Depth-Efficient Quantum Topological Data Analysis for Regime-Specific Detection of Financial Stress

Arul Rhik Mazumder
Carnegie Mellon University
Pittsburgh, PA, USA
ORCID: 0000-0002-2395-4400

Shreyan Ronit Mazumder
Cambridge Rindge and Latin School
Cambridge, MA, USA

***Abstract*—We present, to our knowledge, the first adaptation of Pauli Correlation Encoding (PCE) to quantum topological data analysis, reformulating Betti number estimation as a depth-efficient variational optimization over a compressed qubit register. From a Takens embedding and Vietoris–Rips filtration of S&P 500 returns, we extract combinatorial Laplacians and recast null-space counting as a continuous-PCE Rayleigh-quotient minimization with variational deflation, encoding n_k simplex indices into $O(n_k^{1/\kappa})$ qubits with shallow, ancilla-free circuits. Because the resulting loss is rational rather than bilinear in the correlators, the barren-plateau bound of [1] does not transfer; empirically the gradient variance decays only polynomially, with no exponential barren plateau, over $n = 4$–12 qubits. The classical stage matches `ripser` [2] on all 190 sliding windows (2007–2009). On the real market Laplacians ($\beta_1 = 1$–22), warm-starting from a classical null-space surrogate allows PCE-VQE to recover β_1 exactly at every scale, placing the obstacle in the optimisation landscape rather than the encoding. Chronologically split classification gives in-regime ROC AUC 0.818, but out-of-distribution evaluation on the 2020 COVID shock and 2022 rate cycle (AUC 0.009, 0.515) shows the calibration does not generalize across crisis regimes.**

***Index Terms*—quantum topological data analysis, Betti numbers, Pauli correlation encoding, variational quantum eigensolver, financial crash detection, combinatorial Laplacian**

I. Introduction

Financial markets exhibit complex nonlinear dynamics that defy classical statistical characterization. The 2008 financial crisis—which erased over $10 trillion in US household wealth—emerged from correlations and structural instabilities invisible to conventional volatility measures such as variance or VaR. During stable regimes a time series $\{x_t\}$ exhibits small, roughly Gaussian fluctuations; approaching a crash, the underlying dynamical attractor undergoes topological deformation that can precede the dislocation. In particular, Gidea and Katz [3] showed that persistence-based topological features of a multi-index point cloud (S&P 500, DJIA, NASDAQ, and Russell 2000 viewed jointly as coordinates in $\mathbb{R}^4$) exhibit a sustained rise in the spectral density of L^p-norms of persistence landscapes for roughly 250 trading days prior to the Lehman bankruptcy. This motivates the search for early warning signals that operate on the geometry of market data rather than its pointwise statistics, and in particular on its topology, which is stable under perturbations and captures qualitative structural change that purely local statistics miss [4]–[6]. Our pipeline is inspired by this line of work but adopts a different (single-index Takens delay) embedding, described in Section IV.

Topological Data Analysis (TDA) offers a robust framework for extracting such geometric features. The central invariants are Betti numbers β_k, which count the number of k-dimensional holes in the data manifold: β_0 counts connected components, β_1 counts loops, and β_2 counts voids. However, computing Betti numbers classically requires constructing and diagonalizing combinatorial Laplacians whose dimensions grow combinatorially with the number of simplices—an $O(n_k^3)$ cost per window that becomes prohibitive for real-time deployment at scale.

Quantum algorithms for TDA (qTDA) have been proposed to address this computational bottleneck, beginning with the seminal LGZ algorithm [7] and its subsequent refinements [8]–[13]. These approaches use quantum phase estimation (QPE) to estimate the eigenspectrum of the combinatorial Laplacian, from which the Betti number is recovered as the nullity. However, standard QPE requires deep circuits with many ancilla qubits, making it impractical for near-term and early fault tolerant hardware. Moreover, recent complexity-theoretic analyses [13]–[15] have shown that superpolynomial quantum advantage requires exponentially large Betti numbers and specific spectral gap conditions, raising the bar for practical near-term quantum TDA.

In parallel, Pauli Correlation Encoding (PCE) [1] has emerged as a qubit-efficient variational framework for combinatorial optimization, encoding $m = O(n^k)$ binary variables into n qubits via k-body Pauli correlators. PCE has demonstrated competitive performance on MaxCut [1], portfolio optimization [16], and the LABS problem [17], with provable super-polynomial barren plateau mitigation for losses of a specific bilinear form in Pauli expectations, and shallow circuit depths scaling sublinearly in the number of variables.

Scope

Because the gap between methodological novelty and end-to-end deployment is the central scientific risk of this work, we state the scope explicitly. **What is demonstrated:** (1) a correct classical TDA pipeline (CE benchmark) verified against `ripser` on 190 sliding windows; (2) a PCE-VQE adaptation that recovers Betti numbers on six toy Laplacians ($\beta_1 \leq 2$, 6/6) and on a $\beta_1 = 4$ four-cycle benchmark under tuned pa-

rameters; (3) a gradient variance that decays only polynomially (no exponential barren plateau) over the tested $n = 4$–12 qubit range; (4) a chronologically split out-of-sample evaluation of the β_1 signal as a crash predictor; (5) a resource model for the proposed near-term implementation; and (6) an end-to-end run of PCE-VQE on the real market Laplacians ($\beta_1 = 1$–22): from a random start it recovers no null vector, but a classical-warm-started hybrid recovers β_1 exactly (loss $\sim 10^{-13}$) at every scale, isolating the obstacle as the optimisation landscape rather than the encoding (Sec. VII-C). **What is not demonstrated:** *independent* (non-warm-started) PCE-VQE recovery of β_1 at real-data scale ($n_k \in [31, 429]$, $\beta_1 \in [0, 22]$) is not achieved; the recovery above uses the classical null space to warm-start and is a classical-quantum hybrid (Sec. VII-C); the cost-crossover at $n_k \gtrsim 10^4$ is an *extrapolation*, not a measured threshold; the noise-robustness comparison (Fig. 4) is not at matched encoded problem size; and the β_1 classifier does *not* generalize across distinct crisis regimes.

In this work, we bridge the qTDA and PCE directions by adapting the PCE framework to the spectral problem of Betti number estimation. Two adaptations distinguish our approach from canonical PCE: (i) we use the continuous expectation values $c_i = \langle \Pi_i \rangle \in [-1, +1]$ as eigenvector coefficients, rather than the binary $x_i = \mathrm{sgn}(\langle \Pi_i \rangle)$ outputs that PCE was designed to produce; and (ii) the resulting Rayleigh-quotient loss with variational deflation is rational in the correlators, so the formal trainability proof of [1] does not directly transfer. Our contributions are: (1) a complete qTDA pipeline blueprint from raw financial time series to Betti number estimation; (2) to our knowledge, the first application of PCE to eigenvalue counting — reformulating Betti number estimation as a variational optimization over $O(n_k^{1/\kappa})$ qubits with depth polynomial in $n_k^{1/\kappa}$, trading the logarithmic qubit count of LGZ for shallower, ancilla-free circuits; (3) a continuous-PCE deflation protocol with explicit discussion of why the formal barren-plateau guarantee of [1] does not transfer, and an empirical finite-size gradient-variance study; and (4) numerical experiments covering classical CE/`ripser` validation on real S&P 500 data, PCE-VQE validation on toy Laplacians, chronologically split classification with independent OOD evaluation, $\kappa = 2$ vs. $\kappa = 3$ resource comparison, and matched-scale noise simulation at $n_k \leq 16$. We treat the OOD generalization failure (Sec. VII-H) as a substantive finding rather than a footnote. For reproducibility, the full code, figures, and result artifacts used throughout this paper are publicly available at https://github.com/arulrhikm/Quantum-Market-Crash-TDA.

The remainder of this paper: Sec. II introduces background; Sec. III surveys related work; Sec. IV details the four-stage pipeline; Sec. V presents the classical benchmark; Sec. VI–VII cover experimental setup and results; Sec. IX concludes.

II. Background

A. Topological Data Analysis and Betti Numbers

TDA provides tools for extracting the shape of data from point clouds via simplicial complexes [4]–[6]. A k-simplex is a collection of $k+1$ vertices with all $\binom{k+1}{2}$ pairwise edges; a simplicial complex $\mathcal{K}$ is a collection of simplices closed under taking faces. Given a point cloud $\{z_i\}_{i=1}^n$ in $\mathbb{R}^m$ with distance function d, the Vietoris–Rips complex at scale ε includes a simplex $\sigma = \{z_{i_0}, \ldots, z_{i_k}\}$ whenever $\|z_{i_a} - z_{i_b}\| \leq \varepsilon$ for all pairs a, b. Varying ε yields a filtration of nested complexes whose persistent homology captures multiscale topological structure [6], [18].

The boundary operator $\partial_k : C_k \to C_{k-1}$ maps k-chains to $(k-1)$-chains. The k-th combinatorial (Hodge) Laplacian is defined as

$$\Delta_k \;=\; \partial_{k+1}\partial_{k+1}^\top + \partial_k^\top \partial_k \;\in\; \mathbb{R}^{n_k \times n_k}, \tag{1}$$

where $n_k = |\mathcal{C}_k|$ is the number of k-simplices. By the Hodge Decomposition Theorem, $H_k(\mathcal{K}) \cong \ker(\Delta_k)$, so

$$\beta_k = \dim \ker(\Delta_k), \tag{2}$$

i.e., β_k equals the number of zero eigenvalues of Δ_k [5], [18].

B. Quantum Phase Estimation for TDA

The LGZ algorithm [7] and its variants estimate Betti numbers by applying QPE to the unitary $U = e^{i\Delta_k}$. Since zero eigenvalues of Δ_k map to eigenvalue $e^{i\cdot 0} = 1$ of U, the probability of measuring phase $\theta = 0$ in the QPE register, when starting from a maximally mixed state, yields [19]:

$$\tilde{\beta}_k = 2^q \cdot p(0), \tag{3}$$

where $q = \lceil \log_2 n_k \rceil$ is the number of system qubits. The combinatorial Laplacian must be padded to dimension 2^q; following [19], we use identity padding with $\tilde{\lambda}_{\max}/2$ on the diagonal to avoid introducing spurious zero eigenvalues. Standard QPE requires $O(q)$ ancilla qubits and circuit depths scaling with the desired precision, which limits near-term applicability.

C. Time-Delay Embedding

Given a scalar time series $X = (x_0, \ldots, x_{L-1})$, Takens' embedding [20] forms delay vectors

$$\mathbf{z}_t = (x_t,\; x_{t+\tau},\; \ldots,\; x_{t+(m-1)\tau}) \in \mathbb{R}^m \tag{4}$$

for embedding dimension m and delay τ. By Takens' theorem, for generic smooth dynamical systems on a d-dimensional attractor, this map is a diffeomorphism when $m \geq 2d + 1$, preserving the topology of the attractor. The delay τ is selected as the first local minimum of the mutual information $I(x_t; x_{t+\tau})$, and m as the smallest dimension for which the false nearest neighbor fraction falls below 5% [21]; the concrete application of these criteria to our S&P 500 series is described in Section IV-A.

D. Pauli Correlation Encoding

Pauli Correlation Encoding (PCE) [1] is a qubit-efficient variational framework that encodes $m = O(n^k)$ classical binary variables into n qubits via k-body Pauli correlators. Given an n-qubit parameterized state $|\Psi(\vec{\theta})\rangle$, a set of m Pauli strings $\Pi^{(k)} = \{\Pi_1^{(k)}, \ldots, \Pi_m^{(k)}\}$ is chosen, where each

$\Pi_i^{(k)}$ is a permutation of $P_1^{\otimes k} \otimes \mathbb{1}^{\otimes(n-k)}$ for Pauli matrices $P_1 \in \{X, Y, Z\}$. The i-th binary variable is then decoded via

$$x_i = \operatorname{sgn}\Big(\langle\Psi(\vec{\theta})|\Pi_i^{(k)}|\Psi(\vec{\theta})\rangle\Big). \tag{5}$$

Our instantiation follows the standard, reproducible enumeration strategy used in introductory PCE implementations [1], [22]: enumerate κ-local positions and operator assignments from $\{X, Y, Z\}^\kappa$, place identities elsewhere, and truncate to the required count. More sophisticated, topology-aware assignments and commuting-group measurement reductions are compatible with PCE but are left to future work in this application. For $k = 2$ (quadratic PCE), this yields $m = O(n^2)$ variables from n qubits; for $k = 3$ (cubic PCE), $m = O(n^3)$. A hardware-efficient ansatz (HEA) with alternating single-qubit rotation and entangling layers is trained to minimize a loss function $\mathcal{L}(\vec{\theta})$ constructed from the Pauli expectations.

The key properties of PCE that motivate our adaptation are:

1) **Polynomial qubit compression:** m problem variables are encoded in $O(m^{1/k})$ qubits, yielding quadratic or cubic space savings.
2) **Built-in barren plateau mitigation:** Sciorilli et al. [1] prove that the variance of the gradient of cost functions of the form $\mathcal{L} = \sum_{ij} w_{ij}\langle\Pi_i\rangle\langle\Pi_j\rangle$ is lower-bounded by a super-polynomial function of n, avoiding the exponential gradient vanishing of generic HEA+global cost function combinations [23]. Note that this guarantee applies to the canonical bilinear PCE loss; the Rayleigh-quotient loss with deflation introduced in this paper is rational in the correlators and is not directly covered by the proof. We discuss this carefully in Section IV-D.
3) **Shallow circuits:** Circuit depth scales sublinearly in m, and the number of parameters scales linearly.
4) **Commuting measurement groups:** The Pauli set $\Pi^{(k)}$ decomposes into a small number of mutually commuting subsets, enabling efficient simultaneous measurement.

PCE experiments in literature.: The Sciorilli et al. Nature Communications study [1] reports MaxCut numerical simulations at $m = 7000$ and trapped-ion experiments on IonQ Aria-1 / Quantinuum H1-1 with $n = 17$ qubits and unweighted MaxCut sizes $m \in \{800, 2000\}$ (weighted $m = 512$). The follow-on PCE-LABS paper [17] operates in a different regime: simulations of LABS instances up to $N = 44$–45 with $n = 4$–6 qubits and ~ 30 two-qubit gates, with proof-of-principle experiments on IonQ Forte. The two studies are complementary, not directly comparable in scale; we attribute MaxCut-scale numbers to [1] and LABS-scale numbers to [17].

III. Related Work

A. Quantum Algorithms for TDA

The landscape of quantum TDA algorithms has evolved rapidly since the original LGZ proposal [7].

LGZ and refinements. The original algorithm uses QPE on $e^{i\Delta_k}$ prepared from a maximally mixed state to estimate β_k via (3). Ubaru et al. [8] improved the exponential speedup and depth complexity. McArdle et al. [9] proposed a streamlined algorithm achieving an almost quintic speedup over rigorous classical methods using exponentially fewer qubits, but also introduced a quantum-inspired classical power method with scaling only quadratically worse, casting doubt on exponential advantage for practical instances.

Persistent Betti numbers. Hayakawa [11] gave the first quantum algorithm for persistent Betti numbers of arbitrary dimensions, efficient for Vietoris–Rips complexes but limited to normalized estimates.

Complexity-theoretic limits. Crichigno and Kohler [14] proved that deciding whether a clique complex has a k-dimensional hole is QMA_1-hard and that the exact counting version is #BQP-hard. Schmidhuber and Lloyd [15] separately showed that approximating Betti numbers to multiplicative error is NP-hard for general complexes. Berry et al. [13] showed that superquadratic quantum speedups require multiplicative error approximation with asymptotically growing Betti numbers, and proposed a dequantization showing that exponentially large dimension and Betti number are necessary but insufficient for superpolynomial advantage. Their fault-tolerant resource estimates require tens of billions of Toffoli gates.

Near-term approaches. Scali et al. [24] proposed a thermal-state approach ("thermal QTDA") that reduces Betti number estimation to a purity test on a low-temperature Gibbs state of Δ_k, using $2\lceil\log_2 n_k\rceil$ system qubits plus a SWAP-test ancilla and inheriting an inverse spectral-gap dependence in the state-preparation time. Relative to our PCE approach it uses exponentially fewer qubits ($O(\log n_k)$ vs. $O(n_k^{1/\kappa})$) but needs ancilla overhead and Lindbladian machinery; PCE offers shallower, ancilla-free circuits at the cost of a larger register and no formal trainability guarantee. Akhalwaya et al. [12] compared quantum and classical Monte Carlo estimators for clique-complex Betti numbers.

Hybrid approaches. Nghiem, Lee, and Wei [25] proposed a hybrid quantum-classical algorithm combining classical simplex enumeration with quantum Betti number estimation, using $(|S_r| + |S_{r+1}|) \cdot O(r)$ total qubits—scaling linearly in simplex count, significantly more than LGZ or our approach. Their framework provides provable complexity guarantees for normalized Betti number estimation, whereas our PCE pipeline is a trainable heuristic targeting absolute Betti numbers without such guarantees. Liu [26] proposed bridging from Betti numbers to persistence diagrams via quantum kernel methods.

B. Pauli Correlation Encoding

Beyond the foundational MaxCut work [1], PCE has been extended to portfolio optimization by Soloviev and Krompiec [16], the LABS problem by Sciorilli et al. [17], and budget-constrained optimization by Padín-Martínez et al. [27]; IBM Quantum has published a tutorial implementation [22]. All prior PCE applications target combinatorial optimization, finding a single optimal bitstring via the binary decoding rule in (5). Our work is the first, to our knowledge, to apply PCE to an eigenvalue counting problem, where the goal is to determine

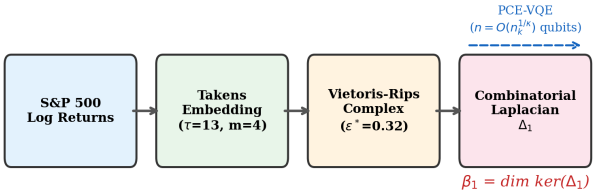

Fig. 1. Overview of the qTDA pipeline with PCE qubit compression: (1) S&P 500 log returns → Takens embedding ($\tau = 13$, $m = 4$) → point cloud in $\mathbb{R}^4$; (2) Vietoris–Rips complex $\mathrm{VR}(\varepsilon^*)$, 461 points per window; (3) combinatorial Laplacian $\Delta_1 = \partial_2\partial_2^\top + \partial_1^\top\partial_1$, $n_k \in [31, 429]$; (4) PCE encoding into $n = O(n_k^{1/\kappa})$ qubits, HEA + COBYLA optimisation with variational deflation → $\tilde{\beta}_1$.

the multiplicity of the zero eigenvalue of a matrix. To do so we depart from canonical PCE in two respects: (i) we use the continuous expectation values $c_i = \langle \Pi_i \rangle \in [-1, +1]$ as eigenvector coefficients without applying the sgn rounding, and (ii) we introduce a Rayleigh-quotient loss with variational deflation. Both adaptations move the cost function outside the bilinear-in-correlators class covered by the trainability proof of [1].

C. Classical TDA for Financial Applications

TDA has been applied to financial market analysis by several groups [3], [6]. Gidea and Katz [3] demonstrated that persistent homology features of a multi-index point cloud constructed from the daily log returns of four major US stock indices (S&P 500, DJIA, NASDAQ, and Russell 2000), each used as a coordinate in $\mathbb{R}^4$, exhibit a sustained rise in the spectral density at low frequencies of L^p-norms of persistence landscapes for roughly 250 trading days prior to the Lehman bankruptcy. Subsequent work has shown that persistence-based features provide complementary information to standard volatility measures for regime classification and crash forecasting. Our pipeline differs from [3] in two respects: we use a single-index S&P 500 Takens delay embedding rather than a multi-index direct embedding, and we replace the classical persistence-based summary by a quantum-accelerated Betti number computation.

IV. Methodology

Our pipeline converts raw financial time series into quantum-estimated Betti numbers through four stages (Fig. 1).

A. Step 1: Time-Delay Embedding

We apply Takens' embedding (4) to daily S&P 500 log returns $r_t = \log(P_t/P_{t-1})$, determining τ and m via mutual information and false nearest neighbors (FNN) directly on the return series.

Selecting τ.: The delay τ is the first local minimum of the mutual information $I(r_t; r_{t+\tau}) = \sum_{x,y} p(x,y)\log[p(x,y)/(p(x)p(y))]$, evaluated at lags $\tau = 1, \ldots, 30$ via a 16×16 histogram of the pairs $(r_t, r_{t+\tau})$. Mutual information is high at small τ (successive coordinates are redundant) and decreases as τ grows; its first local minimum identifies the lag at which $r_{t+\tau}$ is maximally informative without becoming noise. For our return series this minimum occurs at $\tau = 13$ trading days (≈2.5 weeks), reflecting the characteristic decorrelation timescale of daily equity returns.

Selecting m.: With $\tau = 13$ fixed, m is the smallest dimension for which fewer than 5% of nearest neighbors are false. For each embedded point $\mathbf{z}_t$, its nearest neighbor $\mathbf{z}_{t'}$ in the m-dimensional embedding is declared false if the ratio $R_t = |x_{t+m\tau} - x_{t'+m\tau}|/\|\mathbf{z}_t - \mathbf{z}_{t'}\|_m$ exceeds a threshold (we use $R_{\mathrm{tol}} = 15$, within the conventional 10–15 range [28]); $x_{t+m\tau}$ is the candidate $(m{+}1)$-th coordinate. The FNN fraction remains above 5% at $m = 1, 2, 3$ and drops below 5% at $m = 4$, giving the selected embedding dimension. Although Takens' bound formally requires $m \geq 2d + 1$, this bound is conservative and the attractor unfolds sufficiently for downstream TDA at $m = 4$ in practice [21].

Resulting point cloud.: With $\tau = 13$ and $m = 4$, each sliding window of $W = 500$ trading days yields $n_{\mathrm{pts}} = W - (m-1)\tau = 461$ embedded points in $\mathbb{R}^4$.

B. Step 2: Vietoris–Rips Complex Construction

For each windowed point cloud, we construct the Vietoris–Rips complex $\mathrm{VR}(\varepsilon)$ at a range of thresholds ε. Varying ε produces a filtration of nested complexes; tracking how topological features appear and disappear across this filtration yields persistence diagrams—multiscale summaries robust to noise and perturbations [5], [18]. Applied to the sliding window, this procedure produces a time series of Betti curves whose changes signal regime transitions in the underlying financial dynamics [3].

Choosing the working threshold ε^*.: For classical validation with `ripser` we use the full distance-induced Vietoris–Rips filtration. `ripser` returns persistence intervals $\{(b_i, d_i)\}$ for each homological dimension k; we obtain a Betti number at scale ε by counting intervals with $b_i \leq \varepsilon < d_i$ (the standard "slice-of-the-barcode" definition). For the CE benchmark and quantum experiments—which operate at a single scale—we select a working threshold ε^* by scanning a grid constructed from percentiles of the interpoint distance distribution and choosing the ε that maximizes a robust central tendency (median) of β_1 across windows, with a stability tie-break toward smaller ε. This yields $\varepsilon^* = 0.32$ for the 2007–2010 dataset.

C. Step 3: PCE Encoding of the Spectral Problem

This section describes our central methodological contribution: the adaptation of Pauli Correlation Encoding to the problem of estimating $\beta_k = \dim\ker(\Delta_k)$.

1) From eigenvalue counting to variational optimization: Given the combinatorial Laplacian $\Delta_k \in \mathbb{R}^{n_k \times n_k}$, we seek the number of zero eigenvalues. We reformulate this as a sequence of variational feasibility problems. In standard VQE, one would encode Δ_k as a Hamiltonian on $\lceil \log_2 n_k \rceil$ qubits and minimize the expectation value directly. Here, we instead use the PCE correlators to parameterize a trial vector in $\mathbb{R}^{n_k}$

and evaluate the Rayleigh quotient classically (Eq. (8) below). For conceptual motivation, the standard VQE objective for the first eigenvalue is

$$\mathcal{L}_1(\vec{\theta}) = \langle \Psi(\vec{\theta}) | \, \Delta_k \, | \Psi(\vec{\theta}) \rangle \,, \tag{6}$$

where $|\Psi(\vec{\theta})\rangle$ lives in the n_k-dimensional simplex space; the PCE adaptation replaces this with a classical evaluation over the compressed qubit register.

2) Continuous-PCE variable encoding: The Laplacian Δ_k acts on the vector space $\mathbb{R}^{n_k}$ indexed by k-simplices. Rather than encoding this space directly into $\lceil \log_2 n_k \rceil$ qubits (as in standard LGZ), we use PCE to encode n_k basis coefficients into $n = O(n_k^{1/\kappa})$ qubits for a chosen compression order κ.

Concretely, we select a set of n_k Pauli strings $\Pi^{(\kappa)} = \{\Pi_1, \ldots, \Pi_{n_k}\}$ as described in Section II-D. We depart from the canonical PCE framework [1] by treating each Pauli expectation as a continuous eigenvector coefficient,

$$c_i(\vec{\theta}) = \langle \Psi(\vec{\theta}) | \Pi_i | \Psi(\vec{\theta}) \rangle \in [-1, +1], \quad i = 1, \ldots, n_k, \tag{7}$$

rather than rounding to $\mathrm{sgn}(\langle \Pi_i \rangle)$ as is done in the canonical optimization setting. The reachable set $\{\mathbf{c}(\vec{\theta})\}$ is therefore a strict subset of $[-1, +1]^{n_k}$, and the Rayleigh quotient evaluated below operates on this restricted ansatz rather than on all of $\mathbb{R}^{n_k}$. We discuss the resulting expressivity question explicitly in Section VIII.

In our experiments, we instantiate $\Pi^{(\kappa)}$ via a simple deterministic enumeration consistent with the PCE framework [1], [22]: for each κ-tuple of distinct qubit indices and each operator assignment in $\{X, Y, Z\}^\kappa$, we place the chosen Paulis on those indices and identity elsewhere (reversing bit order to match simulator conventions), appending strings until $|\Pi^{(\kappa)}| = n_k$. We choose n so that $\binom{n}{\kappa} \cdot 3^\kappa \geq n_k$. This provides a reproducible mapping from simplex indices to Pauli strings. We did not optimize this assignment for the Laplacian's topology, nor apply commuting-group measurement reductions in our implementation; these are promising directions for future work. The variational loss function becomes

$$\mathcal{L}(\vec{\theta}) = \frac{\mathbf{c}(\vec{\theta})^\top \Delta_k \, \mathbf{c}(\vec{\theta})}{\|\mathbf{c}(\vec{\theta})\|^2} = \frac{\sum_{i,j} [\Delta_k]_{ij} \, c_i \, c_j}{\sum_i c_i^2}, \tag{8}$$

where $\mathbf{c}(\vec{\theta}) = (c_1, \ldots, c_{n_k})^\top$. This is a Rayleigh quotient over the restricted set of vectors parameterized by PCE correlators. Minimizing (8) yields the smallest eigenvalue of Δ_k accessible from within the reachable set; if this is below threshold δ, an approximate null vector has been found.

3) Measurement protocol: The loss (8) requires estimating second-order correlators $c_i \, c_j = \langle \Pi_i \rangle \langle \Pi_j \rangle$. Since the Pauli set $\Pi^{(\kappa)}$ decomposes into a small number G of mutually commuting groups (typically $G = 3$ for $\kappa = 2$ [1]), all n_k expectations can be estimated from G measurement bases per optimization step. The total per-step measurement overhead is $O(G/\epsilon_{\text{shot}}^2)$ shots, independent of n_k.

We note that the Vietoris–Rips Laplacian may admit additional commutation structure from the simplicial complex topology, potentially allowing further grouping reduction beyond generic PCE; we leave the derivation of tighter bounds on G for structured Laplacians to future work.

4) Evaluating the Rayleigh quotient: Evaluating (8) requires computing $\sum_{ij} [\Delta_k]_{ij} \, c_i \, c_j$ from the measured expectations $\{c_i\}$. Since Δ_k is sparse (each row has at most $O(n_{\text{pts}})$ nonzero entries for the VR complex, though empirically far fewer), this sum is evaluated classically in the optimization loop. Across all 190 windows the mean number of nonzero entries per row of Δ_1 is 3.48 (window-averaged maximum 8.8), confirming that the classical post-processing cost per optimization step is negligible compared to the quantum measurement cost.

D. Step 4: Variational Deflation for Betti Number Counting

Estimating β_k requires counting the multiplicity of the zero eigenvalue, not just finding a single ground state. We employ variational deflation [29]: after finding the j-th approximate null vector $\mathbf{c}^{(j)}$, we add a penalty term to the loss for subsequent searches:

$$\mathcal{L}_{j+1}(\vec{\theta}) = \mathcal{L}(\vec{\theta}) + \mu \sum_{\ell=1}^{j} \left| \frac{\mathbf{c}(\vec{\theta})^\top \mathbf{c}^{(\ell)}}{\|\mathbf{c}(\vec{\theta})\| \|\mathbf{c}^{(\ell)}\|} \right|^2, \tag{9}$$

where $\mu > 0$ is a penalty strength (default $\mu = 5.0$ unless stated otherwise). This encourages the optimizer to find states orthogonal to previously identified null vectors. The Betti number is estimated as the number of successful null-space searches before the minimum of $\mathcal{L}_{j+1}$ exceeds the threshold δ (default $\delta = 0.01$):

$$\tilde{\beta}_k = \max\{j : \mathcal{L}_j^* < \delta\}. \tag{10}$$

This deflation scheme carries an intrinsic cost. Each null vector is recovered by its own optimization, so counting the β_k zero eigenvalues requires β_k sequential PCE-VQE runs, and the total circuit-optimization budget grows linearly with the Betti number being estimated. This is a real weakness in exactly the high-β_k regime that motivates quantum TDA [13], and in practice it limits the multiplicity the scheme can reach. Subspace-search variants that recover several null vectors in a single optimization (Sec. IX) are the natural route to removing this linear factor.

1) Why the Sciorilli barren-plateau guarantee does not directly transfer: The super-polynomial barren plateau lower bound proved in [1] applies specifically to loss functions of the bilinear form $\mathcal{L}_{\text{PCE}} = \sum_{ij} w_{ij} \langle \Pi_i \rangle \langle \Pi_j \rangle$ with bounded coefficient weights. Our loss (8) has two structural differences that move it outside the class covered by that proof:

1) the Rayleigh-quotient denominator $\sum_i c_i^2$ introduces a normalization that makes the loss rational, rather than bilinear, in the correlators;
2) the deflation penalty in (9) contributes additional ratio terms whose Hessian structure is not captured by the bilinear analysis.

We therefore make no formal trainability claim. Instead, we report empirical gradient-variance behavior over $n = 4$–12

qubits and $j = 0, 1, 2$ deflation rounds in Section VII-G. Across the tested range the gradient variance decays only polynomially with system size (log–log slope ≈ -1), far from the exponential vanishing of a barren plateau and consistent with the trainability advantages of PCE-style encodings, but this should not be read as a verification of the asymptotic Sciorilli bound.

V. Classical Exact-Eigensolver Benchmark

To contextualise the quantum pipeline, we introduce a classical exact-eigensolver (CE) benchmark that computes Betti numbers by direct dense diagonalisation of Δ_k. The CE benchmark serves two roles: it provides exact ground-truth Betti curves free of the approximation errors of QPE/VQE, and it establishes the cost regime at which quantum methods become advantageous.

A. Algorithm and Threshold

Given $\Delta_k \in \mathbb{R}^{n_k \times n_k}$, the CE benchmark computes the full eigendecomposition $\Delta_k = Q\Lambda Q^\top$ via `numpy.linalg.eigh` (LAPACK `dsyevd`). The Betti number is recovered as the nullity

$$\beta_k^{(\mathrm{CE})} = |\{i : \lambda_i < \delta\}|, \quad \delta = c_\delta \cdot \varepsilon_{\mathrm{mach}} \cdot n_k \cdot \|\Delta_k\|_2, \quad (11)$$

with $c_\delta = 10$ and $\varepsilon_{\mathrm{mach}} \approx 2.2 \times 10^{-16}$. Across all 190 sliding windows we observe a spectral gap of at least three orders of magnitude between true-zero and nonzero eigenvalues; the empirical gap distribution is recorded in `results/verification_results.json`.

B. Complexity and Implementation

The dominant cost is dense eigendecomposition, $O(n_k^3)$ time / $O(n_k^2)$ memory. For our pipeline ($W = 500$, $m = 4$) the typical Laplacian dimension grows as $n_1 = \binom{n_{\mathrm{pts}}}{2}\rho(\varepsilon)$. At ε^* we find $n_1 \in [31, 429]$ (mean 136, median 116), giving eigendecomposition times of $\approx$ 41–228 ms per window (mean $\approx$ 89 ms) on a single CPU; total runtime for all 190 windows is $\approx$ 25.5 s. Across the 190 windows the mean number of nonzero entries per row of Δ_1 is 3.48 (window-averaged maximum 8.8), confirming that classical post-processing of the Rayleigh quotient (8) is negligible relative to the quantum measurement cost. Mean number of 2-simplices per window is 60 (max 476); edge sparsity relative to the complete graph on 461 points is 0.13%, reflecting the conservative ε^* choice and the principal reason β_1 stays moderate (mean 4.72).

C. Validation Against `ripser`

We verify that the CE benchmark agrees with `ripser` [2] on all 190 sliding windows in the 2007–2009 dataset. `ripser` returns persistence intervals for H_0 and H_1; we extract Betti numbers at scale ε^* by counting intervals (b_i, d_i) with $b_i \leq \varepsilon^* < d_i$, and compare against the nullity computed by (11). Table I reports the comparison for β_0 and β_1 at the selected filtration scale $\varepsilon^* = 0.32$. We emphasize that this validation concerns only the classical Betti extractor used as ground-truth in this paper; the PCE-VQE quantum pipeline is validated separately in Section VII-B on toy Laplacians.

TABLE I
CE vs. `ripser` Agreement (2007–2009, $W = 500$, $\varepsilon^* = 0.32$)

Window period	β_0^{rip}	β_0^{CE}	β_1^{rip}	β_1^{CE}
Jan 2007 – Dec 2008	284	284	15	15
Mar 2007 – Feb 2009	296	296	9	9
Jun 2007 – May 2009	327	327	6	6
Sep 2007 – Aug 2009	335	335	3	3
Dec 2007 – Nov 2009	338	338	3	3
All 190 windows	**100%**	—	**100%**	—

Both β_0 and β_1 agreement with `ripser`: 190/190 (100%). $\beta_1 \in [0, 22]$ across all windows (mean 4.72). *This table reports classical-vs-classical agreement; PCE-VQE results are reported in Section VII-B.*

VI. Experimental Setup

A. Dataset

We use daily S&P 500 closing prices downloaded from Yahoo Finance (ticker: `^GSPC`) covering January 2, 2003 to December 30, 2010, encompassing the 2007–2009 financial crisis. The dataset contains 2,014 trading days. Log returns $r_t = \log(P_t/P_{t-1})$ yield 2,013 return observations. After Takens embedding with $\tau = 13$ and $m = 4$, the embedded point cloud contains 1,974 vectors in $\mathbb{R}^4$. Sliding windows of $W = 500$ trading days are extracted from the raw return series with step size 8 days, producing 190 windows; each window of 500 raw days yields $500-(m-1)\tau = 461$ embedded points in $\mathbb{R}^4$. Points are standardised per-dimension before distance computation. For out-of-distribution evaluation (Section VII-H) we additionally use S&P 500 daily data from June 2019–December 2020 (COVID shock) and January 2022–June 2023 (rate-cycle drawdown); the Takens pipeline is applied identically. The 2003–2010 dataset is archived in this repository (`data/sp500_2003_2010.csv`); OOD episodes are downloaded on-demand via `yfinance` with fixed date ranges and the resulting metrics are saved to `results/classification_ood.json` for reproducibility.

B. Classical Baselines

Two classical baselines are employed:

1) `ripser` [2]: exact persistence diagrams via boundary-matrix reduction (primary ground truth).
2) CE benchmark (Section V): exact Betti numbers via dense diagonalisation (secondary baseline and resource reference).

C. Quantum Simulation

1) Small-scale exact simulation: The PCE-variational circuits are implemented in Qiskit and simulated on BlueQubit statevector backends for $n = 4$–12 qubits, encoding Laplacians of dimension $n_k = 16$–259 (quadratic PCE, $\kappa = 2$). These simulations validate the Betti number estimation protocol on synthetic Laplacians; the full real-data range $n_k \in [31, 429]$ has not been run end-to-end through the quantum pipeline,

and we make no quantitative claim about PCE-VQE accuracy at those scales. All statevector-only experiments (convergence, encoding-order comparison, barren-plateau sweep) use BlueQubit; noise experiments use Qiskit Aer density matrices by construction.

2) Resource extrapolation: For larger scales ($n_k = 10^3$–10^5), we report *analytical* resource estimates: qubit count, circuit depth, gate count, and measurement shots, extrapolated from the small-scale empirical data and the known scaling properties of PCE [1]. These are extrapolations, not measurements, and should be read as research targets rather than achieved capabilities.

3) Noise simulation: We simulate depolarizing noise at error rates $p \in \{0, 10^{-4}, 5\times10^{-4}, 10^{-3}, 5\times10^{-3}, 10^{-2}\}$ per gate using Qiskit Aer's density-matrix backend to assess noise robustness. For LGZ-QPE, we report both matched-scale simulated accuracy at $n_k \in \{8, 16\}$ and analytical extrapolation for larger n_k. The matched-scale comparison and its limitations are discussed in detail in Section VII-F.

D. Hardware-Efficient Ansatz

We employ a hardware-efficient ansatz [30] with alternating layers of single-qubit rotations (R_Y, R_Z) and entangling CNOT gates in a linear connectivity topology. The number of layers is set to $L = \lceil 2n \rceil$, scaling linearly in the qubit count n rather than in $m = n_k$ (in contrast, [1] reports brickwork depths sublinear in m; the two parameterizations are related by $n = O(m^{1/\kappa})$, so $L = O(n) = O(m^{1/\kappa})$ here is consistent with the $O(m^{1/\kappa})$ depth class but slightly more conservative). With $O(n)$ two-qubit depth per layer, this yields total circuit depth $O(n^2)$, which we use as the proxy $d_{\text{proxy}} = 2n^2$ in resource comparisons. The classical optimizer is COBYLA, which typically converges in 80–130 iterations on the toy and small-scale benchmarks reported here; the hardware resource estimate in Section VII-I uses 200 iterations as a conservative upper bound.

E. Resource Metrics

We track the following quantities as functions of the Laplacian dimension n_k and PCE order κ:

- Number of qubits: $n = O(n_k^{1/\kappa})$
- Number of variational parameters: $N_{\text{param}} = O(n \cdot L)$
- Two-qubit gate count per circuit evaluation
- Circuit depth
- Number of measurement bases G
- Measurement shots per optimization step
- Number of deflation rounds (equals β_k)
- Total optimizer iterations across all deflation rounds

These are compared between PCE-variational (this work), standard LGZ-QPE, the streamlined algorithm of [9], and the thermal QTDA approach of [24].

VII. Results

A. Betti Curve Comparison (Classical Pipeline)

The first-Betti-number time series $\beta_1(t)$ across all 190 sliding windows, computed by the CE benchmark, is shown in

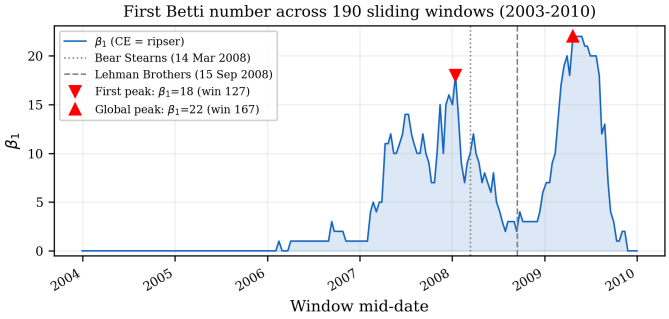

Fig. 2. First Betti number β_1 over 190 sliding windows (2003–2010), computed by the classical CE benchmark and verified against `ripser`. Vertical lines mark Bear Stearns (14 Mar. 2008, dotted) and Lehman Brothers (15 Sep. 2008, dashed). $\beta_1 \in [0, 22]$; mean 4.72.

Fig. 2. CE agrees with `ripser` on 190/190 windows (100%) for both β_0 and β_1, confirming that the classical stage of the pipeline is correct. The β_1 signal rises from zero in late 2006, reaches a first peak at window 127 ($\beta_1 = 18$, January 2007 to January 2009), dips briefly, and then attains the global maximum $\beta_1 = 22$ at window 167 (April 2008 to April 2010). The topological complexity begins building more than a year before the Bear Stearns collapse (March 2008) and Lehman Brothers bankruptcy (September 2008), consistent with the hypothesis that topological loop formation provides a leading indicator of systemic stress [3]. We emphasize that the curves in this section are produced by the classical CE benchmark; PCE-VQE has not been run end-to-end on these windows and we make no claim about quantum-pipeline accuracy at these n_k and β_1 values.

B. PCE Convergence Validation (Toy Laplacians)

We validate the PCE-VQE pipeline on six synthetic Laplacians with analytically known Betti numbers: path graph (β_1=0), hollow triangle (β_1=1), filled triangle (β_1=0), two disjoint hollow triangles (β_1=2), square/4-cycle (β_1=1), and figure-eight (β_1=2). All six are recovered correctly with κ=2, δ=0.01, μ=5.0 (6/6 correct). Loss converges to $\approx 10^{-10}$ for null vectors and jumps to $\approx$3.0 when no further null vector exists (Fig. 3).

The gradient-free deflation protocol correctly recovers $\beta_1 \in \{1, 2, 3\}$ on synthetic Laplacians (3/3), but at $\beta_1 = 4$ recovery becomes parameter-sensitive. This is the same random-start landscape wall we isolate in Sec. VII-C: from a random start even an exact-gradient deflation finds no null vector at $\beta_1 = 4$, yet with a classical-null-space warm start the full $\beta_1 = 4$ and $\beta_1 = 6$ null spaces are recovered to machine precision (`results/kappa_deflation_v2.json`). The sensitivity is therefore a symptom of the hardware-efficient-ansatz landscape, not of penalty tuning, and is removed by the warm-started optimizer.

Toy-to-real gap.: The validated PCE-VQE regime ($n_k \leq 64$, $\beta_1 \leq 4$) is roughly an order of magnitude smaller in both dimension and null-space multiplicity than the real S&P 500 regime ($n_k \in [31, 429]$; $\beta_1 \leq 22$). We probe this gap

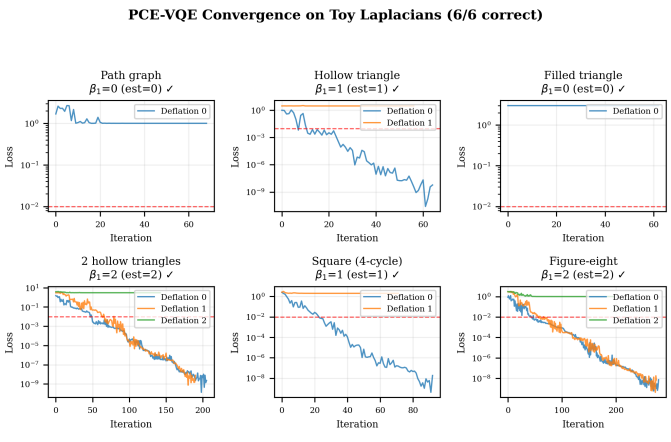

Fig. 3. PCE-VQE convergence on six toy Laplacians (6/6 correct). Each panel is one toy Laplacian, labelled with its true β_1 and the recovered estimate. Within a panel, each coloured curve is one deflation round (round 0 blue, round 1 orange, round 2 green), that is, one search for a null vector orthogonal to those already found; the red dashed line is the null-space threshold $\delta = 0.01$. A curve settling near 10^{-10} has located a null vector, while a round whose best loss jumps to $\approx$3.0 has none left to find and terminates the count. Each panel carries its own legend listing the deflation rounds it contains.

TABLE II
PCE-VQE ON THE REAL S&P 500 WINDOW LAPLACIANS (SMALLEST-n_k WINDOW PER β_1; $\kappa = 2$, $\delta = 10^{-2}$). $\min \mathcal{L}_0$ IS THE MINIMUM ROUND-0 LOSS FROM A RANDOM START (HEA $L = 2n$, COBYLA), WHICH NEVER REACHES δ, SO $\hat{\beta}_1 = 0$. $\min \mathcal{L}_{\Delta_1}$ IS THE MINIMUM LOSS AFTER WARM-STARTING FROM THE CLASSICAL NULL-SPACE SURROGATE (EXACT-GRADIENT L-BFGS, DEPTH $6n$; ONE EXTRA QUBIT FOR WINDOW 67), WHICH REACHES MACHINE ZERO AND RECOVERS β_1.

Window	β_1	n_k	n	$\min \mathcal{L}_0$	$\min \mathcal{L}_{\Delta_1}$
67	1	76	5	1.29	1.4×10^{-13}
185	2	105	6	1.27	1.5×10^{-12}
85	3	99	6	1.21	9.2×10^{-13}
98	4	160	7	1.74	4.5×10^{-13}
157	6	167	7	1.47	4.8×10^{-13}
117	10	176	7	1.80	4.9×10^{-13}
170	22	259	9	1.48	4.8×10^{-13}

directly in Sec. VII-C, running PCE-VQE on the real market Laplacians rather than on toy proxies of them.

C. Real-Window Validation: PCE-VQE on Market Laplacians

The toy Laplacians of Sec. VII-B sit an order of magnitude below the real-data regime. Since the 190 windows already span every first Betti number from 0 to 22, we test transfer directly: for each β_1 we take the smallest-n_k real window with that value (Table II), build Δ_1 with the `ripser`-verified construction of Sec. V, and run PCE-VQE with the toy-validated configuration ($\kappa = 2$, HEA depth $L = 2n$, COBYLA, $\delta = 10^{-2}$). We report round 0, the search for the first null vector; if its minimum loss stays above δ then $\hat{\beta}_1 = 0$ regardless of later deflation. All seven windows use $n \leq 9$ qubits, so each runs in under two minutes on a CPU.

Random-start failure is a landscape problem, not an encoding one.: PCE-VQE recovers no null vector on any window: the minimum round-0 loss stays in $[1.21, 1.80]$ across $\beta_1 = 1$ to 22 (Table II), so $\hat{\beta}_1 = 0$ everywhere, breaking already at $\beta_1 = 1$. The encoding is not the cause. Optimising over a fully expressive n-qubit state (all 2^n amplitudes free) drives the correlator vector onto the null space, with overlap 1.00 and Rayleigh quotient 0.00 at $\beta_1 = 2$ (window 185) and overlap 0.99 at $\beta_1 = 1$ (window 67, exactly representable with one extra qubit). The null vectors lie inside the reachable set; what fails from a random start is the optimisation. The tiny-gap Δ_1 Rayleigh landscape traps COBYLA near $\mathcal{L} \approx 1.3$, and even exact-gradient Adam and L-BFGS plateau near $\mathcal{L} \approx 0.02$–0.05.

Recovery by warm-starting.: This points to a fix. Given a classical approximate null basis N of Δ_1, the surrogate $B = I - NN^\top$ shares its null space but has spectral gap 1, so its Rayleigh quotient is benign to optimise. We minimise the PCE Rayleigh quotient of B with exact adjoint gradients and L-BFGS, then polish on the true Δ_1, using a depth-$6n$ ansatz and one extra qubit where the minimal register cannot represent the null vector exactly. This recovers a null vector on every window, driving the Δ_1 loss to $\sim 10^{-13}$ across $\beta_1 = 1$ to 22 (Table II, last column). The warm start uses the classical null space, so this is a classical-quantum hybrid, not an independent quantum determination of β_1: it shows the null vectors are representable and reachable once the optimiser reaches the right basin, and that the random-start failure is one of landscape, not of encoding or hardware cost. A quantum-independent pipeline would need a quantum route to that warm start, such as a problem-informed or adaptive ansatz (Sec. IX).

D. Encoding-Order Comparison ($\kappa = 2$ vs. $\kappa = 3$)

On a benchmark Laplacian ($n_k = 64$, $\beta_1 = 2$; two disjoint 32-cycles), cubic PCE ($\kappa = 3$, $n = 5$) against quadratic ($\kappa = 2$, $n = 8$) cuts the qubit count by 37.5% and the depth proxy $d_{\text{proxy}} = 2n^2$ by 60.9% (128 $\rightarrow$ 50), with both satisfying $|\Pi^{(\kappa)}| \geq n_k$. With the exact-gradient, warm-started optimizer of Sec. VII-C, both encodings now converge, recovering $\beta_1 = 2$ to machine precision ($\mathcal{L} \approx 5 \times 10^{-15}$ for $\kappa = 2$, 2×10^{-11} for $\kappa = 3$), whereas the earlier gradient-free budgets left both short of δ. So the cubic encoding attains the same recovery at strictly lower qubit count and depth.

E. Resource Scaling Analysis

Table III presents analytical resource estimates comparing PCE (κ=2) with LGZ-QPE across a range of Laplacian dimensions.

Table III also contextualises the depth–qubit trade-off against other quantum TDA methods. LGZ-QPE and the streamlined algorithm of [9] use only $O(\log n_k)$ qubits but require deep circuits ($O(N_P \cdot r \cdot 2^p)$ for standard QPE). Thermal QTDA [24] matches the logarithmic qubit scaling with $2\lceil \log_2 n_k \rceil$ system qubits and a single ancilla, with depth governed by the Lindbladian mixing time. PCE uses more qubits ($O(n_k^{1/\kappa})$) but eliminates ancilla overhead entirely and operates at circuit depths polynomial in $n_k^{1/\kappa}$—the key advantage for near-term hardware where depth, not qubit count, is the binding constraint.

TABLE III
RESOURCE ESTIMATES: PCE (κ=2) VS. LGZ-QPE

n_k	PCE qubits	PCE depth	LGZ total qubits	LGZ depth
64	8	128	13	$> 10^4$
256	16	512	15	$> 10^4$
1 024	32	2 048	17	$> 10^5$
4 096	64	8 192	19	$> 10^5$
16 384	128	32 768	21	$> 10^6$
65 536	256	131 072	23	$> 10^6$

Analytical estimates. PCE depth $= 2n^2$; LGZ depth $= O(N_P \cdot r \cdot 2^p)$ with $p = \lceil \log_2(1/\epsilon) \rceil$ ancilla qubits (not shown). PCE uses more qubits but far shallower circuits and no ancilla overhead.

F. Noise Robustness

We compare PCE-VQE against LGZ-QPE under depolarizing noise at error rates $p \in \{0, 10^{-4}, 5{\times}10^{-4}, 10^{-3}, 5{\times}10^{-3}, 10^{-2}\}$. PCE-VQE: density-matrix simulation at $n = 6$ qubits on 4 test cases ($\beta_1 \in \{0, 1, 2, 3\}$, 5 trials each). LGZ-QPE: matched-scale simulation at $n_k \in \{8, 16\}$ plus an analytical extrapolation based on worst-case depolarizing accumulation over the QPE depth.

What Fig. 4 shows: At matched $n_k = 16$, the LGZ-QPE simulated accuracy drops from 1.00 at $p = 0$ to 0.316 at $p = 10^{-4}$ and below 10^{-2} at $p = 10^{-3}$, while PCE-VQE maintains accuracy above 0.90 through $p = 5{\times}10^{-3}$. This comparison is not at matched encoded problem size: PCE is at fixed $n = 6$, encoding $n_k \approx 63$ simplex dimensions under $\kappa = 2$, while LGZ is at fixed $n_k = 16$. The figure therefore reflects the depth difference between the two encodings at moderate scale, not a like-for-like accuracy-vs-noise comparison. We claim that the shallower PCE-VQE circuits accumulate substantially less depolarizing error per shot than depth-heavy LGZ-QPE; but not that "three-orders-of-magnitude" advantage in tolerable per-gate error at matched task accuracy.

G. Barren Plateau Analysis

We probe trainability through the gradient variance $\mathrm{Var}[\partial\mathcal{L}/\partial\theta_\ell]$ of the true PCE Rayleigh-quotient loss, evaluated exactly through the hardware-efficient ansatz with the adjoint gradient of Sec. VII-C over 100 random initializations, taken as a median over sparse random combinatorial Laplacians (Fig. 5). Over $n = 4$ to 12 the variance decays only *polynomially*: the log–log slopes of variance vs. n are -0.97 (j=0), -1.28 (j=1), and -1.36 (j=2), with the variance still of order 0.2 at $n = 12$. This mild, roughly $1/n$ decrease is far from the exponential vanishing that defines a barren plateau and is consistent with the super-polynomial variance lower bound PCE enjoys for its bilinear loss [1]. That bound is not proved for our rational deflation loss (Sec. IV-D), so this is empirical evidence of trainability, not verification of the asymptotic bound; we make no claim beyond $n = 12$.

H. Classification Performance and Generalization Failure

We evaluate whether the β_1 signal provides predictive value for crash-risk screening using a chronological train/test split. A crash event is defined as a drawdown exceeding 10% within 90 calendar days; under this definition, 37 of the 190 windows are positive. The dataset is split chronologically: windows whose start dates fall in 2003–2006 form the training set (used to select the classifier threshold), and windows starting in 2007–2010 form the held-out test set. Threshold selection maximizes F_1 on the training split only. The volatility-proxy baseline — the 30-day rolling standard deviation of log returns — is evaluated under the identical protocol.

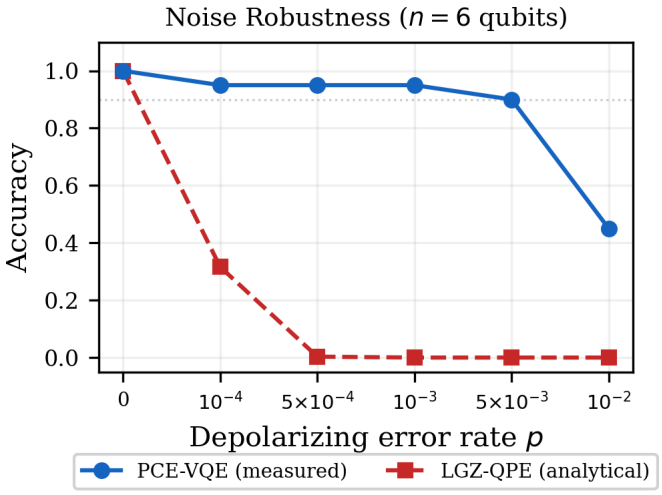

Fig. 4. Noise robustness (illustrative; not a matched-scale comparison): PCE-VQE accuracy (measured, blue, $n = 6$ qubits, $n_k \approx 63$ under $\kappa = 2$) vs. LGZ-QPE at $n_k = 16$ (measured, red solid) and analytical extrapolation (red dashed). PCE maintains >90% accuracy up to $p{=}5{\times}10^{-3}$; simulated LGZ-QPE collapses above $p{=}10^{-4}$, matching the analytical bound. The two curves are not at matched encoded problem size: PCE encodes $n_k \approx 63$ basis dimensions while LGZ operates on $n_k = 16$. The figure illustrates the depth/noise trade-off at moderate scale rather than a like-for-like comparison; see Section VII-F for explicit caveats.

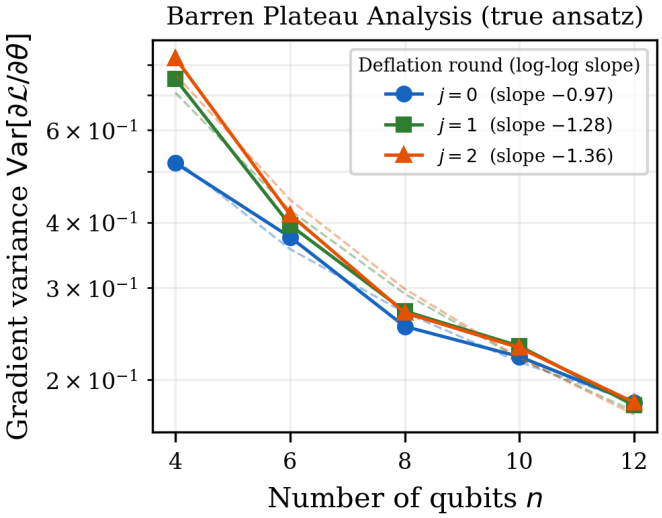

Fig. 5. Gradient-variance scaling of the true PCE Rayleigh-quotient loss through the hardware-efficient ansatz (exact adjoint gradient; deflation rounds $j \in \{0, 1, 2\}$, $n = 4$–12). The variance decays only polynomially (log–log slopes ≈ -1), staying order 0.2 at $n = 12$: a mild decrease, not the exponential vanishing of a barren plateau.

TABLE IV
CLASSIFICATION PERFORMANCE: CHRONOLOGICAL SPLIT

Classifier / split	Precision	Recall	F_1	ROC AUC
β_1 (train, 2003–2006)	0.870	0.741	0.800	0.943
β_1 (test, 2007–2010)	0.278	1.000	0.435	0.818
Volatility (train)	0.545	0.889	0.676	0.905
Volatility (test)	0.146	0.700	0.241	0.652

Test ROC AUC for β_1: bootstrap 95% CI $[0.71, 0.91]$ (1000 resamples); label-permutation $p = 6 \times 10^{-4}$. Test set has 10 positive windows of 64; metrics should be interpreted with attention to small-sample variance.

TABLE V
OUT-OF-DISTRIBUTION EVALUATION (FIXED 2003–2006 THRESHOLD)

Episode	Precision	Recall	F_1	ROC AUC
2020 COVID shock	0.000	0.000	0.000	0.009
2022 rate-cycle drawdown	0.333	0.083	0.133	0.515

On the held-out test set, the β_1-threshold classifier achieves precision 0.278, recall 1.000, F_1 0.435, and ROC AUC 0.818, compared to ROC AUC 0.652 for the volatility proxy (Table IV). We caution that the test split contains only 10 positive windows (of 64) over 4 years; the recall of 1.000 at precision 0.278 means the classifier fires on a large fraction of test windows rather than acting as a sharp early-warning signal. A bootstrap 95% confidence interval on the test ROC AUC (resampling test windows, 1000 resamples) spans $[0.71, 0.91]$, wide enough that the point estimate should not be over-read. To test significance despite the small sample, we permute the held-out labels 20,000 times and recompute the test AUC: the observed β_1 AUC of 0.818 is reached by only a $p = 6 \times 10^{-4}$ fraction of permutations, so it is well beyond chance, whereas the volatility proxy (0.652, $p = 0.07$) is not significant at the 5% level. In-sample metrics on the training split are reported for reference. For robustness, we additionally evaluate on two independent out-of-distribution crisis episodes not used for training or threshold selection: the 2020 COVID shock and the 2022 rate-cycle drawdown, with the fixed 2003–2006-trained threshold (Table V).

The 2020 OOD failure is a primary empirical finding about the β_1 feature.: A test ROC AUC of 0.009 on the COVID episode is essentially the *inverted* classifier: under the fixed 2003–2006 threshold the β_1 signal is *anti-correlated* with crash labels during the COVID shock. Two readings are consistent: a mechanism mismatch (the 2008 GFC was a slow build-up of correlated stress that loop-formation tracks, whereas the 2020 shock was a fast exogenous dislocation with different return-series topology), or a threshold artifact (thresholds fixed on 2003–2006 give arbitrarily bad fixed-threshold AUC under distributional shift). The 2022 chance-level AUC of 0.515 fits either, and disentangling them needs a regime-conditional classifier, left to future work. The in-regime AUC of 0.818 is thus evidence that β_1 carries structural information about GFC-style correlated-stress build-up, not that a single fixed-threshold β_1 classifier is a universal early-warning indicator.

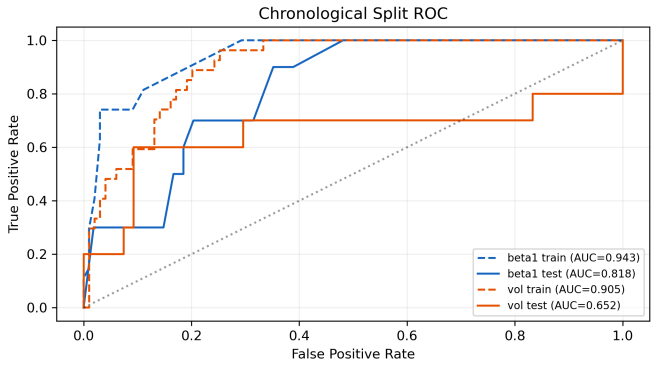

Fig. 6. ROC curves for the β_1 classifier and volatility proxy under the chronological train/test split. Training curves (dashed) show in-sample performance on 2003–2006; test curves (solid) show held-out performance on 2007–2010.

One clarification on scope: this classification study uses the classical exact-eigensolver β_1 of Sec. V, not a quantum estimate. The role of the quantum pipeline is to compute that same topological invariant, so the finding here is a property of the β_1 feature and its regime dependence, and it holds whether β_1 is obtained classically or by PCE-VQE. It is therefore a statement about the topological signal, not a claim about quantum learning.

I. Near-Term Hardware and Cost Crossover (Estimates)

To gauge executability, we estimate the cost of PCE-VQE at $n_k = 256$ ($n = 16$, $\kappa = 2$) on an IBM Heron-class device ($\epsilon_{2Q} = 3 \times 10^{-3}$, $\epsilon_{1Q} = 2 \times 10^{-4}$). Transpiled to $\{CX, SX, R_Z\}$, the $L = 32$ ansatz has $\approx$480 two-qubit and 2640 single-qubit gates, for a per-shot fidelity $(1-\epsilon_{2Q})^{480}(1-\epsilon_{1Q})^{2640} \approx 0.14$. At 10^4 shots per basis, $G = 3$ bases, 200 iterations per round, and $\beta_1 = 2$ rounds, the shot budget is $\approx 1.2 \times 10^7$, within the throughput of current cloud devices; scaling to $n_k \geq 10^3$ would need error mitigation or early fault-tolerant hardware. We have not run this on hardware. Extrapolating the same model against the CE benchmark (all 190 windows in $\approx$ 25.5 s) puts the cost crossover near $n_k \gtrsim 10^4$, where $O(n_k^3)$ diagonalisation overtakes the projected per-evaluation cost. That threshold compounds several extrapolations and is a research target, not a demonstrated crossover; datasets that reach it (multi-asset networks, tick-level embeddings) are plausible.

VIII. DISCUSSION

Structural advantages.: The PCE-variational approach offers no ancilla overhead (vs. $O(\log(1/\epsilon))$ for QPE) and shallow $O(n^2)$ circuits whose gradient variance decays only polynomially (not exponentially) across $n = 4$–12, consistent with PCE trainability [1] even though our rational loss is not formally covered by that proof. The matched-scale noise simulations at $n_k \in \{8, 16\}$ show PCE-VQE holding >90% accuracy where LGZ-QPE collapses; as Sec. VII-F notes, this

illustrates the depth/noise trade-off rather than a calibrated matched-encoding benchmark.

Limitations.: For $\kappa = 2$ PCE uses $O(\sqrt{n_k})$ qubits vs. $O(\log n_k)$ for QPE, so the advantage is in depth, not qubit count. The variational deflation protocol has no performance guarantee and spends one optimization per null vector, so its total budget grows linearly in β_k (the earlier $\beta_1 = 4$ "sensitivity" being a symptom of the random-start landscape rather than a separate issue, Sec. VII-B). This points to subspace methods such as SSVQE or weighted subspace search that target several eigenvectors at once, which we expect to be necessary for higher β_k. We also make no universal-approximation claim for the encoding. The null-space task needs something weaker than full expressivity, namely that the reachable set $\{\mathbf{c}(\vec{\theta})\}$ intersect $\ker(\Delta_k)$, and for deflation contain a spanning set of it, rather than cover all of $\mathbb{R}^{n_k}$. The real-window experiment of Sec. VII-C separates the failure modes. For a fully expressive state the $\kappa = 2$ correlators reach the null vectors of the real Laplacians (overlap ≈ 1, with a one-qubit margin at the minimal register size) even at $\beta_1 = 22$, so the encoding satisfies the feasibility condition. From a random start the hardware-efficient ansatz nonetheless recovers no null vector, but warm-starting from a classical null-space surrogate recovers β_1 exactly at every rung. The binding limitation is therefore the optimisation landscape of the hardware-efficient ansatz, not encoding expressivity; a fully quantum-independent recovery still needs a quantum route to that warm start.

When does this approach make sense?: If the toy-to-real gap is closed, PCE-variational would suit near-term hardware where depth is the binding constraint and moderate-scale problems ($n_k \sim 10^3$–10^5) where classical $O(n_k^3)$ diagonalisation is slow but full fault-tolerant QPE is not yet justified. We claim no quantum advantage at the scales demonstrated: at $n_k \in [31, 429]$ classical methods solve each window in ≤ 228 ms, so the contribution is methodological.

IX. Conclusion

We presented, to our knowledge, the first application of Pauli Correlation Encoding to quantum topological data analysis: a depth-efficient pipeline that recasts Betti-number estimation as a continuous-PCE Rayleigh-quotient minimization with deflation over shallow, ancilla-free circuits, at the cost of a rational loss outside the trainability proof of [1], for which the gradient variance decays only polynomially (no exponential barren plateau) over $n = 4$–12 qubits. Numerical experiments on S&P 500 data validate the classical stage against `ripser`. Run end-to-end on the real market Laplacians ($\beta_1 = 1$–22), the quantum stage recovers no null vector from a random start, but warm-starting from a classical null-space surrogate recovers β_1 exactly at every scale (Sec. VII-C); the obstacle is thus the optimisation landscape of the hardware-efficient ansatz, not the encoding, which does contain the null vectors. Chronologically split classification gives in-regime ROC AUC ≈ 0.818, but the calibration fails to generalize to the 2020 COVID and 2022 rate-cycle regimes, indicating that any deployment would require regime-aware calibration. Resource estimates extrapolate a crossover at $n_k \gtrsim 10^4$ (target, not demonstrated capability). The principal open question is therefore whether a quantum-native warm start, from a problem-informed or adaptive ansatz, can replace the classical null-space guidance used here and recover β_1 independently at real-data scale. We present this work as a methodology that points toward depth-efficient quantum TDA for financial stress, not as a demonstrated capability at that scale.

Future work, ordered by priority: (i) a quantum-native warm start, via problem-informed or adaptive ansätze, so that the real-window recovery of Sec. VII-C becomes independent rather than a classical-quantum hybrid; (ii) subspace deflation (SSVQE, weighted subspace search) for higher-multiplicity null spaces; (iii) regime-aware calibration of the β_1 classifier; (iv) persistent Betti numbers via [11]; (v) hardware execution following [1], [17]; and (vi) formal expressivity and trainability statements for the continuous-PCE loss.

Data Availability

The S&P 500 dataset, pipeline code, and PCE-VQE simulation scripts are available at https://github.com/arulrhikm/Quantum-Market-Crash-TDA.

AI Use Disclosure

In accordance with IEEE policy, we disclose that large language models were used only for writing assistance (clarifying prose and phrasing). All technical content, results, code, and figures were produced and verified by the authors; no AI system generated experimental data or numerical results. The authors take full responsibility for the accuracy, originality, and integrity of the submission, including every cited reference.

Acknowledgment

Classical computations were performed on Google Colab. Quantum statevector simulations were performed on BlueQubit.

References

[1] M. Sciorilli, L. Borges, T. L. Patti, D. García-Martín, G. Camilo, A. Anandkumar, and L. Aolita, "Towards large-scale quantum optimization solvers with few qubits," *Nature Communications*, vol. 16, p. 476, 2025, arXiv:2401.09421.

[2] U. Bauer, "Ripser: efficient computation of Vietoris–Rips persistence barcodes," *Journal of Applied and Computational Topology*, vol. 5, pp. 391–423, 2021.

[3] M. Gidea and Y. Katz, "Topological data analysis of financial time series: Landscapes of crashes," *Physica A: Statistical Mechanics and its Applications*, vol. 491, pp. 820–834, 2018.

[4] G. Carlsson, "Topology and data," *Bulletin of the American Mathematical Society*, vol. 46, no. 2, pp. 255–308, 2009.

[5] H. Edelsbrunner and J. L. Harer, *Computational Topology: An Introduction*. American Mathematical Society, 2010.

[6] N. Otter, M. A. Porter, U. Tillmann, P. Grindrod, and H. A. Harrington, "A roadmap for the computation of persistent homology," *EPJ Data Science*, vol. 6, no. 17, 2017.

[7] S. Lloyd, S. Garnerone, and P. Zanardi, "Quantum algorithms for topological and geometric analysis of data," *Nature Communications*, vol. 7, p. 10138, 2016.

[8] S. Ubaru, I. Y. Akhalwaya, M. S. Squillante, K. L. Clarkson, and L. Horesh, "Quantum topological data analysis with linear depth and exponential speedup," *arXiv preprint arXiv:2108.02811*, 2021.

[9] S. McArdle, A. Gilyén, and M. Berta, "A streamlined quantum algorithm for topological data analysis with exponentially fewer qubits," *arXiv preprint arXiv:2209.12887*, 2022.

[10] C. Gyurik, C. Cade, and V. Dunjko, "Towards quantum advantage via topological data analysis," *Quantum*, vol. 6, p. 855, 2022.

[11] R. Hayakawa, "Quantum algorithm for persistent Betti numbers and topological data analysis," *Quantum*, vol. 6, p. 873, 2022.

[12] I. Y. Akhalwaya, S. Ubaru, K. L. Clarkson, M. S. Squillante, V. Jejjala, Y.-H. He, K. Naidoo, V. Kalantzis, and L. Horesh, "Towards quantum advantage on noisy quantum computers," *arXiv preprint arXiv:2209.09371*, 2022.

[13] D. W. Berry, Y. Su, C. Gyurik, R. King, J. Basso, A. Del Toro Barba, A. Rajput, N. Wiebe, V. Dunjko, and R. Babbush, "Analyzing prospects for quantum advantage in topological data analysis," *PRX Quantum*, vol. 5, p. 010319, 2024.

[14] M. Crichigno and T. Kohler, "Clique homology is QMA_1-hard," *Nature Communications*, vol. 15, p. 9846, 2024.

[15] A. Schmidhuber and S. Lloyd, "Complexity-theoretic limitations on quantum algorithms for topological data analysis," *PRX Quantum*, vol. 4, p. 040349, 2023.

[16] V. P. Soloviev and M. Krompiec, "Large-scale portfolio optimization using Pauli Correlation Encoding," *arXiv preprint arXiv:2511.21305*, 2025.

[17] M. Sciorilli, G. Camilo, T. O. Maciel, A. Canabarro, L. Borges, and L. Aolita, "A competitive NISQ and qubit-efficient solver for the LABS problem," *arXiv preprint arXiv:2506.17391*, 2025.

[18] A. Zomorodian and G. Carlsson, "Computing persistent homology," *Discrete & Computational Geometry*, vol. 33, no. 2, pp. 249–274, 2005.

[19] A. Khandelwal and M. G. Chandra, "Quantum-enhanced topological data analysis: A peep from an implementation perspective," *arXiv preprint arXiv:2302.09553*, 2023.

[20] F. Takens, "Detecting strange attractors in turbulence," in *Dynamical Systems and Turbulence, Warwick 1980*, ser. Lecture Notes in Mathematics. Springer, 1981, vol. 898, pp. 366–381.

[21] H. Kantz and T. Schreiber, *Nonlinear Time Series Analysis*, 2nd ed. Cambridge University Press, 2004.

[22] IBM Quantum, "Pauli correlation encoding to reduce MaxCut requirements," 2025.

[23] M. Cerezo, A. Sone, T. Volkoff, L. Cincio, and P. J. Coles, "Cost function dependent barren plateaus in shallow parametrized quantum circuits," *Nature Communications*, vol. 12, p. 1791, 2021.

[24] S. Scali, C. Umeano, and O. Kyriienko, "The topology of data hides in quantum thermal states," *APL Quantum*, vol. 1, no. 3, p. 036106, 2024, arXiv:2402.15633.

[25] N. A. Nghiem, J. Lee, and T.-C. Wei, "Hybrid quantum-classical framework for Betti number estimation with applications to topological data analysis," *arXiv preprint arXiv:2508.01516*, 2025.

[26] D. Liu, "From Betti numbers to persistence diagrams: A hybrid quantum algorithm for topological data analysis," *arXiv preprint arXiv:2512.02081*, 2025.

[27] J. Padín-Martínez, V. P. Soloviev, A. Borrallo-Rentero, A. Rodríguez-Otero, R. Alfonso-Rodríguez, and M. Krompiec, "Pauli correlation encoding for budget-constrained optimization," *arXiv preprint arXiv:2602.17479*, 2026.

[28] M. B. Kennel, R. Brown, and H. D. I. Abarbanel, "Determining embedding dimension for phase-space reconstruction using a geometrical construction," *Physical Review A*, vol. 45, no. 6, pp. 3403–3411, 1992.

[29] O. Higgott, D. Wang, and S. Brierley, "Variational quantum computation of excited states," *Quantum*, vol. 3, p. 156, 2019.

[30] A. Kandala, A. Mezzacapo, K. Temme, M. Takita, M. Brink, J. M. Chow, and J. M. Gambetta, "Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets," *Nature*, vol. 549, pp. 242–246, 2017.