Gabriel Kaptchuk — Research Statement

My Motivations. My research is motivated by the desire to increase the number of privacy-preserving applications accessible to individuals and communities. To pursue this goal, I cross traditional discipline boundaries within the cryptography, computer security, and privacy research communities. I use my expertise and research skills to shepherd ideas across existing discipline boundaries, with the ultimate goal of promoting cryptographic deployments that center users' needs. Consequently, my published and ongoing work leverages a wide variety of techniques. For instance, I have conducted and analyzed surveys, detailed attacks on encrypted messaging systems, refined our understanding of cryptographic primitives (*i.e.* building blocks), and designed new formal models. My breadth of experience makes me a uniquely valuable researcher, capable of drawing from many methodologies and theoretical models to investigate new projects. Moreover, the ambitious nature of my primary research goal—applying cryptographic techniques to socially impactful, real-world problems—necessitates a wide variety of approaches.

Cryptography and Privacy. Cryptography is the study of algorithms and protocols that formally define who is able to access and control digital information. Although the academic study of cryptography grew from theoretical origins, it has taken on tremendous practical and social importance. Encryption now secures the vast majority of internet browsing traffic, and end-to-end encrypted messaging apps, like WhatsApp and Signal, are ubiquitous. Even advanced cryptographic primitives, like zero-knowledge proofs and secure multiparty computation (MPC), are used by industry and end-users. Cryptographic deployments are fundamental to securing business in the digital era and providing much-needed privacy to billions of people.

Despite the social importance of cryptography, ideas in modern cryptographic research usually trickle down an abstraction hierarchy. That is, the most abstract parts of the field—theoretical cryptography—drive the creation of new primitives, which are refined and optimized by applied cryptographers. Once a cryptographic primitive is concretely efficient, systems security researchers can incorporate it into actual systems proposals. Usable security researchers interview and survey users to evaluate the ways in which they interact with systems, and by extension the embedded cryptography. While there are exceptions to this pattern, the process of percolating ideas back up this hierarchy is rare, slow, and belabored.

While there is tremendous value in divided subdomains, this structure results in fewer (or mismatched) privacy-preserving tools for end users. Hierarchical research communities are hyper-focused, allowing the development of deep intuition and analysis methodologies. On the other hand, division of labor and the existing abstraction hierarchy leads to a gap between researchers developing new cryptography and end users; even the most important user ideas can be lost or distorted when traversing this gap.

My Research. I conduct research across all levels of this abstraction hierarchy. Rather than organize my research contributions and directions according to existing research communities or methodologies, I group them in a way that highlights how individual papers further my primary research goal of designing socially impactful cryptography. Concretely, I see my work as falling into three main categories:

(1) *Improving Foundational Cryptographic Primitives.* My first line of work aims to improve state-of-the-art foundational cryptographic primitives, including MPC and zero-knowledge proofs. My existing work on MPC has addressed issues of scalability [BGJK21] and fairness [CGJ+17]. My work on zero-knowledge proofs has focused on reducing the size of set membership proofs [GGHAK22a] and disjunctive proofs [GGHAK22b, GHAKS22]. These primitives can then be slotted into existing cryptographic use cases with improved efficiency and stronger security properties.

(2) *New Applications of Cryptography.* I'm passionate about designing cryptographic protocols that showcase novel use cases of cryptography for new applications. This process tests the limits of what is practically achievable with existing cryptographic techniques and explores intriguing ways to compose these building blocks. My research in this area includes using zero-knowledge to produce efficient proofs of software vulnerability [GHAH+23] and designing MPC protocols that could underpin dynamic MPC-as-a-service networks [CGG+21]. I have demonstrated ways to use cryptography to prevent rewind attacks on hardware-backed trusted execution environments like Intel SGX [KGM19] and use fully homomorphic encryption to power medical dataset analysis [KGR17]. Finally, my work has identified

vulnerabilities in deployed, end-to-end encrypted messaging systems and explored potential remedies [MKA+21, GGK+16].

(3) *Human-centered Cryptography.* I study cryptography through an explicitly socially aware lens, by designing and deploying cryptographic protocols that meet user's needs. For example, I am currently partnering with an activist organization to conduct a privacy-preserving survey on workplace culture problems in art museums, powered by MPC.[1] I have also studied censorship resistant communication to individuals communicating in closely monitored digital environments [KJGR21, Kap22], the technically feasibility of deploying law enforcement access systems that are resistant to abuse [GKL21], and technical tools that could provide increased security and privacy for sex workers [HKMR22]. Additionally, I endeavored to understand user perceptions of privacy-preserving systems, including deployments of differential privacy [CKR21] and COVID-19 contact tracing apps [KGH+20].

These three research directions are deeply linked and synergistic, and much of my work contains aspects from multiple of these directions. My research on foundational cryptographic primitives provides me with the technical knowledge required to build protocols for new applications. The process of designing new applications of cryptography gives me a unique perspective when it comes to building systems for concrete problems, and an understanding of the feasibility of using different cryptographic primitives. Moreover, understanding how cryptographic protocols are actually deployed and existing social needs provides clarity on the ways in which cryptographic primitives should be improved.

My research agenda is inherently interdisciplinary. While many of my techniques come from Computer Science, exploring new applications of cryptography and the social impact of cryptographic deployments requires building relationships with researchers and stakeholders outside of Computer Science. For instance, I was one of the first cryptographers to spend time embedded in the United States Senate, working closely with lawyers and policy makers on issues of privacy and technology [CK19]. The relationships and experiences from my time in the Senate help to shape my technical research, and I am currently organizing a series of workshops aimed at connecting interdisciplinary privacy researchers with policy experts. The process of building these interdisciplinary relationships can be slow, but it is crucial to doing impactful research.

Organization. In the sections below, I detail my research in each of these directions. I highlight some results in each direction to illustrate my work and briefly discuss ongoing and potential future work.

1 Improving Foundational Cryptographic Primitives

The first part of my work focuses on increasing the performance of MPC and zero-knowledge proofs.

Secure Multiparty Computation. Secure multiparty computation (MPC) [Yao86, GMW87, CCD88, BGW88] allows mutually distrusting parties to compute an arbitrary function of their secret inputs without disclosing anything about their inputs, beyond the output of the function. The security of the protocol relies on the assumption that a subset of the parties participating in the protocol do not collude. Since initial feasibility results for MPC 35 years ago, MPC has recently become practically efficient. Nonetheless, we have yet to see MPC reach its full potential. My research aims to improve state-of-the-art MPC protocols:

– *Scaling MPC To More Parties [BGJK21].* Having more parties participate in an MPC is highly desirable, as adding more parties to an MPC protocol dilutes the powers of each individual party. Unfortunately, the communication complexity and computation complexity of MPC protocols generally increases as the number of players increases. For large computational tasks, this will prevent many kinds of computational devices (*eg.* mobile devices, older computers) from participating in the protocol. In my work, presented at Eurocrypt 2021 [BGJK21], I demonstrated a new MPC protocol with per-party computation and communication costs that actually *decreased* as the number of parties grew. This work was the first time this property has been achieved for a non-trivial class of functionalities, including important applications like machine learning. My work also showed that the resulting protocol is concretely efficient by running it with hundreds of computational parties.

[1] For more information, see `https://museumsmovingforward.com/data-study`.

I recently followed up on this work by showing how to achieve the same asymptotic results in constant-round multiparty garbling setting, a family of protocols suitable for global scale deployments. I have also studied achieving fairness in MPC [CGJ+17] when the parties have access to some append-only ledger, like a blockchain. Without fairness, a protocol participant can learn the output of the computation and then cause the protocol to abort for the remaining parties—which can be a deal breaker in some scenarios.

Zero-Knowledge. Zero-knowledge proofs[2] [GMR85] allow a prover to convince a verifier that some instance x is in an NP language $\mathcal{L}$ without revealing any other information. Zero-knowledge proofs are an integral part of larger cryptographic protocols, providing a convenient way for protocol participants to prove that they are following the protocol honestly while maintaining privacy and anonymity. My recent work focuses on reducing the communication complexity of efficient zero-knowledge proofs:

– *Communication Efficient Disjunctive Composition of Σ-protocols [GGHAK22b].* Disjunctive statements compose a set of clauses with logical OR. Such proofs are critical for systems that require anonymity: simply being able to produce an accepting proof for a certain statement might allow a verifier to infer the prover's identity—even if the proof itself is zero-knowledge. Disjunctive composition allows a prover to "hide in a crowd," as the proof will not reveal which clause in the disjunction is associated with this particular prover. Existing generic disjunctive composition techniques for Σ-protocols, a practically important class of zero-knowledge proofs, produce proofs that grow linearly with the number of clauses. In my work, featured at EUROCRYPT 2022 [GGHAK22b], I demonstrated a new generic composition technique that produces proofs with communication complexity proportional only to the largest Σ-protocol. The technique relies on a new type of cryptographic commitment and is concretely efficient.

I recently built on this work by extending this composition technique to succinct proofs [GHAKS22]. In this case, the advantage of composing succinct proofs in this way is to significantly improve the prover's runtime, owing to a speed gap between the prover algorithm and the simulator. My work [GGHAK22a] also looks at reducing communication when proving set membership using the MPC-in-the-head paradigm [IKOS07].

Ongoing and Future Work. MPC protocols and zero-knowledge are both critical cryptographic building-blocks which are nearing the practicality required for widespread deployment. As such, I anticipate continuing my research in both areas. For example, I am working to improve the concrete efficiency of our multiparty garbling scheme and applying our disjunctive zero-knowledge techniques to produce more efficient zero-knowledge protocols for RAM programs.

As I continue to pursue my larger research goal, I anticipate my interests and expertise in other cryptographic primitives will grow. As mentioned above, I allow my work in this area to be guided by my findings in my other two directions. Thus, when I discover applications or social-implications of a different cryptographic primitive, I anticipate delving deeply into the technical research on that primitive.

2 New Applications of Cryptography

My second research direction is exploring new applications of cryptographic techniques. For instance, I studied deployed, end-to-end encrypted messaging systems iMessage [GGK+16] and Signal [MKA+21], which resulted in uncovering vulnerabilities and proposing mitigations. I also identified ways that fully homomorphic encryption could improve medical data analyst's existing workflows [KGR17] and spent a summer working inside Intel Research, after which I published a paper exploring how public ledgers (*eg.* blockchains) could be used to protect Intel Software Guard Extensions (SGX) from rewinding attacks [KGM19]. By finding new ways to compose cryptographic primitives into full systems, I uncover what is practically achievable with existing cryptographic techniques. Two of my more recent papers are also in this area:

– *MPC-as-a-service [CGG+21].* While cloud computing platforms have allowed the internet to flourish, this architecture requires trusting a few major companies with potentially sensitive data. MPC provides an opportunity to get the best of both worlds—outsourced computation with trust distributed across a large number of organizations. Ideally, existing computational resources could contribute to this privacy-preserving computational service, reducing waste and further distributing trust. Unfortunately, modern

[2] I refer to zero-knowledge proofs and zero-knowledge arguments interchangeably, as is common in practical work.

MPC protocols are inflexible when it comes to their participation models. In general, the computational parties participating in the protocol are required to remain online throughout the execution, and if parties need to drop out of the protocol before it completes—possibly because of an error or other computational obligations—the protocol will fail. In my work, presented at CRYPTO 2021, I explored a new participation model called Fluid MPC, in which one can design MPC protocols that allow parties to leave and join the protocol execution. This provides a formal underpinning that would power an MPC-as-a-service platform.

– *Zero-knowledge Proofs of Exploit [GHAH+23].* The advent of concretely efficient and scalable zero-knowledge proof systems enable new, complex use cases for zero-knowledge. In a work that will appear in the proceedings of the Privacy Enhancing Technologies Symposium 2023, I explore constructing zero-knowledge proofs of exploit, a demonstration that a prover possesses some input to a program that puts the program in an invalid state. These proofs can operate over unmodified binaries, meaning that the prover does not need access to source code—a common strategy in prior work. We implemented such proofs for unmodified MSP430 (a common micro-controller) binaries and demonstrated that existing zero-knowledge proof protocols are efficient enough to make such proofs of exploit usable today.

Ongoing and Future Work. I continue to find new applications of cryptographic primates that push the boundaries of what is practically achievable. I'm currently working with a group of IoT researchers to investigate different ways MPC-as-a-service could be run on lightweight IoT devices and studying the applications that this could power. Additionally, I am studying the concrete costs associated with using zero-knowledge, in a low-effort, black box manner, to improve the accountability and transparency of existing cryptographic protocols, like Apple's recent CSAM scanning proposal [BBM+21].

3 Human-centered Cryptography

My final research direction is studying cryptography through an explicitly socially aware lens, a processing which requires building cross-disciplinary relationships. My work in this area falls into two categories: (1) building cryptographic systems for social good, and (2) studying how end-users understand cryptographic deployments, with the goal of improving the utility of future deployments.

Cryptography for Social Good. My work has investigated cryptography for social good on many fronts. I have studied the feasibility of designing law enforcement access mechanisms with strong, formal notions of abuse resistance into encrypted communication platforms. In my recent work [GKL21], presented at Eurocrypt 2021, I found that it is technically infeasible to meet the demands of law enforcement while also protecting against insider threats and foreign, state-sponsored, attacks. Additionally, I am working closely with experts in sex work to design tools that would improve the digital privacy of sex workers. We recently produced a paper containing 11 concrete security project ideas that we believe would have substantial real-world impact [HKMR22]; our hope is that by releasing these ideas, we can garner wider research effort in this area. Two of my other projects in this area are:

– *Deploying MPC for Social Good.* I am currently partnering with an organization called Museums Moving Forward (MMF), an initiative to support greater equity and accountability in art museum workplaces, to conduct large-scale surveys of workplace culture problems in art museums. Because of widespread distrust between employees, directors, and funders within the art museum work, MMF wanted to run the survey in a privacy-preserving way which minimized the risk that individual's responses might be disclosed while still allowing for high-quality, aggregate results; we have designed an MPC-based system to fit their needs. The use of MPC is allowing MMF to gather data about highly sensitive workplace problems, including pay, harassment, and promotion rates—issues that prior survey efforts have been unable to adequately study. Data collection began in November 2022 and nearly 10,000 individuals around the country have been encouraged to participate. We expect to release a public report of the findings in May 2023.

– *Practical Steganography For Censorship Resistance [KJGR21].* The ubiquitous adoption of digital communication platforms has amplified the surveillance capabilities of nation states. Moreover, countries with repressive regimes regularly practice network censorship, blocking user access to content (*eg.* Wikipedia),

platforms (*eg.* Facebook, WhatsApp), and protocols (*eg.* Tor) that the state deems subversive. This situation is extraordinarily dangerous for activist and community organizers whose communications are always the subject of close scrutiny. As government antipathy towards encryption continues to grow globally, there will be a desperate need to develop tools that can resist this strong form of censorship. In my recent work [KJGR21], presented at ACM CCS 2021 and featured at Real World Cryptography 2022 [Kap22], I explored creating strong censorship resistance tools, focusing specifically on steganography. Steganography is a technique that facilitates hiding a sensitive message within a mundane message, such that a censor cannot distinguish between the communication that it deems subversive and permissible. While steganography has been well studied by the theory community for over 30 years, practical steganography capable of hiding messages within normal looking human communications has never been deployed. In our work, we identified why prior attempts have failed, demonstrated how to overcome these difficulties, and concretely instantiated practical steganography. We implemented our protocol, facilitating steganographic encoding into English text, and evaluated our implementation across multiple computational environments, including mobile.

End-user's Understanding of Cryptography. To maximize the chances that a cryptographic deployment realizes its potential, it is critical that the protocol designers meaningfully understand how end-users will view the system. To that end, I study user perceptions of privacy-preserving systems. For example, I studied how users understand privacy-accuracy tradeoffs in COVID-19 contact tracing apps [KGH+20]. My more recent work in this area focuses on differential privacy:

– *End-users Understanding of Differential Privacy [CKR21].* Differential privacy [DMNS06] is a mathematically rigorous definition of privacy that limits the amount of information about individuals that can be disclosed by aggregate statistics computed over a dataset. Since its introduction, differentially privacy has been swiftly embraced and has recently been deployed by companies and governments. However, understanding the privacy guarantees of a differentially private system can be challenging. Spurred by a concern that users might be misled by descriptions of differential privacy, I designed, conducted, and analyzed two large surveys that shed light on how the descriptions of differential privacy users might encounter in the wild affect user's data sharing habits and privacy expectations. The results, which received a Best Paper Runner-up Award at ACM CCS 2021 [CKR21], show that users are more likely to share their data when it will be protected with differential privacy, but common descriptions of differential privacy do a poor job of setting user's privacy expectations. Additionally, the results suggest a critical factor in determining a user's willingness to share is the alignment between that user's privacy concerns and the way the description of differential privacy raises their expectations.

Ongoing and Future Work. My work on all the highlighted projects is ongoing. I am a PI on a NSF funded program in which we are building an end-to-end tool for censorship resistant communication. My experiences deploying MPC have highlighted audit-ability and usability needs for MPC-based systems that I plan to study in the coming years. Finally, one of my co-authors on [CKR21] was awarded a multi-year grant to extend our work (part of which will fund my effort). We are studying better ways to describe both differential privacy and its security parameters to users. Finally, I was recently awarded a seed-grant from the BU Center for Anti-racist Research to study impact of differential privacy in the census, and am developing a framework to more rigorously reason about the risks associated with forgoing differential privacy.

4 Conclusion

My research is driven by my desire to create more privacy-preserving cryptographic applications that promote the social good of individuals and communities. The need for privacy has never been more significant, making this agenda particularly timely. My approach spans multiple subdisciplines and methodologies, including cryptography, systems security, and usable security. I have demonstrated my ability to pursue this ambitious research agenda by producing top-tier publications that engage with my core research goals from several angles and successfully applying for funding to continue my work. As a tenure-track faculty member, I will continue my work and build out a research team capable of increasing my impact.

References

[BBM+21] 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`, July 2021. (Accessed on Nov 11, 2021). Apple has since removed this document; it is cached at `https://cs-people.bu.edu/kaptchuk/filedump/Apple_PSI_System_Security_Protocol_and_Analysis.pdf`.

[BGJK21] Gabrielle Beck, Aarushi Goel, Abhishek Jain, and Gabriel Kaptchuk. Order-C secure multiparty computation for highly repetitive circuits. In Anne Canteaut and François-Xavier Standaert, editors, *EUROCRYPT 2021, Part II*, volume 12697 of *LNCS*, pages 663–693. Springer, Heidelberg, October 2021.

[BGW88] Michael Ben-Or, Shafi Goldwasser, and Avi Wigderson. Completeness theorems for non-cryptographic fault-tolerant distributed computation (extended abstract). In *20th ACM STOC*, pages 1–10. ACM Press, May 1988.

[CCD88] David Chaum, Claude Crépeau, and Ivan Damgård. Multiparty unconditionally secure protocols (abstract) (informal contribution). In Carl Pomerance, editor, *CRYPTO'87*, volume 293 of *LNCS*, page 462. Springer, Heidelberg, August 1988.

[CGG+21] Arka Rai Choudhuri, Aarushi Goel, Matthew Green, Abhishek Jain, and Gabriel Kaptchuk. Fluid MPC: Secure multiparty computation with dynamic participants. In Tal Malkin and Chris Peikert, editors, *CRYPTO 2021, Part II*, volume 12826 of *LNCS*, pages 94–123, Virtual Event, August 2021. Springer, Heidelberg.

[CGJ+17] 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*, pages 719–728. ACM Press, October / November 2017.

[CK19] Shaanan Cohney and Gabriel Kaptchuk. The hill we must die on: Cryptographers and congress. Real World Cryptography 2019. Available on YouTube, 2019. `https://youtu.be/AyeuIfGjBg4?t=1180`.

[CKR21] 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*, pages 3037–3052. ACM Press, November 2021.

[DMNS06] 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*, volume 3876 of *LNCS*, pages 265–284. Springer, Heidelberg, March 2006.

[GGHAK22a] Aarushi Goel, Matthew Green, Mathias Hall-Andersen, and Gabriel Kaptchuk. Efficient set membership proofs using MPC-in-the-head. *PoPETs*, 2022(2):304–324, April 2022.

[GGHAK22b] Aarushi Goel, Matthew Green, Mathias Hall-Andersen, and Gabriel Kaptchuk. Stacking sigmas: A framework to compose Σ-protocols for disjunctions. In Orr Dunkelman and Stefan Dziembowski, editors, *EUROCRYPT 2022, Part II*, volume 13276 of *LNCS*, pages 458–487. Springer, Heidelberg, May / June 2022.

[GGK+16] 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*, pages 655–672. USENIX Association, August 2016.

[GHAH+23] Matthew Green, Mathias Hall-Andersen, Eric Hennenfent, Gabriel Kaptchuk, Benjamin Perez, and Gijs Van Laer. Efficient proofs of software exploitability for real-world processors. *PoPETs*, 2023(1), April 2023.

[GHAKS22] Aarushi Goel, Mathias Hall-Andersen, Gabriel Kaptchuk, and Nicholas Spooner. Speed-stacking: Fast sublinear zero-knowledge proofs for disjunctions. Cryptology ePrint Archive, Paper 2022/1419, 2022. `https://eprint.iacr.org/2022/1419`.

[GKL21] Matthew Green, Gabriel Kaptchuk, and Gijs Van Laer. Abuse resistant law enforcement access systems. In Anne Canteaut and François-Xavier Standaert, editors, *EUROCRYPT 2021, Part III*, volume 12698 of *LNCS*, pages 553–583. Springer, Heidelberg, October 2021.

[GMR85] Shafi Goldwasser, Silvio Micali, and Charles Rackoff. The knowledge complexity of interactive proof-systems (extended abstract). In *17th ACM STOC*, pages 291–304. ACM Press, May 1985.

[GMW87] Oded Goldreich, Silvio Micali, and Avi Wigderson. How to play any mental game or A completeness theorem for protocols with honest majority. In Alfred Aho, editor, *19th ACM STOC*, pages 218–229. ACM Press, May 1987.

[HKMR22] Vaughn Hamilton, Gabriel Kaptchuk, Allison McDonald, and Elissa M. Redmiles. Safe digital intimacy: A research agenda, 2022. `https://cs-people.bu.edu/kaptchuk/publications/sw-threat-modeling.pdf`.

[IKOS07] Yuval Ishai, Eyal Kushilevitz, Rafail Ostrovsky, and Amit Sahai. Zero-knowledge from secure multiparty computation. In David S. Johnson and Uriel Feige, editors, *39th ACM STOC*, pages 21–30. ACM Press, June 2007.

[Kap22] Gabriel Kaptchuk. Commit acts of steganography——before it's too late. Real World Cryptography 2022. Available on YouTube, 2022. `https://youtu.be/6Gb0xO8csVU?t=54`.

[KGH+20] Gabriel Kaptchuk, Daniel G. Goldstein, Eszter Hargittai, Jake Hofman, and Elissa M. Redmiles. How good is good enough for covid19 apps? the influence of benefits, accuracy, and privacy on willingness to adopt, 2020.

[KGM19] Gabriel Kaptchuk, Matthew Green, and Ian Miers. Giving state to the stateless: Augmenting trustworthy computation with ledgers. In *NDSS 2019*. The Internet Society, February 2019.

[KGR17] Gabriel Kaptchuk, Matthew Green, and Aviel D. Rubin. Outsourcing medical dataset analysis: A possible solution. In Aggelos Kiayias, editor, *FC 2017*, volume 10322 of *LNCS*, pages 98–123. Springer, Heidelberg, April 2017.

[KJGR21] Gabriel Kaptchuk, Tushar M. Jois, Matthew Green, and Aviel D. Rubin. Meteor: Cryptographically secure steganography for realistic distributions. In Giovanni Vigna and Elaine Shi, editors, *ACM CCS 2021*, pages 1529–1548. ACM Press, November 2021.

[MKA+21] Ian Martiny, Gabriel Kaptchuk, Adam J. Aviv, Daniel S. Roche, and Eric Wustrow. Improving signal's sealed sender. In *NDSS 2021*. The Internet Society, February 2021.

[Yao86] Andrew Chi-Chih Yao. How to generate and exchange secrets (extended abstract). In *27th FOCS*, pages 162–167. IEEE Computer Society Press, October 1986.