Research Statement

Mohammad Mahmoody (August 2018)

My research is focused on **foundations of cryptography**, which is the the science of designing *provably* secure protocols based computationally intractable problems. My specialization is in studying the power and limitations of **computational assumptions** in cryptography. A main theme in my research is to identify barriers against basing cryptographic protocols on well-studied computational assumptions. I have also studied *trade-offs* that emerge between efficiency of cryptographic protocols and assumptions behind their designs. My research has helped develop the field of *separations* (i.e., barriers) in cryptography with a focus on more modern primitives such as program obfuscation and functional encryption. The second theme in my research investigates the power and limitations of **physical assumptions** in cryptography that could potentially help us avoid computational assumptions. In this line of work, I have studied the possibility of basing cryptographic tasks on stateless tamper-proof hardware and related mathematical tools that arise in this area such as zero-knowledge (interactive) probabilistically checkable proofs (PCPs). Finally, a more recent direction in my work investigates the power and limitations of **tampering attacks in cryptography** and **learning theory**. This line of work studies vulnerabilities that could be exploited by adversaries who tamper with the randomness or the inputs of the algorithms involved. In the field of adversarial machine learning, this tampering could happen over the the training data of the learner (a poisoning attack) or the test example (an evasion attack). My research investigates **provable bounds** in both of these settings.

Research group and support. My research group currently consists of three PhD students and two post-doctoral researchers. One PhD student and two Masters students have previously graduated from my group. I support my group's research through NSF CAREER award (on cryptographic separations), a subcontract from UC Berkeley (on extending non-black-box separations), two University of Virginia's SEAS Innovation Award (on adversarial learning and adversarial algorithmic fairness) and external fellowships.

Publishing venues and alphabetical order of authors. My works in cryptography appear in top conferences dedicated for foundations and theory of cryptography such as Crypto and TCC. As it is the tradition in Theoretical Computer Science, in almost all of my papers authors are listed in the alphabetical order.

Computational Assumptions in Cryptography

Modern cryptography has given rise to protocols whose security is based on well-defined and well-studied mathematical puzzles (e.g., factoring large integers) that are believed to be computationally hard to solve. Unfortunately, however, we are yet far from *proving* such hardness results in mathematics. Yet, proving the security of almost all cryptographic protocols requires resolving the notorious $\mathbf{P} \neq \mathbf{NP}$ question. As a result, many current cryptographic security proofs heavily rely on computational hardness **assumptions**. A core theme in my research aims at identifying the power of computational intractability assumptions in cryptography. Over the years, my research has taken major steps toward understanding the *limitations* of computational assumptions in cryptography through a theory of *cryptographic separations*. At a more technical level, many of my works help us getting insight into relations between Impagliazzo's worlds [45].

Private-key encryption vs. $\mathbf{NP}$-hardness [44, 65]. Since $\mathbf{P} \neq \mathbf{NP}$ is a *necessary* requirement for the security of almost all cryptographic tasks, a holy grail in cryptography, partly motivated by worst-case to average-case reductions [1, 68, 69] for lattice problems, is to base security of protocols *solely* on the assumption that $\mathbf{P} \neq \mathbf{NP}$ [25]. With Haitner and Xiao [44, 65], I came up with mathematical explanations as to why so far, despite tremendous effort by researchers, designing encrypting schemes (or strong hash functions) with $\mathbf{NP}$-hard security has remained elusive. In particular, we show that in order to achieve $\mathbf{NP}$-hard one-way functions, we need to first resolve long-standing open questions in computational complexity about the existence of program-checkers for Boolean satisfiability, or alternatively, we have to construct interactive proofs with "low complexity" provers in forms that are not known to exist yet.

Public-key vs. private-key cryptography. Private-key encryption can be based on secure hash functions, while building public-key encryption schemes are designed based on structured hard problems from mathematics [5, 10]. Such structured problems, however, also have the potential for enabling *attacks* that exploit this very structure. Therefore, one of the most fundamental questions in cryptography is whether we can base public-key encryption on private-key primitives such as secure hash functions. In 1989, Impagliazzo and Rudich [46] showed that a large class of techniques, called *black-box*, are incapable of achieving this goal. However, black-box techniques are only part of cryptographic toolkit, so extending [46] to non-black-box settings remained a major problem in foundations of cryptography.

- **Power of non-black-box garbling techniques for public-key encryption [36].** With Garg, our joint postdoc Hajiabadi and my graduate student Mohammed [36], I showed that even a popular and powerful *non-black-box* technique in cryptography based on garbled circuits [81] is not capable of basing public-key encryption on private key encryption. We proved our result in a model that was developed in [2, 16] by allowing oracle calls inside given garbled circuits. This model includes a large class of natural non-black-box tricks built into it beyond the fully-black-box framework of Reingold et al. [76]. In another work [38] (explained below), we further expand this non-black-box model.
- **Tight security reductions between public-key and private-key cryptography [7, 8].** Together with Barak [7, 8], I studied whether relaxed forms of public-key encryption (with only a weak *polynomial* security) or fully secure public-key authentication schemes could be based solely on secret-key encryption. We proved *optimal* bounds for the exact achievable security of those fundamental tasks based on ideal hash functions. Our results showed that seminal works of Merkle [67] and Lamport [53] were indeed optimal. Our work [8] resolved a long standing open question of [46].

Complexity of recently-developed powerful encryption primitives [39, 43, 61]. During the twenty first century, cryptography has gone through a revolution of exploring the *feasibility* of highly structured tasks at the cost of relying on newly introduced and less studied assumptions. One such success story led to development of strong encryption primitives such predicate encryption [51] and functional encryption [15, 72]. Understanding the computational assumptions necessary for achieving these powerful primitives is of great importance. With Garg and my PhD student Mohammed [39] we identified which forms of functional encryption could be obtained from code obfuscation (further discussed below). Together with Goyal, Kumar and Lokam [43] and my PhD student Mohammed [61], we studied the power of *identity-based encryption* (IBE) [13, 79]. In an IBE scheme, knowing a single master public key is sufficient to encrypt to each identity, and it is known that IBE is more complex than basic semantically secure public-key encryption [12, 14]. In [61], we proved *limitations* on the power of IBE by showing that hash functions or homomorphic encryption [17, 40, 80] cannot be based on IBE in a black-box way.

Complexity of assumptions for program obfuscation [38, 59, 60]. Program obfuscation was first formally studied in the theory community by Barak et.al [6] where they showed that very strong forms of obfuscation are impossible, but it took till the work of Garg et al. [35] where a powerful form of obfuscation, called indistinguishability obfuscation (IO), was proposed based on the existence of multi-linear maps [34]. In a series of works [38, 59, 60], all co-authored with my graduate students Ameer Mohammed (whose thesis was on this very subject) and student Nematihaji as well as Garg, Pass, and shelat, I proved strong lower-bounds on standard assumptions that can be used to construct IO. Roughly speaking, building upon lower bounds for VBB obfuscating [19] and developing ideas for IO itself, we showed hat IO is too complex to be built solely from any encryption primitive that is of "all-or-nothing" access structure.

Modeling *non-black-box* constructions and reductions [21, 38, 64]. With Garg and my PhD student Mohammed [38], we proved strong *non-black-box* impossibility results for IO from powerful assumptions such as predicate encryption. A major contribution of this work was to introduce a *new model* for framing a broad class of non-black-box techniques as "monolithic" subroutine calls. This model is non-black-box under the original framework developed by Reingold et al [76] and the subsequent extensions of [2, 16], yet it includes natural techniques widely used in cryptography. This new model enables a more meaningful study of separations with respect to known non-black-box techniques. With Pass [64] we had previously shown that, although non-black-box constructions can usually be made black-box [20, 74], there are cases where *only* a non-black-box construction could base on on the other one. When it comes to *security reductions*, with Chung, Lin and Pass [21], I initiated a formal study of *mildly* non-black-box *proofs* of security in which non-uniform advice about the adversary is the source of non-black-box nature of the security reduction.

Complexity of time-lock puzzles and proofs of work [62, 63]. Time-lock puzzles, first constructed by [77], allow encrypting messages that are only decryptable after a specified time has passed. The tightly related notion of *proof of work* [18, 31, 32] and its variations have recently got more attention [11, 75] after finding applications in crypto-currencies [71]. In two works with Moran and Vadhan [62, 63], I study whether time-lock puzzles and (sequential) proofs of work could be based on the mere assumption that one-way functions exist, or more strongly using any random oracle. In [62], we showed that time-lock puzzles cannot be constructed in the random oracle model, and in [63], we showed that proofs of work with *multiple* correct solutions can indeed be constructed using random oracles and depth robust graphs [33].

Separations and assumption trade-offs for secure computation. In a body of works discussed below, I studied the complexity of the assumptions that are necessary for secure computation protocols. In such a protocol, mutually distrustful entities engage to compute a joint function on their local private inputs.

- **Complexity of two-party computation [57, 58].** In two works with Maji and Prabhakaran [57, 58] we characterized the black-box power of private-key as well as public-key cryptography in secure computation. We showed that random oracles (or even strong forms of public-key encryption) cannot help weaken assumptions behind secure computation systems while using a black-box construction.

- **Complexity of fair coin tossing [23, 24].** Secure coin tossing is a basic task in secure computation. It was shown by Cleve [22] that in any two party coin-tossing protocol one party can bias the output by $\Omega(1/r)$ where r is the round complexity of the protocol. Despite constructions of [4, 70] it remained open whether optimally fair coin tossing (with bias $O(1/r)$) could be based on the minimal assumption of one-way functions. In two works with Dachman-Soled, Lindell and Malkin [23, 24], we identified a key barrier against this goal by proving lower bounds on the round complexity.

- **Complexity of round-preserving OT extension [37].** Oblivious transfer (OT) is the building block of secure computation [49, 52], but since it is a costly operation, researchers have suggested efficient ways to *extend* a few base OT operations into many constructed OT operations [47] while only using *cheap symmetric-key* cryptographic operations that can be obtained from a random oracle. The solution of [47], however, adds one more round of communication between the parties. With Garg, Masny, and Meckler [37], I studied the possibility of *very efficient oblivious transfer (OT) extension* protocols and proved in that a cost in round complexity of OT extension is indeed inherent so long as do not want to pay the inefficiency costs of *non-black-box* constructions [9].

Tampering Attacks in Cryptography and Learning Theory

In the fields of algorithm design and cryptography, we typically assume that honest parties have access to uniform and independent randomness, and indeed many tasks (e.g., secure multi-party encryption) are otherwise impossible [26]. In particular, "standard" security proofs no longer hold if adversaries can tamper with the randomness of honest parties. Such attacks also emerge in the area of adversarial machine learning where we also deal with tampering attacks of various forms to the training process or the final classification.

Tampering with randomness in Cryptography [3, 56]. It is known that imperfect sources of randomness, in general, cannot be used for cryptography [27, 28, 29]. With Austrin, Chung, Pass and Karn [3], we studied the possibility of achieving security in cryptography if the randomness of the parties might be under attack by *efficient* viruses who can read everything but can only change the randomness of the system in a limited way. In [3], we demonstrated some basic cryptographic tasks that are impossible to achieve in this setting. Motivated by the fact that randomness is usually generated in blocks rather than bits, with my PhD student Mahloujifar [56], I extended our previous results of [3] by providing attacks even if the randomness is generated in multiple chunks, while each of the blocks is independently tamperable. Our works [3, 56] also gave *algorithmic* proofs for impossibility of extracting randomness from (blockwise) SV sources [78].

Power of tamper-proof hardware in cryptography [41, 48, 66]. Assuming *strong* forms of tamper proof assumptions about hardware lead to positive results [42, 50] that are otherwise impossible without computational assumptions. A candidate approach for achieving *unconditional* security without relying on computational assumptions is to use alternative *physical* models of interaction [26]. In a sequence of works with Goyal, Ishai, Sahai and Xiao [41, 48, 66], I showed how to build cryptography, and in particular secure computation, on *physical* assumptions through *resettable* tamper-proof hardware (that are even allowed to be reset by the adversary) rather than using unproven computational hardness assumptions.

Provable bounds against tampering adversaries in machine learning [55, 56]. In two works with my PhD student Saeed Mahloujifar [55, 56], I studied tools and techniques for tampering attacks that are powerful enough to be applied to domains outside cryptography such as *adversarial machine learning*. In particular, we showed [56] how to increase the error in a learning algorithm *in polynomial time* by tampering only with p fraction of the training data. As opposed to many heuristic attacks in this area, our work led to *provable* bounds of success by efficient attackers. In a follow-up work [54], we show how similar ideas can also be applied to prove inherent bounds for *evasion* attacks where the goal of the adversary is to find adversarial examples that are close to honestly generated ones, but are misclassified by the trained model.

Future Plans

At a high level, my plans for future research is to further develop the theory of assumptions in cryptography and build connections between this field and the younger field of adversarial machine learning. Below, I briefly explain some of my goals in these two research directions.

Exploring the power of non-black-box techniques. My future goal is to continue studying the limitations of computational assumptions and techniques in foundations of cryptography. Towards this goal, I plan to expand my research group and collaborate with new colleagues in other schools. On the other hand, when it comes to *positive* results, I also plan to study the *power* of non-black-box techniques such a garbling schemes in cryptography. Pursuing this goal is ever more motivated by recent results in cryptography [30] showing the inherent power of non-black-box constructions for natural cryptographic tasks [73].

Understanding the power of adversarial agents, beyond attacking learners. My second plan for the future is to expand my research in provable bounds in adversarial machine learning. Leveraging on our initial results [54, 55, 56] in adversarial learning, I plan to study, from a *provable* perspective, how strategic agents (i.e., adversaries, as we call them in cryptography) can affect decision making processes in which utility function models measures *other* than the optimality of the choices. In particular, I plan to study the power and limitation of attackers who target the *fairness* of sequential decision making algorithms. Due to the ever increasing role of automatic decision making systems, understanding the answer to this question is more important than ever.

References

[1] Miklós Ajtai. Generating hard instances of lattice problems. In *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*, pages 99–108. ACM, 1996.

[2] Gilad Asharov and Gil Segev. Limits on the power of indistinguishability obfuscation and functional encryption. *SIAM Journal on Computing*, 45(6):2117–2176, 2016.

[3] Per Austrin, Kai-Min Chung, Mohammad Mahmoody, Rafael Pass, and Karn Seth. On the impossibility of cryptography with tamperable randomness. *Advances in Cryptology - CRYPTO*, 2014.

[4] B. Awerbuch, M. Blum, B. Chor, S. Goldwasser, and S. Micali. Limits on the security of coin flips when half the processors are faulty. In *Unpublished manuscript*, 1985.

[5] Boaz Barak. The complexity of public-key cryptography. In *Tutorials on the Foundations of Cryptography*, pages 45–77. Springer, 2017.

[6] Boaz Barak, Oded Goldreich, Russell Impagliazzo, Steven Rudich, Amit Sahai, Salil Vadhan, and Ke Yang. On the (im) possibility of obfuscating programs. *Journal of the ACM (JACM)*, 59(2):6, 2012.

[7] Boaz Barak and Mohammad Mahmoody. Lower bounds on signatures from symmetric primitives. In *IEEE Symposium on Foundations of Computer Science (FOCS)*, 2007.

[8] Boaz Barak and Mohammad Mahmoody. Merkle puzzles are optimal - an $O(n^2)$-query attack on any key exchange from a random oracle. In *Advances in Cryptology - CRYPTO*, volume 5677 of *Lecture Notes in Computer Science*, pages 374–390. Springer, 2009.

[9] Donald Beaver. Correlated pseudorandomness and the complexity of private computations. In *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*, pages 479–488. ACM, 1996.

[10] Nir Bitansky, Akshay Degwekar, and Vinod Vaikuntanathan. Structure vs. hardness through the obfuscation lens. In *Annual International Cryptology Conference*, pages 696–723. Springer, 2017.

[11] Dan Boneh, Joseph Bonneau, Benedikt Bünz, and Ben Fisch. Verifiable delay functions. In *Annual International Cryptology Conference*, pages 757–788. Springer, 2018.

[12] Dan Boneh, Ran Canetti, Shai Halevi, and Jonathan Katz. Chosen-ciphertext security from identity-based encryption. *SIAM Journal on Computing*, 36(5):1301–1328, 2006.

[13] Dan Boneh and Matthew Franklin. Identity-based encryption from the Weil pairing. *SIAM Journal on Computing*, 32(3):586–615, June 2003. Preliminary version in CRYPTO '01.

[14] Dan Boneh, Periklis Papakonstantinou, Charles Rackoff, Yevgeniy Vahlis, and Brent Waters. On the impossibility of basing identity based encryption on trapdoor permutations. In *IEEE 49th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 283–292. IEEE, 2008.

[15] Dan Boneh, Amit Sahai, and Brent Waters. Functional encryption: Definitions and challenges. In *Theory of Cryptography Conference*, pages 253–273. Springer, 2011.

[16] Zvika Brakerski, Jonathan Katz, Gil Segev, and Arkady Yerukhimovich. Limits on the power of zero-knowledge proofs in cryptographic constructions. In *Theory of Cryptography Conference*, pages 559–578. Springer, 2011.

[17] Zvika Brakerski and Vinod Vaikuntanathan. Efficient fully homomorphic encryption from (standard) LWE. *SIAM Journal on Computing*, 43(2):831–871, 2014.

[18] J-Y Cai, Richard J Lipton, Robert Sedgewick, and AC-C Yao. Towards uncheatable benchmarks. In *Structure in Complexity Theory Conference, Proceedings of the Eighth Annual*, pages 2–11. IEEE, 1993.

[19] Ran Canetti, Yael Tauman Kalai, and Omer Paneth. On obfuscation with random oracles. In *Theory of Cryptography Conference*, pages 456–467. Springer, 2015.

[20] Seung Geol Choi, Dana Dachman-Soled, Tal Malkin, and Hoeteck Wee. Black-box construction of a non-malleable encryption scheme from any semantically secure one. In *Theory of Cryptography Conference*, pages 427–444. Springer, 2008.

[21] Kai-Min Chung, Huijia Lin, Mohammad Mahmoody, and Rafael Pass. On the power of nonuniformity in proofs of security. In *Innovations in Theoretical Computer Science - ITCS*, pages 389–400. ACM, 2013.

[22] Richard Cleve. Limits on the security of coin flips when half the processors are faulty. In *Proceedings of the eighteenth annual ACM symposium on Theory of computing*, pages 364–369. ACM, 1986.

[23] Dana Dachman-Soled, Yehuda Lindell, Mohammad Mahmoody, and Tal Malkin. On the black-box complexity of optimally-fair coin tossing. In *Theory of Cryptography Conference*, pages 450–467. Springer, 2011.

[24] Dana Dachman-Soled, Mohammad Mahmoody, and Tal Malkin. Can optimally-fair coin tossing be based on one-way functions? In *Theory of Cryptography Conference*, pages 217–239. Springer, 2014.

[25] Whitfield Diffie and Martin Hellman. New directions in cryptography. *IEEE transactions on Information Theory*, 22(6):644–654, 1976.

[26] Dodis, Ong, Prabhakaran, and Sahai. On the (im)possibility of cryptography with imperfect randomness. In *FOCS: IEEE Symposium on Foundations of Computer Science (FOCS)*, 2004.

[27] Yevgeniy Dodis, Shien Jin Ong, Manoj Prabhakaran, and Amit Sahai. On the (im) possibility of cryptography with imperfect randomness. In *45th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 196–205. IEEE, 2004.

[28] Yevgeniy Dodis, David Pointcheval, Sylvain Ruhault, Damien Vergniaud, and Daniel Wichs. Security analysis of pseudo-random number generators with input:/dev/random is not robust. In *Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security*, pages 647–658. ACM, 2013.

[29] Yevgeniy Dodis and Yanqing Yao. Privacy with imperfect randomness. In *Annual Cryptology Conference*, pages 463–482. Springer, 2015.

[30] Nico Döttling and Sanjam Garg. Identity-based encryption from the Diffie-Hellman assumption. In *Annual International Cryptology Conference*, pages 537–569. Springer, 2017.

[31] Cynthia Dwork and Moni Naor. Pricing via processing or combatting junk mail. In *Annual International Cryptology Conference*, pages 139–147. Springer, 1992.

[32] Cynthia Dwork, Moni Naor, and Hoeteck Wee. Pebbling and proofs of work. In *Annual International Cryptology Conference*, pages 37–54. Springer, 2005.

[33] Paul Erdoes, Ronald L Graham, and Endre Szemeredi. On sparse graphs with dense long paths. 1975.

[34] Sanjam Garg, Craig Gentry, and Shai Halevi. Candidate multilinear maps from ideal lattices. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*. Springer, 2013.

[35] Sanjam Garg, Craig Gentry, Shai Halevi, Mariana Raykova, Anant Sahai, and Brent Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. In *Foundations of Computer Science (FOCS), 2013 IEEE 54th Annual Symposium on*, pages 40–49. IEEE, 2013.

[36] Sanjam Garg, Mohammad Hajiabadi, Mohammad Mahmoody, and Ameer Mohammed. Limits on the power of garbling techniques for public-key encryption. In *Annual International Cryptology Conference (CRYPTO)*, 2018.

[37] Sanjam Garg, Mohammad Mahmoody, Daniel Masny, and Izaak Meckler. On the round complexity of OT extension. In *Annual International Cryptology Conference (CRYPTO)*, 2018.

[38] Sanjam Garg, Mohammad Mahmoody, and Ameer Mohammad. Lower bounds on indistinguishability obfuscation from all-or-nothing encryption primitives. In *Annual International Cryptology Conference (CRYPTO), Springer, Cham, pp. 661–695*, 2017.

[39] Sanjam Garg, Mohammad Mahmoody, and Ameer Mohammad. When does functional encryption imply obfuscation? In *Theory of Cryptography Conference (TCC), Springer, Cham, pp. 82–115*, 2017.

[40] Craig Gentry. *A fully homomorphic encryption scheme*. Stanford University, 2009.

[41] Vipul Goyal, Yuval Ishai, Mohammad Mahmoody, and Amit Sahai. Interactive locking, zero-knowledge PCPs, and unconditional cryptography. In *Advances in Cryptology - CRYPTO*, volume 6223 of *Lecture Notes in Computer Science*. p. 173–190, Springer, 2010.

[42] Vipul Goyal, Yuval Ishai, Amit Sahai, Ramarathnam Venkatesan, and Akshay Wadia. Founding cryptography on tamper-proof hardware tokens. In *Theory of Cryptography Conference*, pages 308–326. Springer, 2010.

[43] Vipul Goyal, Virendra Kumar, Satyanarayana V. Lokam, and Mohammad Mahmoody. On black-box reductions between predicate encryption schemes. In Ronald Cramer, editor, *Theory of Cryptography (TCC)*, volume 7194 of *Lecture Notes in Computer Science*, pages 440–457. Springer, 2012.

[44] Iftach Haitner, Mohammad Mahmoody, and David Xiao. A new sampling protocol and applications to basing cryptographic primitives on the hardness of NP. In *IEEE Conference on Computational Complexity*, 2010.

[45] Russell Impagliazzo. A personal view of average-case complexity. In *Proceedings of the Tenth Annual Structure in Complexity Theory Conference*, pages 134–147. IEEE Computer Society, 1995.

[46] Russell Impagliazzo and Steven Rudich. Limits on the provable consequences of one-way permutations. In *Proceedings of the 21st Annual ACM Symposium on Theory of Computing (STOC)*, pages 44–61. ACM Press, 1989.

[47] Yuval Ishai, Joe Kilian, Kobbi Nissim, and Erez Petrank. Extending oblivious transfers efficiently. In *Annual International Cryptology Conference*, pages 145–161. Springer, 2003.

[48] Yuval Ishai, Mohammad Mahmoody, and Amit Sahai. On efficient zero-knowledge PCPs. In Ronald Cramer, editor, *Theory of Cryptography Conference, TCC*, volume 7194. Springer, 2012.

[49] Yuval Ishai, Manoj Prabhakaran, and Amit Sahai. Founding cryptography on oblivious transfer - efficiently. In David Wagner, editor, *CRYPTO*, volume 5157 of *Lecture Notes in Computer Science*, pages 572–591. Springer, 2008.

[50] Jonathan Katz. Universally composable multi-party computation using tamper-proof hardware. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 115–128. Springer, 2007.

[51] Jonathan Katz, Amit Sahai, and Brent Waters. Predicate encryption supporting disjunctions, polynomial equations, and inner products. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 146–162. Springer, 2008.

[52] Joe Kilian. Founding cryptography on oblivious transfer. In *Proceedings of the 20th Annual ACM Symposium on Theory of Computing (STOC)*, 988.

[53] Leslie Lamport. Constructing digital signatures from a one-way function. Technical report, 1979.

[54] Mohammad Mahmoody, Dimitrios I. Diochnos, and Saeed Mahloujifar. Adversarial risk and robustness: General definitions and implications for the uniform distribution. In *Conference on Neural Information Processing Systems (NIPS) – to appear*, 2018.

[55] Mohammad Mahmoody, Dimitrios I. Diochnos, and Saeed Mahloujifar. Learning under p-tampering attacks. In *Algorithmic Learning Theory (ALT)*, 2018.

[56] Mohammad Mahmoody and Saeed Mahloujifar. Blockwise p-tampering attacks on cryptographic primitives, extractors, and learners. In *Theory of Cryptography Conference (TCC), Springer, Cham, pp. 245–279*, 2017.

[57] Mohammad Mahmoody, Hemanta K. Maji, and Manoj Prabhakaran. Limits of random oracles in secure computation. In *Proceedings of the 5th Conference on Innovations in Theoretical Computer Science*, ITCS '14, pages 23–34, New York, NY, USA, 2014. ACM.

[58] Mohammad Mahmoody, Hemanta K Maji, and Manoj Prabhakaran. On the power of public-key encryption in secure computation. In *Theory of Cryptography Conference*, pages 240–264. Springer, 2014.

[59] Mohammad Mahmoody, Ameer Mohammad, and Soheil Nematihaj. On the impossibility of virtual black-box obfuscation in idealized models. In *Theory of Cryptography Conference (TCC), Springer, Berlin, Heidelberg, pp. 18–48*, 2016.

[60] Mohammad Mahmoody, Ameer Mohammad, Soheil Nematihaj, Rafael Pass, and abhi shelat. Lower bounds on assumptions behind indistinguishability obfuscation. In *Theory of Cryptography Conference (TCC), Springer, Berlin, Heidelberg, pp. 49–66*, 2016.

[61] Mohammad Mahmoody and Ameer Mohammed. On the power of hierarchical identity-based encryption. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 243–272. Springer, 2016.

[62] Mohammad Mahmoody, Tal Moran, and Salil Vadhan. Time-lock puzzles in the random oracle model. In *Annual Cryptology Conference*, pages 39–50. Springer, 2011.

[63] Mohammad Mahmoody, Tal Moran, and Salil Vadhan. Publicly verifiable proofs of sequential work. In *Proceedings of the 4th conference on Innovations in Theoretical Computer Science*, pages 373–388. ACM, 2013.

[64] Mohammad Mahmoody and Rafael Pass. The curious case of non-interactive commitments–on the power of black-box vs. non-black-box use of primitives. In *Advances in Cryptology–CRYPTO 2012*, pages 701–718. Springer, 2012.

[65] Mohammad Mahmoody and David Xiao. On the power of randomized reductions and the checkability of SAT. In *IEEE Conf. on Computational Complexity*, 2010.

[66] Mohammad Mahmoody and David Xiao. Languages with efficient zero-knowledge PCPs are in SZK. *Theory of Cryptography Conference (TCC)*, 2013.

[67] R. Merkle. C.s. 244 project proposal. In *Facsimile available at http://www.merkle.com/1974.*, 1974.

[68] Daniele Micciancio. Lattice-based cryptography. In *Encyclopedia of Cryptography and Security*, pages 713–715. Springer, 2011.

[69] Daniele Micciancio and Oded Regev. Worst-case to average-case reductions based on Gaussian measures. *SIAM Journal on Computing*, 37(1):267–302, 2007.

[70] Tal Moran, Moni Naor, and Gil Segev. An optimally fair coin toss. In *Theory of Cryptography Conference*, pages 1–18. Springer, 2009.

[71] Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system. 2008.

[72] Adam O'Neill. Definitional issues in functional encryption. Cryptology ePrint Archive, Report 2010/556, 2010. https://eprint.iacr.org/2010/556.

[73] Periklis A. Papakonstantinou, Charles W. Rackoff, and Yevgeniy Vahlis. How powerful are the ddh hard groups? Cryptology ePrint Archive, Report 2012/653, 2012. https://eprint.iacr.org/2012/653.

[74] Rafael Pass and Hoeteck Wee. Black-box constructions of two-party protocols from one-way functions. In *Theory of Cryptography Conference*, pages 403–418. Springer, 2009.

[75] Krzysztof Pietrzak. Simple verifiable delay functions. *IACR Cryptology ePrint Archive*, 2018:627, 2018.

[76] Omer Reingold, Luca Trevisan, and Salil P. Vadhan. Notions of reducibility between cryptographic primitives. In *Theory of Cryptography, First Theory of Cryptography Conference, TCC 2004*, volume 2951 of *Lecture Notes in Computer Science*, pages 1–20. Springer, 2004.

[77] R. L. Rivest, A. Shamir, and D. A. Wagner. Time-lock puzzles and timed-release crypto. Technical report, Cambridge, MA, USA, 1996.

[78] Miklos Santha and Umesh V. Vazirani. Generating quasi-random sequences from semi-random sources. *J. Comput. Syst. Sci.*, 33(1):75–87, 1986.

[79] Adi Shamir. Identity-based cryptosystems and signature schemes. In *Workshop on the theory and application of cryptographic techniques*, pages 47–53. Springer, 1984.

[80] Marten Van Dijk, Craig Gentry, Shai Halevi, and Vinod Vaikuntanathan. Fully homomorphic encryption over the integers. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 24–43. Springer, 2010.

[81] Andrew Chi-Chih Yao. How to generate and exchange secrets. In *Proceedings of the 27th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 162–167. IEEE Computer Society, 1986.