

Forecasting quantum observables: A compressed sensing approach with performance guarantees

Víctor Valls ^{1,*} Albert Akhriev ¹ Olatz Sanz Larrarte ² Javier Oliva del Moral ^{2,3} Štěpán Šmíd ^{4,5}
 Josu Etxezarreta Martinez ² Sergiy Zhuk ¹ and Dmytro Mishaghi ⁵

¹IBM Quantum, IBM Research Europe, Trinity Business School, Dublin, D02 F6N2, Ireland

²Department of Basic Sciences, Tecnun—University of Navarra, 20018 San Sebastián, Spain

³Donostia International Physics Center, 20018 San Sebastián, Spain

⁴Department of Computing, Imperial College London, London, United Kingdom

⁵IBM Quantum, IBM Research Europe, Hursley, Winchester SO21 2JN, United Kingdom



(Received 24 December 2025; accepted 3 May 2026; published 26 August 2026)

Data-driven extrapolation methods aim to extend the dynamics of quantum observables from measurements, but they often lack guarantees on prediction accuracy. We introduce a framework based on atomic norm minimization that can certify whether the spectral model learned by a forecasting algorithm—i.e., Bohr frequencies and amplitudes—is consistent with unitary quantum time evolution. Certification holds when the dynamics are governed by a small number of well-separated Bohr frequencies. We validate the approach on multiple forecasting algorithms applied to spin-chain Hamiltonians with 8–20 sites. Comparing with exact diagonalization, certified models yield an average forecasting error below 0.1 (observable range $[-1, 1]$) in 97% of cases and below 0.05 in 91%–99% of cases. Even in the presence of noise, certified models remain robust at the 0.1 error threshold.

DOI: [10.1103/9qp3-33cs](https://doi.org/10.1103/9qp3-33cs)

Introduction. Quantum spin Hamiltonians are natural test beds for near-term quantum devices because their local interaction structure maps directly onto existing qubit connectivity [1]. Recent experiments have already probed dynamics with more than 100 qubits and circuit depths of several hundred two-qubit gates [2,3]. However, the time window for reliable quantum time evolution (TE) remains limited: Gate errors accumulate rapidly [4], and error mitigation is effective only for short times, as its cost grows exponentially [2].

To overcome this limitation, a growing body of work seeks to infer long-time quantum dynamics from early, low-noise measurements [5–8]. These *data-driven* approaches exploit the fact that expectation values can be expressed as the sum of harmonics, enabling extrapolation that is independent of the system size and qubit connectivity. Methods such as ESPRIT and DMD leverage this structure and have shown that they can extrapolate observables well beyond the measured time window [5,7].

Despite their empirical success, existing forecasting methods cannot certify that the extrapolated observables are consistent with the underlying dynamics, i.e., whether the spectral model captures the Bohr frequencies and amplitudes that govern the unitary time evolution. In this work, rather than proposing a forecasting algorithm and assessing its accuracy empirically, we ask a more fundamental

question: whether a proposed spectral model captures the underlying unitary quantum time evolution dynamics. Addressing this gap requires moving beyond forecasting alone and introducing a systematic framework for validating candidate models.

In this Letter, we introduce such a framework using tools from atomic norm minimization (ANM) [9,10], a compressed sensing approach for sparse spectral recovery with rigorous theoretical guarantees [9,11]. In particular, we exploit the dual characterization of ANM to construct a certification test: Given early-time measurement data and a candidate spectral model, we certify whether that model is compatible with the underlying spectral structure implied by the data. This certification provides a consistency check for reliable forecasting, where correct certification is guaranteed by ANM when the Bohr frequencies are sparse and sufficiently separated, and early-time measurements are low noise (i.e., the time window where it is possible to obtain reliable measurements from quantum devices and apply error mitigation effectively). Importantly, the certification framework is independent of the forecasting algorithms, and it can be used to validate candidate models obtained by different forecasting methods, including existing methods in the literature (e.g., Refs. [5,7]).

The main contributions of this work are as follows:

(1) We introduce a certification framework based on ANM that tests the consistency of candidate time evolution models used for forecasting quantum observables (see the “Certification via atomic norm minimization: Certification pipeline” section).

(2) We validate the framework on spin-chain time evolutions with 8–20 spin sites, using multiple forecasting algorithms on both noiseless and noisy data (see the “Experiments” section).

*Contact author: victor.valls@ibm.com




FIG. 1. Amplitudes $\Re(c(f))$ for TE under $H = \sum_{i=1}^n \sigma_i^z \sigma_{i+1}^z + \sum_{i=1}^n \sigma_i^x$, with $n = 8$ qubits, two initial states, and periodic boundary conditions (site $n + 1 = 1$). The imaginary part of $c(f)$ is zero for both cases.

Problem statement. Consider a quantum system with n qubits (sites) that evolves according to a time-independent Hamiltonian $H \in \mathbb{C}^{2^n \times 2^n}$. Let $|\psi(0)\rangle \in \mathbb{C}^{2^n}$ be an initial state vector. The quantum time evolution of the state vector is given by $|\psi(\tau_k)\rangle = e^{-iH\tau_k} |\psi(0)\rangle$, where $\tau_k = \Delta(k-1)$ is a discrete time, $\Delta > 0$, and $k \in [m] := \{1, \dots, m\}$, with m the total number of time steps. The k th measurement of an observable $O \in \mathbb{C}^{2^n \times 2^n}$ is given by

$$y_k := \langle \psi(\tau_k) | O | \psi(\tau_k) \rangle = \sum_{f \in \Omega} c(f) e^{i2\pi f \tau_k}, \quad (1)$$

where $c(f) \in \mathbb{C}$ and $f \in \Omega \subset [-\frac{1}{2}, \frac{1}{2}]$ is a Bohr frequency, i.e., normalized difference of the eigenvalues of H .

The set Ω contains only the frequencies with nonzero amplitude, which depend on the spectral information of H , the initial state $|\psi(0)\rangle$, and observable O . Note that Ω is generally unknown for large systems since we cannot compute the spectral information of H . Figure 1 shows the amplitudes $\Re(c(f))$ for all $f \in \Omega$ for a transverse field Ising model (TFIM) Hamiltonian with two initial states and the same observable. Note that the number of nonzero amplitudes varies, with panel (a) having fewer and more spread out frequencies than panel (b). Also, note that the spectrum is symmetric since observables are real valued.

The TE forecasting problem we consider is the following. For m given measurements $y := (y_1, \dots, y_m) \in [-1, 1]^m$ and tolerance $\epsilon \geq 0$, find a sequence x_k with $k \in [m']$ such that

$$\frac{1}{m'} \sum_{k=1}^{m'} |x_k - \langle \psi(\tau_{m+k}) | O | \psi(\tau_{m+k}) \rangle| \leq \epsilon. \quad (2)$$

That is, the goal is to predict the m' future values of the time series based on the m available measurements.

Certification via atomic norm minimization. In this section, we briefly review ANM (see the “Certification via atomic norm minimization: Atomic norm minimization” section) and describe how it can be used to certify the frequencies and amplitudes recovered by any spectral estimation algorithm (see the “Certification via atomic norm minimization: Certification pipeline” section).

Certification via atomic norm minimization: Atomic norm minimization. ANM aims to minimize the *atomic norm*: $\|y\|_{\mathcal{A}} := \inf\{t \geq 0 : y \in t \cdot \text{conv}(\mathcal{A})\}$, where $\mathcal{A}$ is a collec-

tion of “atoms” (vectors) and $\text{conv}(\mathcal{A})$ its convex hull. The atomic norm can be seen as the continuous analog of the ℓ_1 norm, since $\mathcal{A}$ may contain an uncountable set of atoms corresponding to all possible continuous-frequency sinusoidal components.

For our TE forecasting problem, $\mathcal{A}$ contains vectors of the form

$$a(f) := (e^{i2\pi f \tau_1}, \dots, e^{i2\pi f \tau_m}), \quad (3)$$

where $f \in [-\frac{1}{2}, \frac{1}{2}]$, and the ANM finds the sparsest combination of these continuous atoms that match the measurements vector y . In some cases, this representation is *unique*, enabling exact recovery of the true frequencies and amplitudes underlying y . For the representation to be unique, two conditions must be met: The signal is sparse (i.e., consisting of few frequencies with respect to the number of data points m) and the frequencies are sufficiently separated (i.e., $|f_i - f_j| \geq \frac{4}{m}$ for all $f_i, f_j \in \Omega$ with $i \neq j$); see details in Theorem 1.2 of Ref. [11] and Theorem 1.1 of Ref. [12].

In practice, we can solve the ANM with atoms of the form in Eq. (3) via a semidefinite program (SDP):

$$\underset{T, t > 0}{\text{minimize}} \quad \frac{1}{2m} \text{Tr}(T) + \frac{1}{2}t \quad \text{s.t.} \quad \begin{bmatrix} T & y \\ y^* & t \end{bmatrix} \succeq 0, \quad (4)$$

where $T \in \mathbb{C}^{m \times m}$ is a Hermitian Toeplitz matrix that admits a decomposition $T = \sum_{f \in F} w(f) a(f) a(f)^*$ (see Lemma 2.2 of Ref. [12]), where $w(f) > 0$ and $F \subset [-\frac{1}{2}, \frac{1}{2}]$ denotes the support set of active frequencies. In brief, the constraint in Eq. (4) ensures that the matrix is positive semidefinite and that y lies in the subspace spanned by vectors $a(f)$. The objective promotes low-rank T , thus enforcing sparsity in the frequency domain. Once the optimal T is found, the frequencies can be recovered from T using Prony’s method [13].

Crucially, the correctness of these recovered frequencies can be certified by solving the associated dual problem:

$$\underset{q \in \mathbb{C}^m}{\text{maximize}} \quad \langle y, q \rangle \quad \text{s.t.} \quad \sup_{f \in [-\frac{1}{2}, \frac{1}{2}]} |Q(f)| \leq 1, \quad (5)$$

where $Q(f) := \langle a(f), q \rangle$ is the *dual polynomial*. The solution is certified as *unique* when the dual polynomial reaches its maximum value $|Q(f)| = 1$ precisely for all $f \in \Omega$ and $|Q(f)| < 1$ otherwise (see Proposition 2.4 of Ref. [12]). An example of this dual certificate is shown in Fig. 2 for the setting used in Fig. 1(a), where $Q(f)$ is real valued and touches ± 1 precisely at the true frequencies (indicated with markers).

Finally, we note that solving the ANM with the SDP formulation in Eq. (4) is computationally impractical for a large number of measurements (see example in the Supplemental Material [14], which includes Ref. [15]). In the following section, we show how ANM provides a principled way to certify whether, under sparsity and frequency separation conditions, a candidate model accurately reproduces the true dynamics in Eq. (1).

Certification via atomic norm minimization: Certification pipeline. The proposed certification pipeline is illustrated in Fig. 3. An algorithm (e.g., ESPRIT [5]) takes as input a vector y with m measurements and returns candidate frequencies Ω and amplitudes $c(f)$ such that $y_k \approx \sum_{f \in \Omega} c(f) e^{i2\pi f \tau_k}$, $k \in$




FIG. 2. Illustration of the dual polynomial $Q(f)$ (green line) for the setting in Fig. 1(a). The quantum TE is generated over 10 s with a sampling rate of five steps per second (i.e., $m = 50$, $\Delta = 1/5$). The purple bars show $\Re(c(f))$ as in Fig. 1(a), where only the $\Re(c(f))$ larger than 0.005 are shown. The orange markers indicate the sign of $\Re(c(f))$, which lie on values where $Q(f) = \pm 1$.

$\{1, \dots, m\}$. These candidate frequencies and amplitudes are then certified by solving the ANM dual problem in Eq. (5).

Certificates: To certify the model, we evaluate three criteria: the duality gap, frequency separation, and the dual polynomial. These three criteria coming from ANM (see the “Certification via atomic norm minimization: Atomic norm minimization” section) are sufficient for certifying a model.

The duality gap measures the difference between the primal and dual values [Eqs. (4) and (5)]. The gap is always non-negative and, due to strong duality, is zero when both problems are solved optimally. We can obtain an upper bound on the optimal primal objective by constructing $T = \sum_{f \in \Omega} |c(f)|a(f)a(f)^*$, and then setting $t = y^*T^\dagger y$. This construction ensures that the constraints in Eq. (4) are satisfied, by the Schur complement condition. A lower bound on the dual value can be obtained by solving the dual problem approximately (as discussed later). The duality gap is considered admissible if this is below a prescribed tolerance.

For frequency separation, we need to check that the minimum distance between the candidate frequencies in Ω exceeds $4/m$, which is required for unique recovery in ANM.

Finally, the dual polynomial has to satisfy $|Q(f)| \approx 1$ for all $f \in \Omega$ and $|Q(f)| < 1$, otherwise, within numerical tolerances. This requirement comes from the dual problem in Eq. (5). Computing the dual polynomial is straightforward once we have a solution to the dual problem.




FIG. 3. Schematic illustration of the proposed certification pipeline. The process verifies frequency and amplitude estimates obtained from any forecasting algorithm using the dual formulation of ANM.

Solving the dual problem: The dual problem in Eq. (5) is infinite dimensional and must be discretized to make it tractable. If we sample the frequency domain finely, we can transform the problem into a second-order cone program. However, because the measurement vector y is real valued (which is always the case since observables are real), the dual problem can be further simplified to the linear program

$$\max_{q \in \mathbb{R}^m} \langle y, q \rangle \quad \text{s.t.} \quad -1 \leq \langle a(f), q \rangle \leq 1, \quad \forall f \in F, \quad (6)$$

where F is a finely discretized set of frequencies in $[0, \frac{1}{2}]$, since the frequency spectrum is symmetric. Finally, although discretization introduces a small approximation error, it remains negligible when the frequency grid is sufficiently fine. In the Supplemental Material [14], we show how selecting different grids affects the accuracy of solving the dual problem.

Experiments. This section evaluates the proposed ANM certification methodology using several forecasting algorithms and spin-chain time evolutions. Importantly, certification is carried out without knowledge of whether the underlying dynamics consist of well-separated Bohr frequencies, since this is something we typically do not know in practice.

Experiments: Setup. Datasets: We consider quantum TEs generated by 15 spin-chain Hamiltonians (4 TFIM and 11 Heisenberg) with 8–20 qubits and periodic boundary conditions [16–25] (see details in the Supplemental Material [14]). For all the experiments, we have used five initial states that are typically used in the literature [26–31] and are easy to prepare in practice (Néel, dimerized, paramagnetic, ferromagnetic, and the zero state). Also, we measure six observables (X , Y , Z , XX , YY , and ZZ). This results in a total of 5850 TEs. The TEs are generated using Qutip [32] for 40 time units with $dt = 0.2$, yielding $m = 200$ measurements per TE. A time unit is $t \times J$, where J is the coupling strength (interaction energy) in a Hamiltonian.

In the experiments, we consider exact TEs and TEs with added unbiased noise, equivalent to taking 1000 shots. Importantly, the added noise captures uncertainty around the true expectation, and so can be used to model multiple sources of uncertainty, including measurement errors and the residual uncertainty from gate errors after error mitigation. Recall that error mitigation techniques usually trade a reduction in the bias for an increase in variance [33], which is, in essence, the uncertainty around the true value.

Spectral estimation algorithms: We use five algorithms to compute the frequencies and amplitudes: SDP, Prony, OMP, DMD, and ESPRIT. SDP corresponds to solving the problem in Eq. (4). Prony [34] is a method that can provably recover the true frequencies and amplitudes under ideal conditions but is highly sensitive to noise. OMP [35,36] is a widely used heuristic in signal processing for sparse signal recovery. DMD [37] is the forecasting algorithm used in Ref. [7], and ESPRIT [38,39] is a well-known algorithm in spectrum estimation, also applied in observable forecasting [5]. The algorithms are implemented with default or minimal tuning settings, as our focus is on evaluating the certification methodology. A brief description of each algorithm is provided in the Supplemental Material [14] (which includes Refs. [40–44]), and the code is available online [45].

TABLE I. Certification reliability (in percentage) across different forecasting algorithms and tolerances ϵ . The spectral models have been certified with the pipeline in the “Certification via atomic norm minimization: Certification pipeline” section.

ϵ	Exact TE					TE with added noise				
	SDP	PNY	OMP	DMD	ESP	SDP	PNY	OMP	DMD	ESP
0.01	97.7	83.4	75.1	85.5	65.3	37.8	50.2	20.2	41.4	16.0
0.05	99.2	97.4	96.4	96.2	91.5	53.3	58.3	69.9	47.4	58.2
0.10	99.7	98.7	99.2	97.6	98.7	83.3	82.2	95.1	68.1	86.9

Forecasting procedure: We run every forecasting algorithm (e.g., OMP) with the first $k \in \{10, \dots, 160\}$ measurements and evaluate the forecasting accuracy with the remaining $200 - k$ measurements. The forecasting accuracy is measured as indicated in Eq. (2).

Certification: For each algorithm run, ANM certification is performed (150 times per algorithm and TE). This is done by discretizing the constraints in Eq. (6) using a frequency grid with 1000 uniformly spaced points in $[0, \frac{1}{2}]$, augmented with the candidate frequencies estimated by the algorithm—to ensure that the dual polynomial is evaluated at the exact/candidate frequencies. Note that it is sufficient to consider $[0, \frac{1}{2}]$ instead of $[-\frac{1}{2}, \frac{1}{2}]$ since the frequency spectrum is symmetric. To account for numerical inaccuracies, empirically chosen tolerances are applied. Specifically, for exact TEs, the dual polynomial must satisfy $|Q(f)| \geq 0.98$ at all candidate frequencies, and the duality gap must be smaller than 0.05. For TEs with added noise, the thresholds are $|Q(f)| \geq 0.95$ and a duality gap smaller than 0.5. In addition, we only consider amplitudes with magnitude larger than 0.025 for exact TEs and 0.05 for TEs with added noise. The implementation details are available in the Supplemental Material [14], which includes Ref. [46].

Experiments: Results. We evaluate forecasting performance using tolerances $\epsilon \in \{0.01, 0.05, 0.1\}$ [see Eq. (2)], for both exact TEs and TEs with added noise. In both cases, we consider the output of a spectral algorithm (i.e., frequencies and amplitudes) *certified* if it meets the criteria in the “Certification via atomic norm minimization: Certification pipeline” section, and we say a certification is *successful* when the forecasting error [as defined in Eq. (2)] is below the desired tolerance ϵ for the remaining $200 - k$ measurements. We define *certification reliability* as the percentage of successful certification across *all certified runs* (successful or not). Namely, how much we can trust that a certified model will generate forecast values with error smaller than or equal to ϵ . In the following, we present certification reliability depending on the number of measurements used to learn the model.

Certification reliability (average): Table I reports the certification reliability for each forecasting algorithm and tolerance ϵ when algorithms take as input the first $k \in \{10, \dots, 160\}$ measurements. For exact TEs, all algorithms achieve high reliability, exceeding 97% and 91% for ϵ equal to 0.1 and 0.05, respectively. With $\epsilon = 0.01$, the performance of all algorithms decreases notably (especially DMD and OMP) with the exception of SDP, which has a reliability of

TABLE II. Percentage of *successful* certifications over all runs, for different forecasting algorithms and tolerances ϵ . The row with $\epsilon = \infty$ indicates the coverage: the percentage of certifications (successful or not) over all runs.

ϵ	Exact TE					TE with added noise				
	SDP	PNY	OMP	DMD	ESP	SDP	PNY	OMP	DMD	ESP
0.01	16.4	15.8	11.2	16.7	18.6	5.0	6.7	7.1	7.1	7.1
0.05	16.6	18.4	14.3	18.8	26.1	7.1	7.7	24.5	8.0	25.9
0.10	16.7	18.6	14.7	19.1	28.2	11.1	10.9	33.4	11.5	38.8
∞	16.8	18.9	14.9	19.6	28.5	13.3	13.2	35.1	16.9	44.6

97.7%. The better performance of SDP is because it solves the problem in Eq. (4) directly. When noise is introduced, reliability decreases across all tolerances and forecasting algorithms, with only OMP achieving a reliability over 95% for $\epsilon = 0.1$.

One way to increase the certification reliability is to reduce the coverage (the fraction of certified runs), which can be controlled with the numerical tolerance thresholds used in the certification. However, changing these thresholds creates a trade-off: Stricter thresholds improve reliability by excluding borderline cases, but reduce coverage because fewer runs meet the criteria. Table II shows the percentage of successful certifications for different values of ϵ including the coverage ($\epsilon = \infty$). Observe that the successful certification rates decrease as ϵ gets smaller and that the values vary across algorithms. For example, Prony and ESPRIT have both a certification reliability of 98.7% with $\epsilon = 0.1$ and exact TE (see Table I), but their successful certification rate of ESPRIT is 50% higher than Prony (Table II). In practice, this trade-off implies that certification results should always be reported together with coverage, as high reliability at small ϵ may reflect only a subset of instances.

Certification reliability using a small number of measurements: In practice, we are often interested in predicting the TE dynamics using a small number of measurements. Figure 4 shows the distribution of the *average* forecasting error [as defined in Eq. (2)] at the *first* certification (the smallest $k \in \{10, \dots, 160\}$ that meets the criteria in the “Certification via atomic norm minimization: Certification pipeline” section).




FIG. 4. Distribution of the forecasting error upon the first certification (i.e., using the least amount of measurements necessary for the forecasting model to satisfy certification criteria), with exact and noisy TEs.




FIG. 5. Distribution of the forecasting error for certification after at least 30 measurements, with exact and noisy TEs.

Overall, the algorithms exhibit similar behavior. For instance, with exact TEs, Prony and OMP achieve errors below 0.05 and 0.1 in approximately 50% and 75% of cases, respectively. For noisy TEs, the proportion of forecasts with errors below 0.05 drops sharply, and both Prony and OMP achieve errors below 0.1 in only about 50% of cases.

These results contrast with the results in Table I, possibly due to insufficient measurements used during spectral estimation or certification. To investigate this, Fig. 5 shows forecasting errors for first certifications issued only after 30 or more measurements. Performance improves noticeably: For exact TEs, Prony achieves errors below 0.05 and 0.1 in 83.4% and 93.4% of cases, respectively, while OMP achieves the same thresholds in 76% and 100% of cases. For noisy TEs, all algorithms increase the proportion of forecasts with errors exceeding 0.1 (cf. Figs. 4 and 5), with OMP achieving errors below 0.1 and 0.15 in 75% and 97% of cases, respectively.

These results indicate that a sufficiently large number of measurements is crucial for reliable certification; while the minimum is unknown, monitoring stability, as suggested in Sec. III C of Ref. [5], or setting a minimum measurement threshold (e.g., $k \geq 30$) can significantly enhance performance.

Conclusions. We have considered the problem of certifying spectral models used for forecasting quantum observables from measurement data. Our proposed certification framework is based on ANM and has guarantees when the time evolution is governed by a few well-separated Bohr frequencies. Numerical results show that the certified models can forecast reliably on both noiseless and noisy data, even without prior knowledge of the underlying spectral properties of the time evolution. For instance, OMP with a 0.1 error threshold achieves forecasting reliabilities of 99% and 95% on noiseless and noisy data, respectively.

The experiments carried out focus on systems with 8–20 sites (qubits), allowing direct comparison with exact diagonalization, but our certification framework can be applied directly to larger, classically intractable systems. As an example, we show in the Supplemental Material [14] (which includes Refs. [47–52]) that for a tilted TFIM with 100 qubits, OMP yields a spectral model that can be certified and used for reliable forecasting. Conversely, certification may fail for systems of arbitrary size if the dynamics are not sufficiently sparse relative to the number of measurements—a limitation that is connected to the spread of quantum information. An interesting direction for future work is to investigate how partial knowledge of the Hamiltonian spectrum can inform the sparsity of the time evolution and, in turn, the number of measurements required for reliable certification.

Acknowledgments. We would like to thank the members of the Quantum Information Lab at Tecnum, and Sergey Bravyi, Ewout van den Berg, Kristan Temme, Niall Robertson, and Nicola Mariella from IBM Quantum. This work has been supported by the BasQ strategy of the Department of Science, Universities, and Innovation of the Basque Government through the “Extrapolation of Von Neumann Dynamics beyond the reach of current Utility Scale Devices (VNDEXUSD)” project.

Data availability. The code used in the numerical experiments is openly available [45].

- [1] B. Fauseweh, Quantum many-body simulations on digital quantum computers: State-of-the-art and future challenges, *Nat. Commun.* **15**, 2123 (2024).
- [2] Y. Kim, A. Eddins, S. Anand, K. X. Wei, E. van den Berg, S. Rosenblatt, H. Nayfeh, Y. Wu, M. Zaletel, K. Temme, and A. Kandala, Evidence for the utility of quantum computing before fault tolerance, *Nature (London)* **618**, 500 (2023).
- [3] E. D. Switzer, N. Robertson, N. Keenan, Á. Rodríguez, A. D’Urbano, B. Pokharel, T. S. Rahman, O. Shtanko, S. Zhuk, and N. Lorente, Realization of two-dimensional discrete time crystals with anisotropic Heisenberg coupling, *Nat. Commun.* **17**, 605 (2026).
- [4] K. Temme, S. Bravyi, and J. M. Gambetta, Error mitigation for short-depth quantum circuits, *Phys. Rev. Lett.* **119**, 180509 (2017).
- [5] A. Erpenbeck, Y. Zhu, Y. Yu, L. Zhang, R. Gerum, O. Goulko, C. Yang, G. Cohen, and E. Gull, Compact representation and long-time extrapolation of real-time data for quantum systems using the ESPRIT algorithm, *Phys. Rev. B* **113**, 115129 (2026).
- [6] K. Manos, M. Weilenmann, and M. Navascues, Extrapolation of quantum measurement data, [arXiv:2507.06912v1](https://arxiv.org/abs/2507.06912v1).
- [7] R. Kaneko, M. Imada, Y. Kabashima, and T. Ohtsuki, Forecasting long-time dynamics in quantum many-body systems by dynamic mode decomposition, *Phys. Rev. Res.* **7**, 013085 (2025).
- [8] Y. Shen, A. Buzali, H.-Y. Hu, K. Klymko, D. Camps, S. F. Yelin, and R. Van Beeumen, Efficient measurement-driven eigenenergy estimation with classical shadows, [arXiv:2409.13691](https://arxiv.org/abs/2409.13691).
- [9] B. N. Bhaskar, G. Tang, and B. Recht, Atomic norm denoising with applications to line spectral estimation, *IEEE Trans. Signal Process.* **61**, 5987 (2013).
- [10] V. Chandrasekaran, B. Recht, P. A. Parrilo, and A. S. Willsky, The convex geometry of linear inverse problems, *Found. Comput. Math.* **12**, 805 (2012).

- [11] E. J. Candès and C. Fernandez-Granda, Towards a mathematical theory of super-resolution, *Commun. Pure Appl. Math.* **67**, 906 (2014).
- [12] G. Tang, B. N. Bhaskar, P. Shah, and B. Recht, Compressed sensing off the grid, *IEEE Trans. Inf. Theory* **59**, 7465 (2013).
- [13] G. Plonka and M. Tasche, Prony methods for recovery of structured functions, *GAMM-Mitteilungen* **37**, 239 (2014).
- [14] See Supplemental Material at <https://link.aps.org/supplemental/10.1103/9qp3-33cs> for the forecasting algorithms and dataset used in the experiments.
- [15] Mosek ApS, MOSEK optimization software, version 11.0, 2025, <https://www.mosek.com>.
- [16] S. R. White and D. A. Huse, Numerical renormalization-group study of low-lying eigenstates of the antiferromagnetic $s = 1$ Heisenberg chain, *Phys. Rev. B* **48**, 3844 (1993).
- [17] Honecker and Wessel, Magnetocaloric effect in quantum spin- s chains, *Condens. Matter Phys.* **12**, 399 (2009).
- [18] T. Giamarchi, *Quantum Physics in One Dimension* (Oxford University Press, Oxford, 2003).
- [19] H. J. Schulz, Phase diagrams and correlation exponents for quantum spin chains of arbitrary spin quantum number, *Phys. Rev. B* **34**, 6372 (1986).
- [20] I. Affleck, Quantum spin chains and the Haldane gap, *J. Phys.: Condens. Matter* **1**, 3047 (1989).
- [21] M. den Nijs and K. Rommelse, Preroughening transitions in crystal surfaces and valence-bond phases in quantum spin chains, *Phys. Rev. B* **40**, 4709 (1989).
- [22] J. Kurmann, H. Thomas, and G. Müller, Antiferromagnetic long-range order in the anisotropic quantum spin chain, *Physica A* **112**, 235 (1982).
- [23] A. K. Bera, B. Lake, A. T. M. N. Islam, B. Klemke, E. Faulhaber, and J. M. Law, Field-induced magnetic ordering and single-ion anisotropy in the quasi-one-dimensional Haldane chain compound $\text{SrNi}_2\text{V}_2\text{O}_8$: A single-crystal investigation, *Phys. Rev. B* **87**, 224423 (2013).
- [24] B. Nachtergaele, W. Spitzer, and S. Starr, Ferromagnetic ordering of energy levels, *J. Stat. Phys.* **116**, 719 (2004).
- [25] S. M. Giampaolo, G. Adesso, and F. Illuminati, Theory of ground state factorization in quantum cooperative systems, *Phys. Rev. Lett.* **100**, 197201 (2008).
- [26] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information: 10th Anniversary Edition* (Cambridge University Press, Cambridge, 2010), 2nd ed.
- [27] M. Brockmann, B. Wouters, D. Fioretto, J. D. Nardis, R. Vlijm, and J.-S. Caux, Quench action approach for releasing the Néel state into the spin-1/2 XXZ chain, *J. Stat. Mech.* (2014) P12009.
- [28] M. Mestyán, B. Pozsgay, G. Takács, and M. A. Werner, Quenching the XXZ spin chain: Quench action approach versus generalized Gibbs ensemble, *J. Stat. Mech.* (2015) P04001.
- [29] V. Alba and P. Calabrese, Entanglement dynamics after quantum quenches in generic integrable systems, *SciPost Phys.* **4**, 017 (2018).
- [30] B. Žunkovič, M. Heyl, M. Knap, and A. Silva, Dynamical quantum phase transitions in spin chains with long-range interactions: Merging different concepts of nonequilibrium criticality, *Phys. Rev. Lett.* **120**, 130601 (2018).
- [31] J. Richter, T. Heitmann, and R. Steinigeweg, Quantum quench dynamics in the transverse-field Ising model: A numerical expansion in linked rectangular clusters, *SciPost Phys.* **9**, 031 (2020).
- [32] J. R. Johansson, P. D. Nation, and F. Nori, Qutip: An open-source Python framework for the dynamics of open quantum systems, *Comput. Phys. Commun.* **183**, 1760 (2012).
- [33] Z. Cai, Quantum error mitigation, *Rev. Mod. Phys.* **95**, 045005 (2023).
- [34] T. Sauer, Prony’s method: An old trick for new problems, in *Snapshots of Modern Mathematics from Oberwolfach*, edited by D. E. Bruschi, M. Firsching, and C. Cederbaum (Snapshots of modern mathematics from Oberwolfach, Oberwolfach, Germany, 2018).
- [35] S. S. Chen, D. L. Donoho, and M. A. Saunders, Atomic decomposition by basis pursuit, *SIAM Rev.* **43**, 129 (2001).
- [36] F. Locatello, R. Khanna, M. Tschannen, and M. Jaggi, A unified optimization view on generalized matching pursuit and Frank-Wolfe, in *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, edited by A. Singh and J. Zhu (PMLR, 2017), Vol. 54, pp. 860–868.
- [37] A. Beck and M. Teboulle, A fast iterative shrinkage-thresholding algorithm for linear inverse problems, *SIAM J. Imaging Sci.* **2**, 183 (2009).
- [38] A. Paulraj, R. Roy, and T. Kailath, Estimation of signal parameters via rotational invariance techniques—ESPRIT, in *Proceedings of the 19th Asilomar Conference on Circuits, Systems and Computers* (1985), pp. 984–995.
- [39] R. Roy and T. Kailath, ESPRIT—estimation of signal parameters via rotational invariance techniques, *IEEE Trans. Acoust., Speech, Signal Process.* **37**, 984 (1989).
- [40] P. Stoica and R. L. Moses, *Spectral Analysis of Signals* (Pearson Prentice Hall, Upper Saddle River, NJ, 2005), Vol. 452.
- [41] M. Wax and T. Kailath, Detection of signals by information theoretic criteria, *IEEE Trans. Acoust., Speech, Signal Process.* **33**, 387 (1985).
- [42] N. Golyandina and N. Zvonarev, Information criteria for signal extraction using singular spectrum analysis: White and red noise, *Algorithms* **17**, 395 (2024).
- [43] N. Golyandina, V. Nekrutkin, and A. A. Zhigljavsky, Analysis of time series structure—SSA and related techniques, in *Monographs on Statistics and Applied Probability* (Chapman & Hall/CRC, Boca Raton, FL, USA, 2001).
- [44] Y. Hua and T. Sarkar, Matrix pencil method for estimating parameters of exponentially damped/undamped sinusoids in noise, *IEEE Trans. Acoust., Speech, Signal Process.* **38**, 814 (1990).
- [45] V. Valls, Code “quantum-observable-forecasting”, GitHub, 2026, <https://github.com/victor-ibm/quantum-observable-forecasting>.
- [46] D. P. Kingma and J. Ba, Adam: A method for stochastic optimization, *arXiv:1412.6980*.
- [47] J. Karthik, A. Sharma, and A. Lakshminarayanan, Entanglement, avoided crossings, and quantum chaos in an Ising model with a tilted magnetic field, *Phys. Rev. A* **75**, 022304 (2007).
- [48] Y. Liu, M. Li, J. Bao, B. Guo, and Z. Sun, Multipartite nonlocality in an Ising model with a tilted magnetic field at zero and finite temperatures, *Phys. Lett. A* **450**, 128396 (2022).
- [49] A. Dutta, G. Aeppli, B. K. Chakrabarti, U. Divakaran, T. F. Rosenbaum, and D. Sen, *Quantum Phase Transitions in*

- Transverse Field Spin Models: From Statistical Physics to Quantum Information* (Cambridge University Press, Cambridge, 2015).
- [50] M. S. Rudolph, T. Jones, Y. Teng, A. Angrisani, and Z. Holmes, Pauli propagation: A computational framework for simulating quantum systems, [arXiv:2505.21606](#).
- [51] T. Begušić and G. K.-L. Chan, Real-time operator evolution in two and three dimensions via sparse Pauli dynamics, [PRX Quantum](#) **6**, 020302 (2025).
- [52] C.-F. (Anthony) Chen, A. Lucas, and C. Yin, Speed limits and locality in many-body quantum dynamics, [Rep. Prog. Phys.](#) **86**, 116001 (2023).