

# Covert Quantum Learning: Privately and Verifiably Learning from Quantum Data

Abhishek Anand, Matthias C. Caro, Ari Karchmer, and Saachi Mutreja

**Motivation.** As quantum technologies progress beyond the NISQ era [1, 2], learning from quantum data produced by quantum sensors and simulators will become central to scientific discovery. Yet most researchers will access such data only indirectly, delegating computations and measurements on a priori unknown quantum data to *public, untrusted but stronger* quantum devices to tackle problems beyond the reach of *private, trusted but weaker* devices. Verifying the correctness of the data as well as ensuring that the learned target and the conclusions drawn from it remain private thus become pivotal challenges. For instance, a researcher using cloud-based quantum sensors to probe material properties may want to keep their findings confidential, even though the raw measurement data is collected on shared infrastructure.

Such a setting of quantum learning from remote data access is incomparable to that of verified (blind) quantum computation (V(B)QC) [3–6]: In V(B)QC, a (mostly/entirely) *classical* client delegates *known* quantum computations on *trusted* input data to an untrusted quantum server. In contrast, we consider a client that aims to process *untrusted, public* data obtained by delegating data collection, rendering V(B)QC insufficient. While recent work has explored interactively verified learning from remote quantum data access [7–9], these lack privacy guarantees. This necessitates new models and algorithmic tools for verifiable private learning from remote data access.

To address privacy in classical delegated data collection, the *covert learning* model was introduced by [10]. Informally, two core objectives of a *covert learning algorithm* are:

1. **Efficient Learning:** The algorithm should efficiently use the oracle access to the unknown concept to obtain good learning guarantees (e.g., as phrased in agnostic PAC-learning [11, 12]).
2. **Covertness (or Privacy):** The algorithm should prevent an eavesdropper from gaining information about the unknown concept.

This model of [10] comes with two caveats: First, it only instantiated covert algorithms under computational hardness assumptions, Second, it is designed specifically to leverage separations between learning with classical membership queries and classical random examples. However, given the plethora of data oracles considered in classical and quantum learning theory—such as (quantum) examples [11–13], (quantum) statistical queries [14, 15], or copies of a quantum state [16–21]—more broadly applicable models for covert learning are needed. To pave the way for covert learning algorithms in quantum learning theory, we thus view the following question as central:

*Is efficient quantum learning with meaningful covertness possible without computational assumptions?*

We formalize this question and answer it affirmatively in two complementary quantum learning-theoretic scenarios. To endow learning algorithms with privacy and verifiability guarantees, we consider learners with two forms of access to the same unknown object: a *public, stronger* oracle whose transcript is visible to an eavesdropper, and a *private, weaker* oracle hidden from the eavesdropper. This reflects a practical and operational hierarchy: obtaining high-quality quantum data (e.g., from multi-copy measurements or phase states) may require powerful cloud services, while simple data (e.g., from single-copy measurements or classical queries) may be collected locally.

**Covert Quadratic Function Learning, Stabilizer State Learning, and Pauli Shadow Tomography.** Statistical queries (SQs) are a fundamental data access model in classical learning [14], and their quantum analogue, the quantum SQ (QSQ) model, has recently attracted significant interest [15, 22–25]. In the SQ model, given the target function  $f$ , the learner submits a bounded query function  $q$  to receive an estimate of  $\mathbb{E}_x[q(x, f(x))]$ . Similarly, in the QSQ model, for an unknown target state  $\rho$ , the learner submits a bounded observable  $M$  to receive an estimate of  $\text{tr}[M\rho]$ . This data access models a scenario where only approximate aggregate statistics of the data rather than individual samples are available. This is especially well suited to noisy quantum hardware, where individual measurement outcomes are unreliable, but some aggregate statistics can be trusted.

Among the most notable results on (Q)SQs are unconditional, exponential separations: learning parities from random examples vs SQs [14] and learning quadratic functions from quantum measurement outcomes vs QSQs [22]. Thus, to efficiently solve these tasks, a learner that privately only has (Q)SQ-access needs to draw on additional public (measurement) samples. Our first contribution formalizes a notion of covert learning tailored to this remote data-access setting and gives algorithms that achieve it. To ensure that the learned hypothesis remains private despite a part of the data being public, we require a covert learner to satisfy:

1.  **$\delta_c$ -Completeness:** Using both the public and the private data access, the learner should identify the unknown function with success probability  $\geq 1 - \delta_c$ .
2.  **$\delta_p$ -Coverttness:** After seeing only the public data collected by the learner, any eavesdropper's probability of identifying the unknown target, over the target distribution and the learner's randomness, is at most  $\delta_p$ .

We achieve these desiderata for learning from public (measurement) examples and private (Q)SQs:

**Theorem 1:** *We give computationally efficient covert algorithms for learning...*

- ...parities with  $\delta_c = 2^{-n}$  and  $\delta_p = 1/\text{poly}(n)$  from  $\text{poly}(n)$  private SQs and  $\mathcal{O}(n)$  public examples, and
- ...quadratic function states with  $\delta_c = \delta_p = 2^{-n}$  from  $n$  private QSQs and  $\mathcal{O}(n)$  public measurements.

These results should be contrasted with a learner who only uses private (Q)SQs to solve the respective learning task [14, 22]. Theorem 1 shows that the learner can be exponentially more query- and even computationally efficient when making use of the more powerful public random/measurement example oracle, while using private SQ/QSQs to ensure that little information about the learned function leaks to any eavesdropper.

Our algorithms here follow a two-step principle: First use the public oracle to reduce the space of possibilities, then privately solve the learning task in the reduced space. Concretely, for parity learning, we first use public examples to shrink the space of consistent parities to  $\log(\text{poly}(n))$  dimensions, where we can efficiently identify the correct one using SQs. For learning a quadratic function described by some square matrix, we first use public Bell sampling to learn all off-diagonal entries [22]. Then, we efficiently learn the diagonal entries using QSQs. These protocols are covert because after seeing the intentionally-incomplete public data, an eavesdropper still has to guess randomly among the remaining possibilities— $\text{poly}(n)$  consistent parities in the former and  $2^n$  potential diagonals in the latter case. This principle extends to other tasks such as learning stabilizer states [26, 27] and Pauli shadow tomography [28–30]:

**Corollary 1:** *We give covert algorithms for...*

- ...computationally efficiently learning stabilizer states with  $\delta_c = \delta_p = 2^{-n}$  from  $\mathcal{O}(n)$  private 1-copy measurements and  $\mathcal{O}(n)$  public 2-copy measurements, and
- ...performing Pauli shadow tomography with  $\delta_c = \delta_p = 2^{-n}$  to accuracy  $1/\text{poly}(n)$  with  $\text{poly}(n)$  private 1-copy measurements and  $\text{poly}(n)$  public 2-copy measurements.

Contrasting this with the lower bounds faced by a learner that only relies on private 1-copy measurements— $\Omega(n^2)$  copies for stabilizer state learning [21] and  $\Omega(2^n/\epsilon^2)$  copies for  $\epsilon$ -accurate Pauli shadow tomography [31]—our covert learner uses the public data access to achieve a sample complexity improvement, while ensuring that any eavesdropper gains only a negligible chance of solving the learning task. Note that not every separation between a public and private oracle yields a covert algorithm: the public queries must be designed so as not to inadvertently reveal the target while still enabling efficient learning when combined with private queries.

**Covert and Verifiable Quantum Data, Forrelation, and Simon's.** In the covert learning models considered above, the oracle queries and responses are purely classical, despite relating to underlying quantum systems. If we consider a fully quantum learning scenario with quantum queries and responses, however, a fundamental rule of quantum physics comes into play: Observing in general perturbs a quantum system. Thus, we have to consider covert *verifiable* learning, where in addition to completeness and coverttness, we also require soundness against

adversaries that may *modify* the states, potentially causing the learner to output an incorrect result. Our second contribution is to formally define and algorithmically realize covert verifiable learning from quantum data.

We focus on quantum data in the form of phase states  $|\psi_f^{\text{Ph}}\rangle \propto \sum_{x \in \{0,1\}^n} (-1)^{f(x)} |x\rangle$  and quantum examples  $|\psi_f^{\text{Ex}}\rangle \propto \sum_{x \in \{0,1\}^n} |x, f(x)\rangle$ , both of which are ubiquitous in quantum algorithms in the oracle model [16–19] as well as in quantum learning theory [13, 20]. First, we exhibit a simple adversarial strategy to show that, information-theoretically, no protocol can simultaneously achieve correctness and covertness for quantum oracle queries against general quantum adversaries in the worst case. Consequently, we identify a physically motivated class of restricted adversaries, *i.i.d. ancilla-free* adversaries, against which we can efficiently, verifiably, and covertly solve Forrelation [19] and Simon’s problem [18] using public quantum queries and private classical queries. This adversary model captures scenarios where the eavesdropper lacks quantum memory and cannot maintain coherence across rounds.

**Theorem 2:** *We give covert verifiable algorithms for efficiently solving Forrelation and Simon’s problem from  $\text{poly}(n)$  public quantum queries and  $\text{poly}(n)$  private classical queries with the following guarantees:*

1. **Completeness:** *With no adversary, the algorithm accepts and solves the problem with probability  $\geq 1 - 2^{-n}$ .*
2. **Soundness:** *With any adversary, the algorithm accepts and fails to solve the problem with probability  $\leq 2^{-n}$ .*
3. **Coverttness:** *Any i.i.d. ancilla-free adversary that extracts information about the unknown function(s) with probability  $1/\text{poly}(n)$  causes the learner to abort with probability  $\geq 1 - 2^{-n}$ .*

Crucially, the algorithms in Theorem 2 have polynomial query and time complexities. Thus, they overcome the exponential classical query complexity lower bounds for Forrelation and Simon’s by using an untrusted, remote quantum data oracle, all while maintaining privacy of the solution. Moreover, even without covertness, completeness and soundness together as in Theorem 2 constitute a new result in interactively verifying quantum learning.

Our core algorithmic contribution is a subroutine for covertly and verifiably acquiring copies of  $|\psi_f^{\text{Ph}}\rangle$  (or  $|\psi_f^{\text{Ex}}\rangle$ ) from public quantum queries and private classical queries. For verification, we employ classical query access to perform shadow overlap-based phase state certification [32], which we lift from its original i.i.d. context to our adversarial setting using tools from [33]. To ensure covertness, the algorithm maximally entangles the public queries with an auxiliary private register. Any eavesdropping attempt necessarily breaks the entanglement, leading to an irrecoverable loss in global fidelity that the learner can detect via shadow overlap estimation.

**Outlook.** As public and shared quantum computing infrastructure matures, we expect privacy-preserving quantum learning to become increasingly relevant; our models and techniques, building on tools from quantum learning, tomography, and certification, can serve as a starting point for a corresponding theory. Our work opens several directions for future research. On the one hand, extending our results to stronger adversary models in bounded and noisy storage settings [34, 35], to non-uniform or noisy quantum data [36–39], and to channel or Hamiltonian learning problems [24, 40–43] are important technical challenges. On the other hand, given connections between classical covert learning and key exchange [10], there is room to illuminate the connections of covert quantum learning to quantum cryptography. We believe that this combination of a newly proposed framework, novel algorithms, and intriguing future directions will be of interest to the QTML audience.

## References

- [1] John Preskill. “Quantum Computing in the NISQ era and beyond”. In: *Quantum* 2 (Aug. 2018), p. 79. ISSN: 2521-327X. DOI: [10.22331/q-2018-08-06-79](https://doi.org/10.22331/q-2018-08-06-79). URL: <https://doi.org/10.22331/q-2018-08-06-79> (page 1).
- [2] John Preskill. “Beyond NISQ: The Megaquop Machine”. In: *ACM Transactions on Quantum Computing* 6.3 (Apr. 2025). DOI: [10.1145/3723153](https://doi.org/10.1145/3723153). URL: <https://doi.org/10.1145/3723153> (page 1).
- [3] Joseph F. Fitzsimons and Elham Kashefi. “Unconditionally Verifiable Blind Computation”. In: *Physical Review A* 96.1 (July 5, 2017), p. 012303. ISSN: 2469-9926, 2469-9934. DOI: [10.1103/PhysRevA.96.012303](https://doi.org/10.1103/PhysRevA.96.012303). arXiv: [1203.5217](https://arxiv.org/abs/1203.5217) (page 1).
- [4] Alexandru Gheorghiu, Theodoros Kapourniotis, and Elham Kashefi. “Verification of quantum computation: An overview of existing approaches”. In: *Theory of Computing Systems* 63.4 (July 2018), pp. 715–808. DOI: [10.1007/s00224-018-9872-3](https://doi.org/10.1007/s00224-018-9872-3). URL: <http://dx.doi.org/10.1007/s00224-018-9872-3> (page 1).
- [5] Anne Broadbent. “How to Verify a Quantum Computation”. In: *Theory of Computing* 14.11 (2018), pp. 1–37. DOI: [10.4086/toc.2018.v014a011](https://doi.org/10.4086/toc.2018.v014a011). URL: <https://theoryofcomputing.org/articles/v014a011> (page 1).
- [6] Urmila Mahadev. “Classical Verification of Quantum Computations”. In: *SIAM Journal on Computing* 51.4 (2022), pp. 1172–1229. DOI: [10.1137/20M1371828](https://doi.org/10.1137/20M1371828). eprint: <https://doi.org/10.1137/20M1371828>. URL: <https://doi.org/10.1137/20M1371828> (page 1).
- [7] Matthias C. Caro, Marcel Hinsche, Marios Ioannou, Alexander Nietner, and Ryan Sweke. “Classical Verification of Quantum Learning”. In: *15th Innovations in Theoretical Computer Science Conference (ITCS 2024)*. Ed. by Venkatesan Guruswami. Vol. 287. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024, 24:1–24:23. ISBN: 978-3-95977-309-6. DOI: [10.4230/LIPIcs.ITCS.2024.24](https://doi.org/10.4230/LIPIcs.ITCS.2024.24). URL: <https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.ITCS.2024.24> (page 1).
- [8] Yinghao Ma, Jiayi Su, and Dong-Ling Deng. *Classical Verification of Quantum Learning Advantages with Noises*. 2024. arXiv: [2411.09210](https://arxiv.org/abs/2411.09210) [quant-ph]. URL: <https://arxiv.org/abs/2411.09210> (page 1).
- [9] Matthias C. Caro, Jens Eisert, Marcel Hinsche, Marios Ioannou, Alexander Nietner, and Ryan Sweke. *Interactive proofs for verifying (quantum) learning and testing*. 2024. arXiv: [2410.23969](https://arxiv.org/abs/2410.23969) [quant-ph]. URL: <https://arxiv.org/abs/2410.23969> (page 1).
- [10] Ran Canetti and Ari Karchmer. “Covert Learning: How to Learn with an Untrusted Intermediary”. In: *Theory of Cryptography Conference*. Springer, 2021, pp. 1–31. DOI: [10.1007/978-3-030-90456-2\\_1](https://doi.org/10.1007/978-3-030-90456-2_1) (pages 1, 3).
- [11] V. N. Vapnik and A. Ya. Chervonenkis. “On the uniform convergence of relative frequencies of events to their probabilities”. In: *Th. Prob. App.* 16 (1971), pp. 264–280. DOI: [10.1137/1116025](https://doi.org/10.1137/1116025) (page 1).
- [12] Leslie G. Valiant. “A theory of the learnable”. In: *Communications of the ACM* 27.11 (1984), pp. 1134–1142. ISSN: 00010782. DOI: [10.1145/1968.1972](https://doi.org/10.1145/1968.1972) (page 1).
- [13] Nader H. Bshouty and Jeffrey C. Jackson. “Learning DNF over the uniform distribution using a quantum example oracle”. In: *Proceedings of the Eighth Annual Conference on Computational Learning Theory*. COLT ’95. Santa Cruz, California, USA: Association for Computing Machinery, 1995, pp. 118–127. ISBN: 0897917235. DOI: [10.1145/225298.225312](https://doi.org/10.1145/225298.225312). URL: <https://doi.org/10.1145/225298.225312> (pages 1, 3).
- [14] Michael Kearns. “Efficient noise-tolerant learning from statistical queries”. In: *Journal of the ACM (JACM)* 45.6 (1998), pp. 983–1006. DOI: [10.1145/293347.293351](https://doi.org/10.1145/293347.293351) (pages 1, 2).
- [15] Srinivasan Arunachalam, Alex B. Grilo, and Henry Yuen. *Quantum statistical query learning*. 2020. arXiv: [2002.08240](https://arxiv.org/abs/2002.08240) [quant-ph] (page 1).
- [16] David Deutsch and Richard Jozsa. “Rapid solution of problems by quantum computation”. In: *Proceedings of the Royal Society of London. Series A: Mathematical and Physical Sciences* 439.1907 (1992), pp. 553–558. DOI: [10.1098/rspa.1992.0167](https://doi.org/10.1098/rspa.1992.0167) (pages 1, 3).
- [17] Ethan Bernstein and Umesh Vazirani. “Quantum complexity theory”. In: *SIAM J. Comput.* 26.5 (1997), pp. 1411–1473. ISSN: 0097-5397. DOI: [10.1137 / S0097539796300921](https://doi.org/10.1137/S0097539796300921). URL: <https://doi.org/10.1137/S0097539796300921> (pages 1, 3).
- [18] Daniel R Simon. “On the power of quantum computation”. In: *SIAM journal on computing* 26.5 (1997), pp. 1474–1483. DOI: [10.1137/S0097539796298637](https://doi.org/10.1137/S0097539796298637) (pages 1, 3).

- [19] Scott Aaronson and Andris Ambainis. “Forrelation: A problem that optimally separates quantum from classical computing”. In: *Proceedings of the forty-seventh annual ACM symposium on Theory of computing*. 2015, pp. 307–316. doi: [10.1145/2746539.2746547](https://doi.org/10.1145/2746539.2746547) (pages 1, 3).
- [20] Srinivasan Arunachalam and Ronald de Wolf. “Guest column: A survey of quantum learning theory”. In: *ACM Sigact News* 48.2 (2017), pp. 41–67. doi: [10.1145/3106700.3106710](https://doi.org/10.1145/3106700.3106710) (pages 1, 3).
- [21] Srinivasan Arunachalam, Sergey Bravyi, Arkopal Dutt, and Theodore J. Yoder. “Optimal Algorithms for Learning Quantum Phase States”. In: *18th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2023)*. Ed. by Omar Fawzi and Michael Walter. Vol. 266. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023, 3:1–3:24. ISBN: 978-3-95977-283-9. doi: [10.4230/LIPIcs.TQC.2023.3](https://doi.org/10.4230/LIPIcs.TQC.2023.3). URL: <https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.TQC.2023.3> (pages 1, 2).
- [22] Srinivasan Arunachalam, Vojtech Havlicek, and Louis Schatzki. “On the Role of Entanglement and Statistics in Learning”. In: *Advances in Neural Information Processing Systems*. Ed. by A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine. Vol. 36. Curran Associates, Inc., 2023, pp. 55064–55076. URL: [https://proceedings.neurips.cc/paper\\_files/paper/2023/file/acb7ce5aab6e134300a2361dd90a501f-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2023/file/acb7ce5aab6e134300a2361dd90a501f-Paper-Conference.pdf) (pages 1, 2).
- [23] Alexander Nietner. *Unifying (Quantum) Statistical and Parametrized (Quantum) Algorithms*. 2023. arXiv: [2310.17716](https://arxiv.org/abs/2310.17716) [quant-ph]. URL: <https://arxiv.org/abs/2310.17716> (page 1).
- [24] Armando Angrisani. “Learning unitaries with quantum statistical queries”. In: *Quantum* 9 (July 2025), p. 1817. ISSN: 2521-327X. doi: [10.22331/q-2025-07-30-1817](https://doi.org/10.22331/q-2025-07-30-1817). URL: <https://doi.org/10.22331/q-2025-07-30-1817> (pages 1, 3).
- [25] Chirag Wadhwa and Mina Doosti. “Learning Quantum Processes with Quantum Statistical Queries”. In: *Quantum* 9 (May 2025), p. 1739. ISSN: 2521-327X. doi: [10.22331/q-2025-05-12-1739](https://doi.org/10.22331/q-2025-05-12-1739). URL: <https://doi.org/10.22331/q-2025-05-12-1739> (page 1).
- [26] Scott Aaronson and Daniel Gottesman. *Identifying stabilizer states*. 2008. URL: <https://pirsa.org/08080052/> (page 2).
- [27] Ashley Montanaro. *Learning stabilizer states by Bell sampling*. 2017. arXiv: [1707.04012](https://arxiv.org/abs/1707.04012) [quant-ph] (page 2).
- [28] Hsin-Yuan Huang, Richard Kueng, and John Preskill. “Information-theoretic bounds on quantum advantage in machine learning”. In: *Physical Review Letters* 126.19 (2021), p. 190505. doi: [10.1103/PhysRevLett.126.190505](https://doi.org/10.1103/PhysRevLett.126.190505) (page 2).
- [29] Robbie King, David Gosset, Robin Kothari, and Ryan Babbush. “Triply Efficient Shadow Tomography”. In: *PRX Quantum* 6.1 (Feb. 2025). ISSN: 2691-3399. doi: [10.1103/prxquantum.6.010336](https://doi.org/10.1103/prxquantum.6.010336). URL: <http://dx.doi.org/10.1103/PRXQuantum.6.010336> (page 2).
- [30] Sitan Chen, Weiyuan Gong, and Qi Ye. “Optimal tradeoffs for estimating Pauli observables”. In: *2024 IEEE 65th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE. 2024, pp. 1086–1105. doi: [10.1109/FOCS61266.2024.00072](https://doi.org/10.1109/FOCS61266.2024.00072) (page 2).
- [31] Sitan Chen, Jordan Cotler, Hsin-Yuan Huang, and Jerry Li. “Exponential separations between learning with and without quantum memory”. In: *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE. 2022, pp. 574–585. doi: [10.1109/FOCS52979.2021.00063](https://doi.org/10.1109/FOCS52979.2021.00063) (page 2).
- [32] Hsin-Yuan Huang, John Preskill, and Mehdi Soleimanifar. “Certifying almost all quantum states with few single-qubit measurements”. In: *Nature Physics* (Sept. 2025). doi: [10.1038/s41567-025-03025-1](https://doi.org/10.1038/s41567-025-03025-1) (page 3).
- [33] Omar Fawzi, Richard Kueng, Damian Markham, and Aadil Oufkir. “Learning properties of quantum states without the IID assumption”. In: *Nature Communications* 15.1 (2024), p. 9677. doi: [10.1038/s41467-024-53765-6](https://doi.org/10.1038/s41467-024-53765-6) (page 3).
- [34] Ivan B Damgård, Serge Fehr, Louis Salvail, and Christian Schaffner. “Cryptography in the bounded-quantum-storage model”. In: *SIAM Journal on Computing* 37.6 (2008), pp. 1865–1890. doi: [10.1137/060651343](https://doi.org/10.1137/060651343) (page 3).
- [35] Stephanie Wehner, Christian Schaffner, and Barbara M Terhal. “Cryptography from noisy storage”. In: *Physical Review Letters* 100.22 (2008), p. 220502. doi: [10.1103/PhysRevLett.100.220502](https://doi.org/10.1103/PhysRevLett.100.220502) (page 3).

- [36] V. Kanade, A. Rocchetto, and S. Severini. “Learning DNFs under product distributions via  $\mu$ -biased quantum Fourier sampling”. In: *Quantum Information & Computation* 19.15&16 (2019), pp. 1261–1278. DOI: [doi.org/10.26421/QIC19.15-16-1](https://doi.org/10.26421/QIC19.15-16-1) (page 3).
- [37] Andrew W. Cross, Graeme Smith, and John A. Smolin. “Quantum learning robust against noise”. In: *Physical Review A* 92.1 (July 2015). ISSN: 1094-1622. DOI: [10.1103/physreva.92.012327](https://doi.org/10.1103/physreva.92.012327). URL: <http://dx.doi.org/10.1103/PhysRevA.92.012327> (page 3).
- [38] Alex B. Grilo, Iordanis Kerenidis, and Timo Zijlstra. “Learning-with-errors problem is easy with quantum samples”. In: *Physical Review A* 99.3 (Mar. 2019). ISSN: 2469-9934. DOI: [10.1103/physreva.99.032314](https://doi.org/10.1103/physreva.99.032314) (page 3).
- [39] Matthias C. Caro. “Quantum learning Boolean linear functions w.r.t. product distributions”. In: *Quantum Information Processing* 19, 172 (2020). DOI: [10.1007/s11128-020-02661-1](https://doi.org/10.1007/s11128-020-02661-1) (page 3).
- [40] Chirag Wadhwa, Laura Lewis, Elham Kashefi, and Mina Doosti. “Agnostic Process Tomography”. In: *PRX Quantum* 6 (4 Dec. 2025), p. 040371. DOI: [10.1103/q2nb-zg9m](https://doi.org/10.1103/q2nb-zg9m). URL: <https://link.aps.org/doi/10.1103/q2nb-zg9m> (page 3).
- [41] Matthias C. Caro. “Learning Quantum Processes and Hamiltonians via the Pauli Transfer Matrix”. In: *ACM Transactions on Quantum Computing* 5.2 (June 2024). DOI: [10.1145/3670418](https://doi.org/10.1145/3670418). URL: <https://doi.org/10.1145/3670418> (page 3).
- [42] Ainesh Bakshi, Allen Liu, Ankur Moitra, and Ewin Tang. “Structure Learning of Hamiltonians from Real-Time Evolution”. In: *2024 IEEE 65th Annual Symposium on Foundations of Computer Science (FOCS)*. 2024, pp. 1037–1050. DOI: [10.1109/FOCS61266.2024.00069](https://doi.org/10.1109/FOCS61266.2024.00069) (page 3).
- [43] Hong-Ye Hu, Muzhou Ma, Weiyuan Gong, Qi Ye, Yu Tong, Steven T. Flammia, and Susanne F. Yelin. *Ansatz-free Hamiltonian learning with Heisenberg-limited scaling*. 2025. arXiv: [2502.11900](https://arxiv.org/abs/2502.11900) [quant-ph]. URL: <https://arxiv.org/abs/2502.11900> (page 3).