.

[43] For example, it could have been collected illegally in the first place.

[44] For example, access to a dataset about real financial transactions should be limited even if personal data is masked, just like the financial analysis system on which it may reside, or the corporate network to which this system is connected to.

The interrelationships between privacy technologies and cybersecurity can be complex and difficult to unravel. This is on top of the complexities of privacy technologies themselves, often employing advanced cryptography that is an esoteric subject in itself. Arguably, supervisors can eschew most of this complexity by focusing on requirements and outcomes in a technology agnostic approach. In terms of regulation and supervisory expectations, it is unnecessary to require any specific privacy technology. Instead, the desired privacy preserving properties of processing should be specified and supervised institutions allowed to select the technologies that achieve those properties and are the easiest to adopt considering their circumstances. In addition, cybersecurity or technology risk related requirements will take care of the broader control environment. If there is data protection legislation and technology risk regulation in effect, then often it is sufficient to use soft power approaches, such as speeches, outreach, or circulars to provide more guidance on the topic. In terms of supervision, supervisors should aim to understand the nature and extent of using privacy technologies, perhaps focusing on specific cases where personal data is processed. An outcome-based approach works well. The key determination to make is whether the steps and the results of such processing comply with data protection and cybersecurity requirements irrespective of the underlying technologies. Attention should be paid to avoid a false sense of security that may emerge if privacy technologies are not sufficiently backed by a strong cybersecurity control environment.

Conclusion

The growth of the digital economy, including fintech, offers opportunities for economic growth and development. Active and open participation in the digital economy is imperative to ensure all users and the wider economy derives maximum benefits. However, participation in the digital economy is not without risks. Users may withhold or withdraw their active participation in the digital economy for privacy reasons, including limiting digital footprints, shielding personal data from third parties and governments. Firms may also limit their participation to comply with data protection laws, or fear of sharing information with regulatory authorities that could potentially penalize them or put their intellectual property at risk.

Two elements can help build trust in the digital economy: regulation, and technology. First, financial regulation, built on broader national data protection frameworks, can support policy-based infrastructure to facilitate growth in the digital economy. For example, consent based data portability (like Open Banking) can provide data subjects with control of their data and determine what, how much, and with whom they share. Second, privacy technologies that aim to support confidentiality can help build trust in the digital economy. Input privacy technologies that reduce data access like homomorphic encryption, secure multi-party computation, differential privacy, and federated learning, and output privacy technologies that reduce inference from shared data like zero-knowledge proofs, data masking, and synthetic data are in various stages of development and offer considerable promise. They could form the basis of trust in the digital economy, allowing data subjects to interact more with the digital economy with assurance of privacy.

Privacy technologies provide a promising avenue for allowing greater data portability while preserving privacy and complying with national data frameworks. They provide increasing opportunities for regulated firms as well as regulatory authorities in terms of improving efficiency and delivering use cases that can improve outcomes for markets and consumers. The use of privacy technologies relies, fundamentally, on the availability of robust national data frameworks.

Data security, and more broadly cybersecurity, can be considered the base layer, on which data privacy can be built which can facilitate data portability that can unlock tremendous value for firms and regulatory authorities. However, privacy technologies are at various stages of development and while some are theoretically appealing, their practical use might currently be limited. Others are already production grade and being use in the public and private sector. Regulatory authorities must be attuned to the risks of privacy technologies, and they cannot be seen as a panacea for privacy preservation.

There are various trade-offs for different privacy technologies. Generally, the more privacy is preserved, the less data is useful for additional value extraction beyond its original use. This can be an economic disincentive for firms, and so regulation and supervision play an important role to balance the interests of all stakeholders. This involves trade-offs as well: how much privacy can be mandated while ensuring legitimate oversight and control (e.g., AML/CFT) and fostering an innovative and competitive financial system.

On a more operational level, trade-offs can be made depending on the overhead and scalability attributes of the various technologies versus the scope of use and targeted privacy preservation, while staying within the boundaries set by regulatory requirements. For example, where to use zero knowledge proofs or homomorphic encryption and to what extent, how to tune the parameters in differential privacy applications, and so on.

We propose three key considerations for supervisors to consider ensuring they are able to support the opportunities that privacy technologies could generate. Supervisors should understand trade-offs of privacy technologies through outreach and engagement; ensure domestic collaboration and international cooperation to improve knowledge sharing, regulatory certainty, and clarity of mandates; and consider privacy technologies as complementary to a base layer of cybersecurity controls.

References

Agur, I., and others (2023). *Bank Competition and Household Policy Privacy in a Digital Payment Monopoly*. IMF Working Papers 23/121. International Monetary Fund. Retrieved from Bank Competition and Household Privacy in a Digital Payment Monopoly

Armantier, O., Doerr, S., Frost, J., Fuster, A., & Shue, K. (2024). *Nothing to Hide? Gender and Age Differences in Willingness to Share Financial Data*. BIS Working Papers No. 1187. Bank for International Settlements. Retrieved from https://www.bis.org/publ/work1187.pdf.

Apple Inc. (2017). *Differential Privacy Overview*. Retrieved from https://www.apple.com/privacy/docs/Differential_Privacy_Overview.pdf

Babina, T., Bahaj, S., Buchak, G., De Marco, F., Foulis, A., Gornall, W., Mazzola, F., & Yu, T. (2024). *Customer Data Access and Fintech Entry: Early Evidence from Open Banking* (Bank of England Staff Working Paper No. 1059). Bank of England. Retrieved from https://www.bankofengland.co.uk/-/media/boe/files/working-paper/2024/customer-data-access-and-fintech-entry-early-evidence-from-open-banking.pdf

Bains, P., & Wu, C. (2023). *Institutional Arrangements for Fintech Regulation: Supervisory Monitoring* (Fintech Note No. 2023/004). International Monetary Fund. Retrieved from https://www.imf.org/en/Publications/fintech-notes/Issues/2023/06/23/Institutional-Arrangements-for-Fintech-Regulation-Supervisory-Monitoring-534291

Carrière-Swallow, Yan and Vikram Haksar, (2019). *"The Economics and Implications of Data: An Integrated Perspective*," IMF Departmental Paper 19/16. https://www.imf.org/en/Publications/Departmental-Papers-Policy-Papers/Issues/2019/09/20/The-Economics-and-Implications-of-Data-An-Integrated-Perspective-48596

CGAP. (2024) "*Key Considerations for Open Finance | CGAP Research & Publications".* https://www.cgap.org/research/publication/key-considerations-for-open-finance

Dwork, C., McSherry, F., Nissim, K., Smith, A. (2006). *Calibrating Noise to Sensitivity in Private Data Analysis.* In: Halevi, S., Rabin, T. (eds) Theory of Cryptography. TCC 2006. Lecture Notes in Computer Science, vol 3876. Springer, Berlin, Heidelberg. https://doi.org/10.1007/11681878_14

Dwork, C., & Roth, A. (2014). *The Algorithmic Foundations of Differential Privacy*. Retrieved from https://www.cis.upenn.edu/~aaroth/Papers/privacybook.pdf.

Fairfield, J. A. T., & Engel, C. (2015). Privacy as a Public Good. *Duke Law Journal*, 65(3), 385-457. Retrieved from https://scholarship.law.duke.edu/dlj/vol65/iss3/1/.

Fleckenstein, M., Obaidi, A., & Tryfona, N. (2023*). A Review of Data Valuation Approaches and Building and Scoring a Data Valuation Model.* Harvard Data Science Review, 5(1). Retrieved from https://doi.org/10.1162/99608f92.c18db966

Goldfarf, A., and others. (2023). *"The Economics of Digital Privacy".* National Bureau of Economic Research. Working Paper 30943. Retrieved from https://www.nber.org/system/files/working_papers/w30943/w30943.pdf

Haksar, Vikram and others, (2021). *"Toward a Global Approach to Data in the Digital Age,"* IMF Staff Discussion Note 2021/005. https://www.elibrary.imf.org/view/journals/006/2021/005/006.2021.issue-005-en.xml

Hartmann, F., & Kairouz, P. (2023). *Distributed Differential Privacy for Federated Learning.* Google Research Blog. Retrieved from https://research.google/blog/distributed-differential-privacy-for-federated-learning/.

IBM Research. (2022). *Privacy-Preserving Federated Learning in Finance*. Retrieved from https://research.ibm.com/blog/privacy-preserving-federated-learning-finance

Murphy, K., and others (2024). *Central Bank Digital Currency Data Use and Privacy Protection.* (Fintech Note No 2024/004). International Monetary Fund. Retrieved from Central Bank Digital Currency Data Use and Privacy Protection

Organisation for Economic Co-operation and Development. (2013). *OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data*. OECD Publishing. Retrieved from https://www.oecd-ilibrary.org/science-and-technology/oecd-guidelines-on-the-protection-of-privacy-and-transborder-flows-of-personal-data_9789264196391-en

Roessler, B., & DeCew, J. (2023). *Privacy*. In E. N. Zalta & U. Nodelman (Eds.), *The Stanford Encyclopedia of Philosophy* (Spring 2024 Edition). Retrieved from https://plato.stanford.edu/entries/privacy/

Statista. (2022). *Volume of data/information created, captured, copied, and consumed worldwide from 2010 to 2025*. Retrieved from https://www.statista.com/statistics/871513/worldwide-data-created/

Sweeney, L. (2002). *A model for protecting privacy*. *Carnegie Mellon University, Data Privacy Lab*. Retrieved from https://epic.org/wp-content/uploads/privacy/reidentification/Sweeney_Article.pdf

United States Census Bureau. (2021). *Census Bureau Releases 2020 Census Quality Metrics*. Retrieved from https://www.census.gov/newsroom/press-releases/2021/2020-census-key-parameters.html

PUBLICATIONS