Federal Register Notice 87 FR 35250, https://www.federalregister.gov/documents/2022/06/09/2022-12432/request-for-information-on-advancing-privacy-enhancing-technologies, June 9, 2022

Request for Information (RFI) on

Advancing Privacy Enhancing Technologies

Canetti, Ran; Kaptchuk, Gabe; Reyzin, Leonid; Smith, Adam; and Varia, Mayank

Response to the RFI on Advancing Privacy-Enhancing Technologies

Ran Canetti, Gabe Kaptchuk, Leonid Reyzin, Adam Smith, and Mayank Varia
Boston University (academic institution)

July 8, 2022

PETs provide an extremely powerful and versatile toolbox for controlled processing of data, while keeping the data itself hidden. Consequently, PETs allow for a radical rethinking of the possible uses of data, proving fine-grained control over which functions of data are to be exposed (and to whom) and which are to be kept hidden. When properly and responsibly used, this power of PETs can greatly advance democracy for the benefit of all individuals and communities. However, when used carelessly and irresponsibly, PETs can also cause great harm. Therefore, the adoption and development pf PETs should be done in a careful and principled way, while making sure to adequately educate the relevant constituents in the power and dangers of using PETs, and while taking care to provide adequate checks and balances that are appropriate to each and every use case.

More specifically, used properly, PETs can cut the Gordian knot that pits transparency and accountability against privacy and secrecy in a seeemingly inherent way: It provides ways to obtain information that is inherently derived from private data but is crucial for good governance, while keeping the data private and at the same time providing public assurance to the correctness and relevance of the obtained information. But when deployed improperly or recklessly, PETs risk exacerbating inequities, reducing privacy, diminishing transparency and accountability, and undermining public confidence in government agencies and their data products. Furthermore, the power of PETs inevitably leads to a series of policy questions about *when* and *how* PETs should be deployed.

The goal of this report is to advise government agencies regarding the processes and regulations that need to be put in place in order to incentivize responsible development and adoption of PETs, while minimizing the risk of harmful deployments of PETs. For that purpose, we first elaborate on the potential benefits, as well as on the dangers and pitfall. Then we proceed with our recommendations (which are derived from said benefits and dangers). Specifically, we answer four of the questions listed in the Government's Request for Information:

- The democratizing promise of public sector use of PETs to achieve socially beneficial outcomes. (Addressing Q3 about benefits from the adoption of PETs.)
- The individual, group, and societal harms that can arise from misuse of PETs. (Addressing Q7 about risks related to PETs adoption.)
- The benefit of separating the policy questions about a system's goals and values from the technical questions about how to combine PETs to realize it. (Addressing Q8 about existing best practices that are helpful for PETs adoption.)
- A path forward toward developing and deploying PETs that serve the public good. (Addressing Q9 about overcoming sociotechnical and trust barriers, and improving equity.)

About the Authors

We are a group of cryptography and privacy researchers at Boston University who develop and analyze privacy enhancing technologies (PETs). Our areas of expertise cover the mathematical and cryptographic foundations of PETs (including [1–16]), the practical development and deployment of systems involving PETs [17–25], the social and policy contexts in which these technologies are deployed [26–32], and serving on committees that advise national statistics organizations [33, 34]. Members of our group have been at the forefront of the envisioning of PETs, as well as the main algorithmic and analytical tools that build PETs, over the past 30 years.

We are encouraged to see the Office of Science and Technology Policy, the National Artificial Intelligence Initiative Office, and Networking and Information Technology Research and Development take an active interest in furthering the adoption of PETs. We appreciate and welcome the opportunity to provide important information and context about PETs based on our decades of experience in cutting-edge cryptographic research from the early 1990s through today.

Q3. The Promise of Public Sector PETs

In the modern information era, data is powerful. At its best, the explosion in data availability and the massive cost reduction for data analysis promise to bring about a more efficient and democratic world. For governments, data-driven insights can lead to evidence-based policymaking [35] that improves the government's ability to efficiently allocate resources and respond to the needs of its citizens. In the private sector, services offered by companies to people can be effortlessly customized based on individual traits and habits. However, this idealistic promise is undercut by several real-world risks like data breaches, bulk surveillance, and a massive network of data brokers [36] that collectively solidify existing power and exacerbate inequities.

As a consequence, we describe in this section **the democratizing power in *not* having data**. PETs allow for the possibility of having our cake and eating it too: they enable the social benefits of data-driven policymaking while mitigating the harms of social control, censorship, and authoritarianism. These technologies are coming to maturity just as they are most needed and are now fast, simple, and expressive enough to use in many public sector and civil society applications.

To showcase the democratizing and empowering promise of PETs in government and civil society applications, we briefly describe a few of the applications of PETs that we have designed, developed, and deployed [37] together with domain experts and impacted communities. We focus on the human rights and civic values that were addressed using PETs.

- **Transparently describe government actions** [19]: We participated in the design of PETs for the 2020 Decennial Census. The implementation of differential privacy provides strong privacy guarantees, as opposed to heuristic ones. Another major advantage relative to prior disclosure limitation techniques is the ability to *publish the specific methodology performed by the government*, so that researchers can account for this distortion; this has brought the complex debate about statistical disclosure limitation into the public discourse [38].

- **Promote grassroots initiatives and volunteer activism** [20–22]: Jointly with the Boston Women's Workforce Council, we deployed PETs to measure the gender and racial wage gap in the greater Boston area based on real salary data of about 1/6 of the Boston workforce. A key challenge was *designing an easy to understand and easy to use* system so organizations were comfortable contributing their payroll data toward a measurement that they all agreed would have social benefit, but whose value relied on a high rate of voluntary participation.

– **Provide a voice to all** [23]: Inspired by the #MeToo movement, we designed a system together with the non-profit organization Callisto that allows multiple survivors of a repeat sexual offender to be securely matched together so they can act with a stronger voice. Our biggest focus was to design a privacy-respecting system that *protects survivors' assault reports from litigation or investigation discovery requests*, including a subpoena to Callisto itself.

– **Stop bulk surveillance** [25]: PETs allow financial institutions to better detect fraud and money laundering [39], which promotes public safety. Our fraud detection tool developed with Fiverity focuses on the ethical and legal obligations to *safeguard personal information of real people* (i.e., non-fraudulent accounts) even while searching for instances of fraud.

– **Promote constitutional rights** [26, 27]: PETs allow the court system to protect law enforcement's sources and methods while simultaneously providing criminal defendants with their due process rights to inspect and challenge the resulting evidence gathered. A key challenge here was *integrating the use of PETs into the legal process* by which prosecutors and defense attorneys converge on a property that the defense wants the prosecution to prove.

In summary, PETs expand the Pareto frontier of data analyses that can be performed: they enable us to design new methods to conduct statistical analyses that simultaneously *enhance accuracy and quality, respect personal autonomy and consent choices, and enforce data minimization and purpose restrictions.* As a result, PETs offer the promise of conducting socially beneficial and privacy-respecting statistical analyses that would have been difficult—if not impossible—to compute otherwise. But we must also be prudent about the fact that PETs could be used to conduct socially detrimental and privacy-invasive calculations as well; we discuss these risks next.

Q7. The Risk of PETs: *When* Should PETs Be Deployed?

As we have explored above, there are an immense number of opportunities to deploy PETs that will serve the social good, making a clearly positive impact of society. As such, it might be tempting to deploy PETs wherever possible. However, **deploying PETs without careful consideration upfront**, but instead mimicking Silicon Valley's "move fast and break things" mentality, will not only fail to realize the potential of PETs, but **could create tangible *negative* impacts for society**. In particular, we are concerned about the following negative impacts that can arise if PETs are deployed inappropriately or inconsiderately:

– **PETs Deployments that Reduce Privacy.** The title "privacy enhancing technology" is somewhat a misnomer, in that a typical understanding of the word "privacy" might lead the casual observe to assume the wrong properties of a PET. *It is very possible for a PET to reduce the privacy afforded to individuals and organizations when deployed—and, indeed, this will likely be the case if not deployed with care.* The reason for this is threefold:

(1) A single PET on its own may only address a subset of privacy considerations, and may in fact make some participants worse off than if the statistical analysis never occurred. For instance, a PET that provides only input privacy might calculate an accurate statistic but fail to provide disclosure limitation, leading to reconstruction attacks that reveal specific individuals' data.

(2) The baseline privacy expectations of an individual, impacted group, or organization might be very high, such that *any* change to the system might actually reduce their privacy. For example, creating a differentially private release of a database will reduce

the privacy afforded to individuals with personal information in that database if the prior baseline is that no statistical releases from that database occurred.

(3) PETs provide *fine-grained* notions of security. That is, privacy can be defined virtually arbitrarily, within the context of the application, and a PET meeting that definition of privacy can be constructed. As an illustrative example: several proposals for election security include a zero knowledge proof that each voter is included in the final vote tally, but by contrast it would be terrible to produce a zero knowledge proof that a particular citizen voted for a specific candidate.

– **Privacy Washing.** Using a PET to perform a harmful computation does nothing to mitigate the harm. Indeed, using a PET often serves to distract or obfuscate to the casual observer the harm that deploying the computation would cause them. Examples of this concern include proposals to deploy PETs to allow for bulk surveillance of encrypted data [40, 41], scan all cloud uploads against abusive content [42], and perform targeted advertising with private real-time auctions [43–46]. Large parts of the cybersecurity community and civil liberties groups have decried these proposals [47–52]—not because the cryptographic techniques were flawed, but because these specific proposals *obscure from the broader debate over whether these tasks should occur at all, raise questions about whose privacy is being protected and prioritized, and failed to provide public accountability and auditing comparable to existing non-PET systems.*

– **Amplifying Market Concentration.** Many PET deployments (especially secure multi-party computation and differential privacy) allow organizations to learn insights from data to which they would not otherwise have access. While this facilitates some of the best applications of PETs such as data-driven policymaking, sharing insights in this way is not without risks. Industrial organizations can leverage access to these insights to solidify competitive advantages (*e.g.* training better machine learning models using private data); the cost of implementing PETs also adds barriers to entry for small companies. Government bodies can further solidify their control over the population (*e.g.* using private geolocation data to target police deployments in marginalized communities). Such deployments of PETs tend to *further consolidate power in monopolistic businesses or autocratic governments*, rather than uplift democratic principles.

– **Losing the Public Trust.** Trust is built slowly and shattered quickly. Missteps in deploying PETs—either by choosing applications poorly, faulty implementations, or misunderstanding threat models—will likely prevent the carefully considered privacy-enhancing systems that would have the most positive impact. *People and organizations might conflate previous privacy-reducing PETs deployments with future privacy-enhancing PETs deployments*, simply because some of the underlying technical tools are shared and they are worried about rehashing previous debacles. In the worst case, a government misstep in deploying PETs could trigger a backlash that would prevent statistical agencies from access data that they currently have, as the misstep might be seen as a sign of ineptitude.

Q8. Specifying Policy Objectives and Assessing Democratic Values

In this section, we advocate for **regulation that establishes a process for vetting PETs used within the U.S. federal government that transparently specifies its intended functionality, and allows for democratic debate about potential risks and possible mitigations**.

With the excitement surrounding a number of successful deployments of PETs [53], it is sometimes easy to forget that PETs are *tools*—means to an end. *Asking if a PET is appropriate in a specific scenario is like asking if a new kind of asphalt might be helpful for building a highway.* While asphalt is an essential tool, the discussion hides the more important questions such as: who will be harmed and who will be helped by the road's construction? These are questions of *power*, that must be answered in rigorous consultation with impacted communities [54, 55]. Only once these questions have been adequately settled can detailed engineering discussions start.

Put another way, cautious deployment of PETs requires asking the question "when should a PET be deployed." Engaging with this question requires *specifying concretely* what information the envisioned system should release and what it should keep private, and evaluating the merits of such a specification. Importantly and perhaps counterintuitively, *this discussion can happen entirely independently of the actual PET techniques that would be used and combined to realize this specification.* Instead, the discussion should focus on the "input-output" behavior of the system, *i.e.* the information possessed by each party and what each party should learn after interacting with the system. Fortunately, researchers in cryptography and programming languages have already designed the principles necessary to follow this approach.

Existing practice: Specifying ideal functionalities. While discussing the potential deployment of a PET without ever mentioning the PET might seem impossible or too abstract, this is in fact a common practice in cryptography—albeit in a highly formalized context. To capture the wide range of security properties that a system under consideration should achieve (and also the properties that it doesn't claim to have), cryptographers write an *ideal functionality, which is a formal specification of the behavior of a system* that abstracts away all low-level implementation details including which PET, if any, should be used. The ideal functionality can describe the system as though it were purely magic and can leverage fictional, perfectly trustworthy parties. A cryptographer will then prove that some particular set of PETs, properly interleaved, successfully produces a perfect imitation of the ideal functionality, *i.e.* creates something that is "as good as" the idealized specification and moreover is insatiable in the real world.

For example, an (informal) ideal functionality describing an end-to-end encrypted messaging system that scans for abusive content might specify that (1) users can send a messages to a service provider, along with an intended recipient of that message; (2) if the user's message matches a list of prohibited and abusive content, the service provider learns the contents of the message. Otherwise, the service provider learns only the length and timing of the messages, but nothing about the contents; and (3) the service provider can choose to deliver the message to the intended recipient at a time of its choosing. Notice that this specification need never mention the word encryption. Indeed, encryption might be one tool used to realize this functionality, but other approaches could, in principle, exist. Moreover, it is possible to have a policy discussion on the merits of such a system purely by considering the ideal functionality. And by removing any mention of cryptography, the ideal functionality is easier for the public to understand and debate.

A path forward toward the design of PETs. The practice of writing these ideal specifications has technical value [2–4]. Additionally, we believe that it has *tremendous untapped social value.* In some sense, the practice of writing and scrutinizing these specifications independently of the underlying implementation modularizes the process of analyzing deployments of PETs. In particular, it *cleanly separates the social question of what goals a PET deployment should accomplish*—which society should collectively decide through the democratic process—*from the technological question of how these goals should be met*, which can be safely relegated to cryptographers and information

security experts once the values questions have been debated and decided.

There are a number of reasons to promote this two-step process toward the design of PET-enabled systems, including:

- *Receiving feedback and constructive criticism from everyone in a democratic society, including marginalized voices.* History has shown time and again that systems (computational or otherwise) that may have good outcomes for some people can also yield disparate impacts on other groups in the community [54, 55].
- *Inserting useful "friction" into the design process* [56, 57]. As mentioned above in our discussion of risks, making everything computable isn't always a good thing. Whereas people could point in the past toward the "impossibility" of a computation in order to stop it, the existence of PETs forces a more careful introspection upfront about which data analyses will promote rather than hinder human rights.
- *Providing transparency and legitimacy of government actions.* It is particularly important to describe upfront how a PET-enabled system will work (for instance, through a System of Records Notice or similar) because by its very nature the system cannot be radically transparent during use. In fact, auditing the design of the system is just one part of transparency; we must also allow for auditing the execution of the system as we describe next.

Q9. Doing it Right: *How* Should PETs Be Deployed?

If PETs are to serve democratic values, the deployment of PETs must follow democratic principles. Making it easy for people to trust the fidelity of a PET deployment is critical to making the deployment successful. Unfortunately, there are numerous potential pitfalls into which the designer of PET can fall that will quickly shatter any trust that has been developed. To that end, we **suggest a series of guiding principles that can mitigate the potential risks of PET deployments and increase the chances that they will reach their full democratic promise**.

Public Development Process. *Choices about the particular techniques used in a PETs deployment encode the values of its designers.* A particular ideal functionality can be realized with many different combinations of individual privacy enhancing technologies. Choosing one PET over another—or even using different mathematical assumptions or parameters to realize a single PET—can have concrete and disparate consequences for the level of privacy that the PET deployment provides. For example, it is possible to realize secure multiparty computation using either an honest-majority assumption, in which only a colluding majority of the parties participating in the protocol could compromise the privacy of the computation, or a dishonest-majority assumption, in which the data will remain private if even a single party behaves honestly. Similarly, differential privacy puts an upper bound on the amount of information that can leak about an individual during a statistical disclosure, but the parameter choices meaningfully change the qualitative impact of this bound on individuals and communities.

When developing a new PET deployment, it is critical that the choice of techniques and parameters be made *publicly.* NIST has long followed this process with its standardization efforts [58–61], and successfully garnered the trust of both security experts and the general public. At a minimum, having public process ensures that experts can identify vulnerabilities before it is too late. More importantly, disclosing the techniques in use *allows the public to debate if the values encoded in those techniques are appropriate for the application.*

Public Verifiability. The validity of PET deployments should be verifiable—even by an adversarial party who didn't participate in any aspect of the original design or execution, and who has vested interest in demonstrating the opposite. Because PETs are unfamiliar, mathematically complex, and intrinsically *hide something*, a careless deployment of PETs may be used to cast doubt on the results of the system. But, with careful construction, it is possible to *deploy PETs that affirmatively offer public verification of their proper execution and that can provide evidence to counteract any accusation that might discredit the system.*

We believe that there are two crucial types of verifiability that are *necessary* in order to make PET deployments trustworthy:

- **Results Verifiability.** *Privacy must not come at the expense of public oversight and accountability.* Systems that do not leverage PETs can be carefully checked by data scientists, and the results that they produce can be verified by simply looking at the underlying data. In principle, any member of the public who doubts a result could get access to the data and check on their own. Many PET deployments, on the other hand, make this kind of radical transparency impossible for experts and the general public alike—indeed, this is a core motivation for using PETs. As such, PET deployments should use techniques to ensure the results are convincing to *everyone*, like verifiable computation and zero knowledge proofs.

  We remark that this challenge is not new; indeed, election officials have grappled for decades with the challenge of providing public verifiability of the election tallying process, while also guaranteeing voter privacy and protecting against voter coercion. Privacy and anti-coercion threats necessitate that the election system make it impossible to check how each individual person voted, thereby necessitating more onerous methods to provide public verifiability such as risk-limiting audits, cameras to observe the tallying process, and so on. Systems employing PETs must similarly include methods to demonstrate correct behavior and obtain public legitimacy, and new research may be required to achieve this goal.

- **Code Auditability.** It is also necessary to follow good software development practices when developing and deploying PETs: source code should be publicly available and written in a way that meaningful auditing is feasible. The best way to ensure that a software package does not contain security-critical vulnerabilities is to *allow experts ample opportunity to attempt to break it.* Additionally, elegance and simplicity should be the watchwords of PETs software development. If it is too difficult or painful to engage with the software, none of the necessary audits will happen and the organization deploying the PET may have a false sense of security.

Plan for the Long Term. Because PETs limit the ways in which data can be used, it is critical to think through the full lifetime of the data before committing to a particular PET deployment. The use of PETs can be *irrevocable*, so making decisions based only on short-term needs may make it impossible to accomplish critical tasks in the long term. There are three ways in which we anticipate that foresight is necessary before making an initial PET deployment or modifying a system in operation:

- **Understanding Irrevocable Decisions.** *PET systems require careful consideration upfront about which types of extensibility are explicitly desired versus explicitly forbidden.* For example, once data is encrypted, then losing the decryption key is tantamount to deleting the data. This can be a blessing if you *want* to delete the data, for instance when decommissioning a system at its end-of-life, but it can be a substantial issue otherwise. More sophisticated PETs can also embed irreversible decisions. For example, if a survey is conducted using privacy preserving techniques, and a desired correlation is not computed, it may be infeasible to

go back and collect the data again... though once again, we stress that *this may be a desirable feature in some contexts to limit mission and scope creep.* Alternatively, if data is collected using local differential privacy with a particular privacy budget, then later decreasing the amount of noise in the data set (in an effort to get more accurate statistics) would be a complicated process that may require new data collection.

- **System Changes Impacting Public Trust.** When a deployment of PETs does permit changes after initial deployment, any use of this power to change how the system works may *shatter the public's trust in the PET—and in the organization running the privacy-enhanced system.* For example, a government agency might announce that it plans to collect data and only release the results using (central) differential privacy under some published privacy parameter. If the agency then reduces the privacy in the release after data collection has occurred, this would constitute a breach of the norms surrounding the use of the data. As such, *retaining the technical flexibility to change the system later is not always preferable to making irrevocable decisions upfront*, as it brings about its own problems. Recent work has explored how to limit access to this power altogether and to ensure public verifiability for any modifications made to a PET system in use [28, 62].
- **Composing with Future PET Deployments.** As PETs deployments proliferate, *there will be a desire to have deployments interoperate*, which may be impossible without prior planning. For example, the 2020 Census uses differential privacy in its statistical disclosures. The current deployment does not support verifying that the randomness injected into the results was selected honestly—it is technically possible that a malicious data operator manipulated the results by carefully selecting the "random" values given to the disclosure avoidance algorithm. While it is in principle possible to prove that the randomness used was selected independently of the census data (using zero-knowledge to prevent any further privacy leaks), the current system design does not support the composition of the Census differential privacy release and zero-knowledge proofs—and there may be no easy way to retrofit the system later.

The challenges described above are generally not technically complicated; the necessary crypto and privacy tools to avoid these bad outcomes are generally well-understood and sufficiently adaptable to address most realistic deployment scenarios. For instance, a simple countermeasure to the Census issue described above is to post cryptographic commitments to the sensitive data in order to maintain the ability to prove properties about it in the future. Instead, *the rate-limiting resource is human foresight and patience.*

Explainability. Especially for public facing PET deployments, ensuring that the deployment is explained clearly—both to technical and non-technical audiences [63]—is critical to engendering trust. Several recent works in the domain of usable security by us [29,30,64] and others (*e.g.*, [65–68]) have highlighted the need for clearer communication about PETs. Without a clear explanation, people might not trust the system with their data when appropriate, or they might expect more privacy guarantees that the system will actually provide. *Misalignment of expectations and reality will cause distress and mistrust.* As such, developing clear communication about a technological deployment is critical to its long-term success.

Prioritizing Privacy & Functionality over Efficiency. It is prudent to take a realistic evaluation of the efficiency needs of PET deployments. We acknowledge that PETs are not at the point where they can be used for high-performance, low-latency, real-time systems. But those statistical calculations are rare. Particularly for government deployments of PETs, having protocols run for

hours or days may be *completely practical*. And when compared to the legal burdens or outright impossibility of performing a valuable analysis otherwise, the cost of the computing power is often downright cheap.

Research literature into PETs offer a variety of different options for balancing the privacy afforded to the data, the functionality the PET provides, and the concrete efficiency of the resulting system. It is important to be realistic in PET deployments about the level of efficiency that is actually tolerable, as opposed to what would be ideally desired. Prioritizing the efficiency of the system often results is *worse* privacy guarantees and poorer functionality—recall that the most efficient system is one that preserves no privacy at all. Being pragmatic about the acceptable performance in the context of a particular application radically changes the space of PETs that are practical and ready for deployment *today*.

Conclusion

In conclusion, we summarize the most important takeaways from this report:

- Although data is powerful, it can also be toxic. There is thus great value in *not* having access to data so that it cannot be altered, breached, censored, surrepetitiously surveilled, or otherwise abused.
- At its best, PETs offer the promise to improve government transparency and oversight, promote grassroots initiatives, provide a voice to all communities, mitigate bulk surveillance, and promote human rights.
- Deploying PETs carelessly can lead to tangible harms for society, such as reducing the privacy of some people or entrenching existing power structures. We offer several principles to mitigate these risks and allow PET deployments to reach their full promise.
- We recommend distinguishing the objectives of a PET system from its cryptographic implementation. Providing ideal specifications allows everyone to contribute toward the democratic discourse about the goals and values that a PET system *should* imbue.
- All aspects of PET deployments should be verifiable. Furthermore, the verification process should be public and simple-minded, so that even those parties who have a vested interest in demonstrating the opposite will have no choice but to be convinced in the correctness of the results. The public must be able to view the initial PET design and implementation, and separately also verify that it is producing correct results in use.
- When designing PET systems, one must decide upfront which aspects of the system are immutable (and therefore limits scope creep) versus extensible (and therefore adaptable as organizational boundaries and trust decisions change). A "middle ground" here is to allow for changes to be made, but in a manner that enforces public oversight and auditability.
- The ability to explain to the public what a PET algorithm does is crucial, even if it requires sacrificing some accuracy or efficiency. Misalignment of expectations and reality will lead to distress and mistrust.
- For many government applications considering PETs, efficiency is not the main bottleneck to adoption today. Instead the main challenge is human planning and patience throughout the process of PET design, evaluation, deployment, and oversight.

References

[1] Mihir Bellare, Ran Canetti, and Hugo Krawczyk. Keying hash functions for message authentication. In Neal Koblitz, editor, *Advances in Cryptology – CRYPTO'96*, volume 1109 of *Lecture Notes in Computer Science*, pages 1–15, Santa Barbara, CA, USA, August 18–22, 1996. Springer, Heidelberg, Germany.

[2] Ran Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *42nd Annual Symposium on Foundations of Computer Science*, pages 136–145, Las Vegas, NV, USA, October 14–17, 2001. IEEE Computer Society Press.

[3] Ran Canetti. Universally composable security. *J. ACM*, 67(5):28:1–28:94, 2020.

[4] Ran Canetti, Alley Stoughton, and Mayank Varia. EasyUC: Using EasyCrypt to mechanize proofs of universally composable security. In Stephanie Delaune and Limin Jia, editors, *CSF 2019: IEEE 32st Computer Security Foundations Symposium*, pages 167–183, Hoboken, NJ, USA, jun 25-28 2019. IEEE Computer Society Press.

[5] Ran Canetti and Hugo Krawczyk. Analysis of key-exchange protocols and their use for building secure channels. In Birgit Pfitzmann, editor, *Advances in Cryptology – EUROCRYPT 2001*, volume 2045 of *Lecture Notes in Computer Science*, pages 453–474, Innsbruck, Austria, May 6–10, 2001. Springer, Heidelberg, Germany.

[6] Yevgeniy Dodis, Leonid Reyzin, and Adam Smith. Fuzzy extractors: How to generate strong keys from biometrics and other noisy data. In Christian Cachin and Jan Camenisch, editors, *Advances in Cryptology – EUROCRYPT 2004*, volume 3027 of *Lecture Notes in Computer Science*, pages 523–540, Interlaken, Switzerland, May 2–6, 2004. Springer, Heidelberg, Germany.

[7] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In Shai Halevi and Tal Rabin, editors, *TCC 2006: 3rd Theory of Cryptography Conference*, volume 3876 of *Lecture Notes in Computer Science*, pages 265–284, New York, NY, USA, March 4–7, 2006. Springer, Heidelberg, Germany.

[8] Christina Garman, Matthew Green, Gabriel Kaptchuk, Ian Miers, and Michael Rushanan. Dancing on the lip of the volcano: Chosen ciphertext attacks on apple iMessage. In Thorsten Holz and Stefan Savage, editors, *USENIX Security 2016: 25th USENIX Security Symposium*, pages 655–672, Austin, TX, USA, August 10–12, 2016. USENIX Association.

[9] Benjamin Fuller, Mayank Varia, Arkady Yerukhimovich, Emily Shen, Ariel Hamlin, Vijay Gadepally, Richard Shay, John Darby Mitchell, and Robert K. Cunningham. SoK: Cryptographically protected database search. In *2017 IEEE Symposium on Security and Privacy*, pages 172–191, San Jose, CA, USA, May 22–26, 2017. IEEE Computer Society Press.

[10] Joël Alwen, Binyi Chen, Krzysztof Pietrzak, Leonid Reyzin, and Stefano Tessaro. Scrypt is maximally memory-hard. In Jean-Sébastien Coron and Jesper Buus Nielsen, editors, *Advances in Cryptology – EUROCRYPT 2017, Part III*, volume 10212 of *Lecture Notes in Computer Science*, pages 33–62, Paris, France, April 30 – May 4, 2017. Springer, Heidelberg, Germany.

[11] Arka Rai Choudhuri, Matthew Green, Abhishek Jain, Gabriel Kaptchuk, and Ian Miers. Fairness in an unfair world: Fair multiparty computation from public bulletin boards. In Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu, editors, *ACM CCS 2017: 24th Conference on Computer and Communications Security*, pages 719–728, Dallas, TX, USA, October 31 – November 2, 2017. ACM Press.

[12] Ran Canetti, Yilei Chen, Leonid Reyzin, and Ron D. Rothblum. Fiat-Shamir and correlation intractability from strong KDM-secure encryption. In Jesper Buus Nielsen and Vincent Rijmen, editors, *Advances in Cryptology – EUROCRYPT 2018, Part I*, volume 10820 of *Lecture Notes in Computer Science*, pages 91–122, Tel Aviv, Israel, April 29 – May 3, 2018. Springer, Heidelberg, Germany.

[13] Albert Cheu, Adam D. Smith, Jonathan Ullman, David Zeber, and Maxim Zhilyaev. Distributed differential privacy via shuffling. In Yuval Ishai and Vincent Rijmen, editors, *Advances in Cryptology – EUROCRYPT 2019, Part I*, volume 11476 of *Lecture Notes in Computer Science*, pages 375–403, Darmstadt, Germany, May 19–23, 2019. Springer, Heidelberg, Germany.

[14] Ran Canetti, Aloni Cohen, Nishanth Dikkala, Govind Ramnarayan, Sarah Scheffler, and Adam D. Smith. From soft classifiers to hard decisions: How fair can we be? In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, pages 309–318. ACM, 2019.

[15] Ran Canetti, Yael Tauman Kalai, Anna Lysyanskaya, Ronald L. Rivest, Adi Shamir, Emily Shen, Ari Trachtenberg, Mayank Varia, and Daniel J. Weitzner. Privacy-preserving automated exposure notification. Cryptology ePrint Archive, Report 2020/863, 2020. `https://eprint.iacr.org/2020/863`.

[16] Ran Canetti, Benjamin Fuller, Omer Paneth, Leonid Reyzin, and Adam D. Smith. Reusable fuzzy extractors for low-entropy distributions. *Journal of Cryptology*, 34(1):2, January 2021.

[17] Nikolaj Volgushev, Malte Schwarzkopf, Ben Getchell, Mayank Varia, Andrei Lapets, and Azer Bestavros. Conclave: Secure multi-party computation on big data. In *Proceedings of the Fourteenth EuroSys Conference*, pages 3:1–3:18. ACM, 2019. https://arxiv.org/pdf/1902.06288.pdf.

[18] John Liagouris, Vasiliki Kalavri, Muhammad Faisal, and Mayank Varia. Secrecy: Secure collaborative analytics on secret-shared data. *arXiv*, abs/2102.01048, 2021. https://arxiv.org/abs/2102.01048.

[19] Daniel Alabi, Audra McMillan, Jayshree Sarathy, Adam Smith, and Salil Vadhan. Differentially private simple linear regression. *Proceedings on Privacy Enhancing Technologies*, 2022(2):184–204, 2022.

[20] Boston Women's Workforce Council. Data privacy. https://thebwwc.org/mpc.

[21] Museums Moving Forward. Data study. https://museumsmovingforward.com/data-study.

[22] Greater Boston Chamber of Commerce. Pacesetters. https://bostonchamber.com/networks/pacesetters/.

[23] Anjana Rajan, Lucy Qin, David W. Archer, Dan Boneh, Tancrède Lepoint, and Mayank Varia. Callisto: A cryptographic approach to detecting serial perpetrators of sexual misconduct. In *Proceedings of the 1st ACM SIGCAS Conference on Computing and Sustainable Societies*, pages 49:1–49:4. ACM, 2018. https://www.projectcallisto.org/callisto-cryptographic-approach.pdf.

[24] Fireblocks. About fireblocks. https://www.fireblocks.com/about.

[25] Fiverity. A holistic approach to digital fraud detection. https://www.fiverity.com.

[26] Kenneth A. Bamberger, Ran Canetti, Shafi Goldwasser, Rebecca Wexler, and Evan J. Zimmerman. Verification dilemmas in law and the promise of zero-knowledge proofs. *Berkeley Technology Law Journal*, 37(1), 2022. Available at SSRN: https://ssrn.com/abstract=3781082.

[27] Dor Bitan, Ran Canetti, Shafi Goldwasser, and Rebecca Wexler. Using zero-knowledge to reconcile law enforcement secrecy and fair trial rights in criminal cases. In *2nd ACM Symposium on Computer Science and Law (to appear)*, 2022. Available at SSRN: https://ssrn.com/abstract=4074315.

[28] Matthew Green, Gabriel Kaptchuk, and Gijs Van Laer. Abuse resistant law enforcement access systems. In Anne Canteaut and François-Xavier Standaert, editors, *Advances in Cryptology – EUROCRYPT 2021, Part III*, volume 12698 of *Lecture Notes in Computer Science*, pages 553–583, Zagreb, Croatia, October 17–21, 2021. Springer, Heidelberg, Germany.

[29] Rachel Cummings, Gabriel Kaptchuk, and Elissa M. Redmiles. "I need a better description": An investigation into user expectations for differential privacy. In Giovanni Vigna and Elaine Shi, editors, *ACM CCS 2021: 28th Conference on Computer and Communications Security*, pages 3037–3052, Virtual Event, Republic of Korea, November 15–19, 2021. ACM Press.

[30] Gabriel Kaptchuk, Daniel G. Goldstein, Eszter Hargittai, Jake M. Hofman, and Elissa M. Redmiles. How good is good enough for COVID19 apps? The influence of benefits, accuracy, and privacy on willingness to adopt. *Digital Threats: Research and Practice*, 2022. Available at https://arxiv.org/pdf/2005.04343.pdf.

[31] Sarah Scheffler and Mayank Varia. Protecting cryptography against compelled self-incrimination. In Michael Bailey and Rachel Greenstadt, editors, *USENIX Security 2021: 30th USENIX Security Symposium*, pages 591–608. USENIX Association, August 11–13, 2021.

[32] Aloni Cohen, Sarah Scheffler, and Mayank Varia. Can the government compel decryption? Don't trust – verify. In *2nd ACM Symposium on Computer Science and Law (to appear)*, 2022.

[33] U.S. Bureau of Economic Analysis. Advisory committee on data for evidence building. https://www.bea.gov/evidence.

[34] UN Committee of Experts on Big Data and Data Science for Official Statistics. Privacy preserving techniques task team. https://unstats.un.org/bigdata/task-teams/privacy/index.cshtml.

[35] 115th Congress. Foundations for evidence-based policymaking act of 2018. https://www.congress.gov/bill/115th-congress/house-bill/4174.

[36] Shoshana Zuboff. *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. Profile Books, 1st edition, 2018.

[37] Boston University. Open-source libraries, tools, and platforms that enable the implementation and deployment of applications that employ secure multi-party computation. https://github.com/multiparty.

[38] danah boyd and Jayshree Sarathy. Differential perspectives: Epistemic disconnects surrounding the us census bureau's use of differential privacy. In *Harvard Data Science Review (Forthcoming)*, 2022.

[39] U.K. Financial Conduct Authority. 2019 Global AML and Financial Crime TechSprint. https://www.fca.org.uk/events/techsprints/2019-global-aml-and-financial-crime-techsprint.

[40] Dorothy E. Denning and Dennis K. Branstad. A taxonomy for key escrow encryption systems. *Commun. ACM*, 39(3):34–40, 1996.

[41] David Chaum. Privategrity: online communication with strong privacy. In *Real World Cryptography*, 2016.

[42] Abhishek Bhowmick, Dan Boneh, Steve Myers, Kunal Talwar, and Karl Tarbe. The Apple PSI system. `https://www.apple.com/child-safety/pdf/Apple_PSI_System_Security_Protocol_and_Analysis.pdf`, 2021.

[43] Google. Federated learning of cohorts (floc). `https://github.com/WICG/floc`.

[44] Google. Turtledove. `https://github.com/WICG/turtledove`.

[45] Microsoft Corporation. Parakeet. `https://github.com/WICG/privacy-preserving-ads/blob/main/Parakeet.md`.

[46] Microsoft Corporation. Multi-party computation of ads on the web (MaCAW). `https://github.com/WICG/privacy-preserving-ads/blob/main/MACAW.md`.

[47] Harold Abelson, Ross J. Anderson, Steven M. Bellovin, Josh Benaloh, Matt Blaze, Whitfield Diffie, John Gilmore, Matthew Green, Susan Landau, Peter G. Neumann, Ronald L. Rivest, Jeffrey I. Schiller, Bruce Schneier, Michael A. Specter, and Daniel J. Weitzner. Keys under doormats. *Commun. ACM*, 58(10):24–26, 2015.

[48] Susan Landau. Punching the wrong bag: The deputy AG enters the crypto wars, October 2017. `https://www.lawfareblog.com/punching-wrong-bag-deputy-ag-enters-crypto-wars`.

[49] Hal Abelson, Ross J. Anderson, Steven M. Bellovin, Josh Benaloh, Matt Blaze, Jon Callas, Whitfield Diffie, Susan Landau, Peter G. Neumann, Ronald L. Rivest, Jeffrey I. Schiller, Bruce Schneier, Vanessa Teague, and Carmela Troncoso. Bugs in our pockets: The risks of client-side scanning. *arXiv*, abs/2110.07450, 2021. `https://arxiv.org/abs/2110.07450`.

[50] Ran Canetti and Gabriel Kaptchuk. The broken promise of Apple's announced forbidden-photo reporting system – and how to fix it. `https://www.bu.edu/riscs/2021/08/10/apple-csam`, 2021.

[51] Bennett Cyphers. Don't play in google's privacy sandbox. `https://www.eff.org/deeplinks/2019/08/dont-play-googles-privacy-sandbox-1`, 2019.

[52] Bennett Cyphers. Google's floc is a terrible idea. `https://www.eff.org/deeplinks/2021/03/googles-floc-terrible-idea`, 2021.

[53] Mayank Varia. A survey of MPC offerings. `https://drive.google.com/file/d/1NT_vdxRC8YEPlkQa2KHw22ai9IshyU73/view`, 2018.

[54] Catherine D'Ignazio and Lauren F. Klein. *Data Feminism.* MIT Press, 2020.

[55] Sasha Costanza-Chock. *Design Justice: Community-Led Practices to Build the Worlds We Need.* MIT Press, 2020.

[56] Paul Ohm and Jonathan Frankle. Desirable inefficiency. *Fla. L. Rev.*, 70:777, 2018.

[57] Paul Ohm. Statement of Paul Ohm, Professor, Georgetown University Law Center and Member, Commission on Evidence-Based Policymaking. In *Testimony for the Hearing on Protecting Privacy, Promoting Policy: Evidence-Based Policymaking and the Future of Education, Before the Committee on Education and the Workforce, U.S. House of Representatives*, 2018. Available at `https://bipartisanpolicy.org/download/?file=/wp-content/uploads/2019/03/CEP-Comissioner-Paul-Ohm-Testimony-from-1.30.18.pdf`.

[58] National Institute for Science and Technology. Cryptographic standards and guidelines. `https://csrc.nist.gov/projects/cryptographic-standards-and-guidelines/archived-crypto-projects/aes-development`, 2016.

[59] National Institute for Science and Technology. Hash functions. `https://csrc.nist.gov/projects/hash-functions/sha-3-project`, 2017.

[60] National Institute for Science and Technology. Post-quantum cryptography. `https://csrc.nist.gov/Projects/post-quantum-cryptography`, 2017.

[61] National Institute for Science and Technology. Lightweight cryptography. `https://csrc.nist.gov/Projects/lightweight-cryptography`, 2017.

[62] Prabhanjan Ananth, Abhishek Jain, Zhengzhong Jin, and Giulio Malavolta. Pre-constrained encryption. In *13th Innovations in Theoretical Computer Science Conference (ITCS 2022)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2022.

[63] Alexandra Wood, Micah Altman, Aaron Bembenek, Mark Bun, Marco Gaboardi, James Honaker, Kobbi Nissim, David R. O'Brien, Thomas Steinke, and Salil Vadhan. Differential privacy: A primer for a non-technical audience. *Vanderbilt Journal of Entertainment & Technology Law*, 21(1):209–275, 2018. `http://www.jetlaw.org/journal-archives/volume-21/volume-21-issue-1/differential-privacy-a-primer-for-a-non-technical-audience/`.

[64] Lucy Qin, Andrei Lapets, Frederick Jansen, Peter Flockhart, Kinan Dak Albab, Ira Globus-Harris, Shannon Roberts, and Mayank Varia. From usability to secure computing and back again. Cryptology ePrint Archive, Report 2019/734, 2019. `https://eprint.iacr.org/2019/734`.

[65] Ruba Abu-Salma, Elissa M. Redmiles, Blase Ur, and Miranda Wei. Exploring user mental models of End-to-End encrypted communication tools. In *8th USENIX Workshop on Free and Open Communications on the Internet (FOCI 18)*, Baltimore, MD, August 2018. USENIX Association.

[66] Omer Akgul, Ruba Abu-Salma, Wei Bai, Elissa M. Redmiles, Michelle L. Mazurek, and Blase Ur. *From Secure to Military-Grade: Exploring the Effect of App Descriptions on User Perceptions of Secure Messaging*, page 119–135. Association for Computing Machinery, New York, NY, USA, 2021.

[67] Christian Stransky, Dominik Wermke, Johanna Schrader, Nicolas Huaman, Yasemin Acar, Anna Lena Fehlhaber, Miranda Wei, Blase Ur, and Sascha Fahl. On the limited impact of visualizing encryption: Perceptions of E2E messaging security. In *Seventeenth Symposium on Usable Privacy and Security (SOUPS 2021)*, pages 437–454. USENIX Association, August 2021.

[68] Dominik Wermke, Nicolas Huaman, Christian Stransky, Niklas Busch, Yasemin Acar, and Sascha Fahl. Cloudy with a chance of misconceptions: Exploring users' perceptions and expectations of security and privacy in cloud office suites. In *Sixteenth Symposium on Usable Privacy and Security (SOUPS 2020)*, pages 359–377. USENIX Association, August 2020.