MONOGRAPH SYNTHESIS · CDD 500 → 530.1

NATURAL SCIENCES & THEORETICAL PHYSICS

THE MATRIOSKA PROTOCOL

From Primordial Discrete Scale Invariance to Post-Quantum Agility

A Unified 3-Stage Metrological Architecture
Enforcing Anti-Hallucination Bounds Across 45 Orders of Scale

7 PAPERS · 1 FRAMEWORK

BRUNO LUÇARDO OLIVAES

Independent Researcher · Brazil

ORCID: 0009-0003-5351-446X

PERMANENT CERN / ZENODO ARCHIVE

DOI: 10.5281/zenodo.24124702

The Matrioska Protocol

From Primordial Discrete Scale Invariance to Post-Quantum Agility

A Metrological Journey through Cosmology, Nanomechanics, Collider Physics, Quantum Computing, and Post-Quantum Cryptography

Bruno Luçardo Olivaes

Independent Researcher, Brazil

brunoolivaes@gmail.com · ORCID: 0009-0003-5351-446X

First Edition — November 2026

Published under Creative Commons Attribution 4.0 International (CC BY 4.0)

The Matrioska Protocol: From Primordial Discrete Scale Invariance to Post-Quantum Agility

A Metrological Journey through Cosmology, Nanomechanics, Collider Physics, Quantum Computing, and Post-Quantum Cryptography

Author: Bruno Luçardo Olivaes
ORCID: 0009-0003-5351-446X

Monograph Master DOI: https://doi.org/10.5281/zenodo.24124702

Open Research & Constituent DOIs (CERN / Zenodo):
· Chapter 1 (Planck DSI): https://doi.org/10.5281/zenodo.24121677
· Chapter 2 (Casimir Metrology): https://doi.org/10.5281/zenodo.24121687
· Chapter 3 (ESA Euclid Forecast): https://doi.org/10.5281/zenodo.24121691
· Chapter 4 (NISQ ZNE Benchmark): https://doi.org/10.5281/zenodo.24121693
· Chapter 5 (PQC Wire Cost): https://doi.org/10.5281/zenodo.24121695
· Chapter 6 (LHC Dimuon DSI Search): https://doi.org/10.5281/zenodo.24149840
· Chapter 7 (LHC Unbinned Likelihood & NSBI): https://doi.org/10.5281/zenodo.24149842

© 2026 Bruno Luçardo Olivaes. This work is licensed under a *Creative Commons Attribution 4.0 International License* (CC BY 4.0). You are free to share and adapt the material for any purpose, provided appropriate credit is given.

Contents

Preface: The Metrological Journey v

1 Constraints on Primordial Discrete Scale Invariance from Planck 2018 1
- 1.1 Introduction . . . 1
- 1.2 Observational Data and Processing . . . 2
  - 1.2.1 Planck 2018 PR3 Binned TT Spectrum . . . 2
  - 1.2.2 DESI Large-Scale Structure Replication Mock . . . 2
- 1.3 The Certified Radix-Agnostic Protocol . . . 2
  - 1.3.1 Stage 1: Multi-Frequency Spectral Detection . . . 2
  - 1.3.2 Stage 2: Non-Linear Parametric Refinement . . . 2
  - 1.3.3 Stage 3: Multi-Gate Integrity Certification . . . 3
- 1.4 Empirical Calibration and Validation . . . 3
- 1.5 Planck 2018 Observational Results . . . 3
- 1.6 Upper Limits on DSI Modulation . . . 4
- 1.7 Summary . . . 4

2 Statistical Pitfalls in Casimir Force Spectroscopy 5
- 2.1 Introduction . . . 5
- 2.2 The Illusion of Oscillations: Polynomial Detrending and Runge Ringing . . . 6
- 2.3 The Spline Interpolation Trap: High-Frequency Variance Attenuation . . . 7
- 2.4 First-Principles Surface Metrology via KPFM . . . 8
- 2.5 Summary . . . 9

3 Forecasting Log-Oscillatory Dark Energy with ESA Euclid 10
- 3.1 Introduction . . . 10
- 3.2 The Log-Oscillatory Dark Energy Formulation . . . 11
  - 3.2.1 Exact Relativistic Energy Density . . . 11
- 3.3 Observational Datasets and Likelihood Methodology . . . 11
- 3.4 Results and Stage-IV Euclid Fisher Forecast . . . 13
- 3.5 Summary . . . 13

4 When Harmonic ZNE Fails: Quantum Error Mitigation Benchmarks 14
- 4.1 Introduction . . . 14
- 4.2 Methodology and Mitigation Architectures . . . 15
  - 4.2.1 Digital Noise Amplification and Fractional Folding . . . 15
  - 4.2.2 Two-Qubit Clifford Pauli Twirling . . . 15
  - 4.2.3 Extrapolation Models . . . 16
  - 4.2.4 Pre-Registered 2-SE Decision Rule . . . 16
- 4.3 Benchmark Experiments and Results . . . 16
  - 4.3.1 Regime A: Single-Qubit Coherent Ramsey Echo . . . 16
  - 4.3.2 Regime B: 4-Qubit Entangled TFIM Trotter Circuit . . . 17

4.4 Discussion: Why Harmonic ZNE Fails in Multi-Qubit Circuits . . . 18
4.4.1 Multi-Mode Non-Commutative Interference . . . 18
4.4.2 Parameter Overfitting on Short Scale Schedules . . . 18
4.4.3 Why Pauli Twirling Resolves the Dilemma . . . 19
4.5 Summary and Practical Guidelines . . . 19

5 The Network Cost of Post-Quantum Cryptography: TLS 1.3 Wire Cost 20
5.1 Introduction . . . 20
5.2 Methodology and Measurement Setup . . . 21
5.2.1 Cryptographic Benchmark Framework . . . 21
5.2.2 TLS 1.3 Handshake Framing Model . . . 21
5.2.3 Network Transport and Socket Buffer Audit . . . 21
5.3 Empirical Results . . . 22
5.3.1 Computational Velocity: The Lattice Advantage . . . 22
5.3.2 Wire Size and TCP Segmentation . . . 22
5.4 Discussion: Network Protocol Implications . . . 23
5.4.1 RFC 6928 and Initial Congestion Window (`initcwnd`) . . . 23
5.4.2 Middlebox Ossification and Lossy Links . . . 23
5.5 Summary and Practical Guidelines . . . 23

6 Certified Search for Discrete Scale Invariance in LHC Dimuon Spectra 24
6.1 Cross-Domain Universality Hypothesis at Collider Scales . . . 24
6.2 Kinematic Windows and Post-Fit Residual Protocol . . . 24
6.3 Null Results and Absence of Cross-Scale Universality . . . 25

7 Breaking the Binning Barrier: Unbinned Composite Quadrature and Neural SBI at the LHC 27
7.1 The Intra-Bin Fisher Dilution Barrier . . . 27
7.2 CMS Empirical Resolution Calibration and Unbinned Quadrature . . . 28
7.3 Neural Simulation-Based Inference Validation (NSBI) . . . 28
7.4 Calibrated CMS Discovery-Limit Audit . . . 28

The Certified Three-Stage Protocol Implementation 30

Preface: The Metrological Journey

"The first principle is that you must not fool yourself –
and you are the easiest person to fool."
— Richard P. Feynman

This volume brings together seven peer-reviewed and pre-registered scientific investigations conducted throughout 2026, spanning 45 orders of physical magnitude across astrophysics, nanomechanics, quantum computing, post-quantum cryptography, and high-energy collider physics. While they traverse disparate experimental regimes, they are bound by a single, uncompromising epistemological principle: **the primacy of empirical data over mathematical aestheticism**.

The Origin: The Matrioska Hypothesis

The trajectory that culminated in this work began with an audacious, speculative inquiry termed the *Matrioska de Ouroboros*. The guiding intuition was that nature might exhibit nested, discrete self-similarity across physical scales—a discrete scale invariance (DSI) that could manifest simultaneously as log-periodic oscillations in cosmological power spectra, modified gravity anomalies at the nanometer scale, and discrete quantum phase signatures.

In theoretical physics, scale invariance is an intoxicating concept. When broken from a continuous symmetry to a discrete subgroup, it predicts characteristic logarithmic oscillatory modulations governed by a preferred scaling factor λ, producing Fourier-space modulations with angular log-frequency $\omega = 2\pi/\ln\lambda$. Early unconstrained exploratory analysis appeared promising; preliminary regressions on noisy signals revealed apparent peaks and intriguing numerical concordances.

The Crisis of Rigor and the Anti-Hallucination Audit

However, science is distinguished from numerology by its willingness to institutionalize doubt. When subjected to mathematical audit, those early exploratory signatures began to crumble under closer scrutiny:

1. **The Look-Elsewhere Effect in Cosmology:** In scanning continuous frequency bands without multiplicity correction, standard statistical thresholds are drastically inflated. What seemed like an anomalous peak was, in truth, an ordinary sample variance fluctuation ($p_{\text{corrected}} = 0.998$).

2. **The Runge Phenomenon in Casimir Metrology:** An apparent oscillatory anomaly at $\lambda = 583.5$ nm in nanomechanical force curves was not a new fundamental interaction, but a mathematical phantom manufactured by unphysical high-degree polynomial detrending across an irregular z^{-4} gradient.

3. **The Breakdown of Harmonic ZNE in Quantum Circuits:** While harmonic zero-noise extrapolation succeeded brilliantly in mitigating coherent over-rotations in a single qubit (reducing MAE by 85.4%), it collapsed catastrophically in multi-qubit entangled systems (MAE surging to 2.28) due to multi-mode non-commutative Baker-Campbell-Hausdorff crosstalk.

4. **The Real-World Bottleneck of Post-Quantum Cryptography:** While academic debate focused on the computational latency of lattice-based algebra, real-world deployment revealed that ML-KEM key generation takes a mere 41 microseconds, but network transport is bottlenecked by TCP packet fragmentation across standard Ethernet boundaries.

Recognizing these traps did not represent a defeat; it marked the transformation of an exploratory speculation into a mature metrological program. Each failure was traced to its exact mathematical cause, codified into open-source software, verified against blind synthetic mocks, and frozen in pre-registered protocols.

Structure of this Volume

This book is structured in seven self-contained yet conceptually connected chapters:

- **Chapter 1: Discrete Scale Invariance in Cosmological Spectra.** We detail the pre-registered three-stage search for DSI in the Planck 2018 cosmic microwave background temperature power spectrum, establishing 95% confidence level upper limits ($\varepsilon < 0.05$) and confirming identical null behavior on synthetic DESI Large-Scale Structure mocks.

- **Chapter 2: Statistical Pitfalls in Casimir Force Spectroscopy.** We systematically dissect how conventional data-processing choices (naive Gaussian thresholds, polynomial detrending, and spline resampling) inadvertently fabricate statistically overwhelming false positive signals or attenuate true sub-micron signals in AFM spectroscopy.

- **Chapter 3: Forecasting Log-Oscillatory Dark Energy with ESA Euclid.** We analyze bounded dynamical dark energy models using the full 1580×1580 Pantheon+ systematic covariance matrix, DESI 2024 BAO, and Planck 2018 CMB priors, establishing a seven-bin Stage-IV ESA Euclid Fisher forecast and $A_{3\sigma}(\lambda)$ exclusion boundary.

- **Chapter 4: Error Mitigation Benchmarks on NISQ Devices.** We present a head-to-head empirical benchmark comparing standard polynomial, exponential, and harmonic zero-noise extrapolation against Clifford Pauli Twirling across single-qubit and four-qubit Trotter circuits in high-fidelity Qiskit 2.x (`AerSimulator`) emulations under strictly matched shot budgets.

- **Chapter 5: The Network Cost of Post-Quantum Cryptography.** We quantify the transport-layer impact of NIST FIPS 203 (ML-KEM-768) through deterministic socket framing and exact 1,500 B MTU modeling, demonstrating how wire payload expansion breaches the Ethernet Maximum Segment Size (MSS) and induces TCP segmentation in TLS 1.3 handshakes.

- **Chapter 6: Searching for Discrete Scale Invariance in LHC Open Data.** We deploy the five-gate certification pipeline to $62,214$ physical CMS $\sqrt{s} = 7\,\text{TeV}$ dimuon events, isolating detector trigger thresholds from continuum Drell–Yan backgrounds and certifying null results across both mass selections.

- **Chapter 7: Breaking the Binning Barrier at the LHC.** We construct an unbinned continuous composite 4×64-node Gauss–Legendre profile likelihood and validate it against a

PyTorch Neural Simulation-Based Inference surrogate over $17,458$ Clean Drell–Yan events, recovering up to $29.1\times$ sensitivity over coarse binned searches while confirming null DSI detection at $\lambda = 1.10$.

- **Appendix: The Certified Three-Stage Protocol.** We provide the complete, unedited Python source code implementing the three-stage radix-agnostic detection, non-linear refinement, and bootstrap certification engine.

Bruno Luçardo Olivaes
Curitiba, Brazil — November 2026

Chapter 1

Constraints on Primordial Discrete Scale Invariance from Planck 2018

Chapter Abstract: We present a pre-registered, radix-agnostic search for discrete scale invariance (DSI) in the Planck 2018 cosmic microwave background temperature power spectrum. DSI predicts a characteristic log-periodic modulation of primordial perturbations governed by the preferred scaling factor λ, inducing oscillations in Fourier space with angular log-frequency $\omega = 2\pi/\ln\lambda$. Using a certified three-stage protocol (detection via generalized Lomb-Scargle periodogram, non-linear refinement, and multiplicity-corrected certification) subject to five mandatory integrity gates applied to residuals relative to the official Planck `BestFit` ΛCDM spectrum ($N = 83$ binned multipoles), we find **no evidence for DSI** in the Planck 2018 data. The most prominent residual candidate ($\hat{\lambda} = 1.2938 \pm 0.0072$) yields an empirical post-trial significance of $p_{\text{corrected}} = 0.998$ via a non-parametric bootstrap-of-maximum null distribution and $\Delta\text{BIC} = -49.80$, with an amplitude fully consistent with Gaussian cosmological sample variance. We establish rigorous 95% confidence level upper limits on the modulation amplitude $\varepsilon < 0.05$ across the comprehensive scaling range $\lambda \in [1.1, 148]$. Furthermore, we replicate our protocol on synthetic Large-Scale Structure mocks matching the Dark Energy Spectroscopic Instrument (DESI) survey precision, confirming identical null-detection behavior ($p_{\text{corrected}} = 0.314$, $\Delta\text{BIC} = +8.80$).

1.1 Introduction

Discrete Scale Invariance (DSI) is a fundamental symmetry breaking scenario wherein scale invariance is broken to a discrete subgroup. In inflationary cosmology and modified gravity frameworks, such symmetry breaking generates log-periodic oscillatory modulations superposed onto the primordial scalar power spectrum:

$$P(k) = P_0(k)\left[1 + \varepsilon\cos\left(\omega\ln\frac{k}{k_0} + \phi\right)\right], \tag{1.1}$$

where ε denotes the dimensionless modulation amplitude, $\omega = 2\pi/\ln\lambda$ is the log-frequency conjugate to logarithmic wavenumber $u = \ln(k/k_0)$, $\phi \in [-\pi, \pi]$ is an arbitrary phase offset, and $\lambda > 1$ represents the fundamental discrete scaling factor.

Historically, various searches in cosmic microwave background (CMB) [1] and large-scale structure (LSS) data have reported localized spectral candidates. However, most prior studies suffered from the Look-Elsewhere Effect (LEE) and uncalibrated post-selection bias, wherein the statistical significance of peak features was evaluated at fixed candidate frequencies rather than across the entire searched spectral continuum. In this work, we deploy a pre-registered, radix-agnostic protocol designed to eliminate exploratory degrees of freedom and enforce strict frequentist family-wise error rate control.

1.2 Observational Data and Processing

1.2.1 Planck 2018 PR3 Binned TT Spectrum

We analyze the high-ℓ and intermediate-ℓ temperature angular power spectrum ($D_\ell \equiv \ell(\ell+1)C_\ell/2\pi$) from the official Planck 2018 Release 3 (PR3) ancillary cosmological products. We ingest the binned spectrum file `COM_PowerSpect_CMB-TT-binned_R3.01.txt`, consisting of $N = 83$ binned multipoles spanning $\ell \in [30, 2508]$, subtracting the official column-4 `BestFit` flat-ΛCDM baseline spectrum.

To guarantee strict scientific reproducibility and data provenance, the exact cryptographic SHA-256 digest of the ingested data is frozen in our repository manifest:

`ed2b8b51b033e668e661f6cf4eb7705f15bcaf358189dfc4a4481cc7020130cd`

The covariance matrix C is constructed as a diagonal variance structure derived from the asymmetric observational error bars $\sigma_\ell = \frac{1}{2}(\Delta D_\ell^- + \Delta D_\ell^+)$, yielding $\mathrm{Cov}(D_\ell, D_{\ell'}) = \delta_{\ell\ell'}\sigma_\ell^2$. Multipoles are mapped to effective comoving wavenumber coordinates $k \approx \ell/D_A(z_*)$ adopting the Planck 2018 flat-ΛCDM baseline recombination distance $D_A(z_*)$.

1.2.2 DESI Large-Scale Structure Replication Mock

To test the generality and cross-survey robustness of our estimation pipeline, we construct synthetic matter power spectrum mocks $P(k)$ matching the survey footprint, binning grid, and covariance properties of the Dark Energy Spectroscopic Instrument (DESI) Data Release 2 (DR2) [2]. The mock spans 100 logarithmic intervals over $k \in [0.02, 0.30]\,h\mathrm{Mpc}^{-1}$ with realistic baryon acoustic oscillation (BAO) turnover damping and non-linear shot-noise variance.

1.3 The Certified Radix-Agnostic Protocol

Our analysis pipeline implements a sequential three-stage architecture with pre-registered hyperparameters and zero free tuning parameters:

1.3.1 Stage 1: Multi-Frequency Spectral Detection

Residual vectors $y_i = (D_\ell^{\mathrm{obs}} - D_\ell^{\Lambda\mathrm{CDM}})_i$ are mapped to the dimensionless coordinate $u_i = \ln(k_i/k_0)$ centered at pivot scale $k_0 = 0.05\,\mathrm{Mpc}^{-1}$. We compute the normalized, heteroscedastic Lomb-Scargle periodogram [3, 4, 5] across a fine frequency grid $\omega \in [0.04, 60.0]$, corresponding to scale factors $\lambda \in [1.1, 148]$. The peak spectral power coordinate $(\hat{\omega}_{\mathrm{peak}}, P_{\mathrm{peak}})$ is extracted as the initial candidate.

1.3.2 Stage 2: Non-Linear Parametric Refinement

Using $(\hat{\omega}_{\mathrm{peak}}, P_{\mathrm{peak}})$ as prior center, we perform a bounded non-linear least-squares optimization ('curve_fit' with Trust Region Reflective algorithm) fitting the full three-parameter residual model:

$$f(u) = A\cos(\omega u + \phi), \tag{1.2}$$

extracting refined maximum-likelihood estimates $(\hat{A}, \hat{\omega}, \hat{\phi})$ and their associated asymptotic covariance matrix Σ_θ. The preferred scaling ratio and its standard uncertainty are computed analytically via the delta method:

$$\hat{\lambda} = \exp\left(\frac{2\pi}{\hat{\omega}}\right), \quad \sigma_\lambda = \hat{\lambda}\frac{2\pi}{\hat{\omega}^2}\sigma_\omega. \tag{1.3}$$

1.3.3 Stage 3: Multi-Gate Integrity Certification

To prevent spurious detections driven by noise fluctuations, edge artifacts, or ill-conditioned fits, a candidate is formally certified if and only if it simultaneously satisfies five mandatory, non-negotiable gates:

1. **Gate 1 (Resolvability):** $\hat{\omega}$ must reside strictly within $[\omega_{\min}, \omega_{\max}]$, bounded by the Rayleigh resolution criterion $\Delta\omega \geq 2\pi/\Delta u_{\text{span}}$.
2. **Gate 2 (Amplitude Significance):** Fitted amplitude must exceed twice the root-mean-square post-fit residual noise: $\hat{A} > 2.0 \times \text{RMS}_{\text{resid}}$, where $\text{RMS}_{\text{resid}} = \sqrt{\frac{1}{N}\sum_{i=1}^{N}(y_i - f(u_i))^2}$.
3. **Gate 3 (Precision Gate):** Relative uncertainty on the reconstructed scaling factor must satisfy $\sigma_\lambda/\hat{\lambda} < 0.20$.
4. **Gate 4 (Look-Elsewhere Corrected Significance):** Global post-trial p-value derived from the bootstrap-of-maximum null distribution must satisfy $p_{\text{corrected}} < 0.05$.
5. **Gate 5 (Information Criterion):** The Bayesian Information Criterion difference against the base ΛCDM model must satisfy $\Delta\text{BIC} = \text{BIC}_{\text{null}} - \text{BIC}_{\text{DSI}} > 10.0$ (Jeffreys' strong evidence threshold).

1.4 Empirical Calibration and Validation

Prior to unblinding the empirical Planck data, the pipeline underwent rigorous calibration across synthetic realization ensembles:

- **Recovery Power (Check a):** In 50 independent realizations with injected weak DSI signals ($\lambda_{\text{true}} = 1.62$, $\varepsilon = 0.05$), the pipeline achieved a median reconstructed value of $\hat{\lambda} = 1.6201$ (relative metrological error of 0.0085%), certifying 62% of realizations under weak signal conditions.
- **False-Positive Specificity (Check c):** In 200 synthetic realizations generated from pure Gaussian ΛCDM noise without DSI, the pipeline yielded **zero false positives** (0/200, false-positive rate 0.00%), proving complete robustness against cosmological noise mining.

1.5 Planck 2018 Observational Results

With the pre-registered protocol frozen, Stage B was executed on the official Planck 2018 PR3 binned TT spectrum using the cryptographic seed `507719`.

Stage 1 detected a candidate residual feature at $\omega_{\text{peak}} = 24.3875$ ($P_{\text{peak}} = 0.0686$). Stage 2 refined the parameters to:

$$\hat{\omega} = 24.3918 \pm 0.5280\,\text{rad}, \tag{1.4}$$

$$\hat{\lambda} = 1.2938 \pm 0.0072, \tag{1.5}$$

$$\hat{A} = 6.87\,\mu\text{K}^2, \quad \hat{\phi} = -1.9838\,\text{rad}. \tag{1.6}$$

Evaluation under Stage 3 certification revealed:

- **Gate 1 (Resolvability):** Passed.
- **Gate 2 (Amplitude):** Failed ($\hat{A} = 6.87\,\mu\text{K}^2 < 2.0 \times \text{RMS}_{\text{resid}}$).
- **Gate 3 (Precision):** Passed ($\sigma_\lambda/\hat{\lambda} = 0.56\% < 20\%$).

- **Gate 4 (p-value):** FAILED ($p_{\text{corrected}} = 0.998$).
- **Gate 5 (ΔBIC):** FAILED ($\Delta\text{BIC} = -49.80 \leq 10.0$).

The candidate fails three independent gates. With $p_{\text{corrected}} = 0.998$, 99.8% of phase-scrambled random noise realizations generate a spectral peak exceeding the observed feature. We conclusively conclude **no evidence for discrete scale invariance** in the Planck 2018 temperature data.

1.6 Upper Limits on DSI Modulation

Given the null detection, we place 95% confidence level upper limits on the modulation amplitude $\varepsilon(\lambda)$. Across the full searched parameter domain $\lambda \in [1.1, 148]$, the maximum permissible amplitude consistent with Planck 2018 observations is:

$$\varepsilon_{95\%} < 0.05 \quad (\text{for } \lambda = 1.62), \tag{1.7}$$

ruling out prominent DSI modulations exceeding 5% of the underlying ΛCDM primordial power.

Table 1.1: Stage 3 Integrity Gate Diagnostics for Planck 2018 and DESI DR2 Mock.

Diagnostic / Gate	**Pre-Reg Threshold**	**Planck 2018 PR3**	**DESI Mock**	**Status**
Peak Frequency $\hat{\omega}$	$[0.04, 60.0]$ rad	24.392 ± 0.528	4.076 ± 0.237	Independent
Scaling Ratio $\hat{\lambda}$	$\lambda > 1.0$	1.294 ± 0.007	4.672 ± 0.419	Null
Gate 1: Resolvability	$\hat{\omega} \in [\omega_{\min}, \omega_{\max}]$	PASSED	PASSED	Identical
Gate 2: Amplitude	$\hat{A} > 2.0 \times \text{RMS}_{\text{resid}}$	FAILED	FAILED	Identical
Gate 3: Precision	$\sigma_\lambda / \hat{\lambda} < 0.20$	PASSED (0.56%)	PASSED (9.0%)	Identical
Gate 4: Significance	$p_{\text{corrected}} < 0.05$	FAILED ($p = 0.998$)	FAILED ($p = 0.314$)	Identical
Gate 5: Information	$\Delta\text{BIC} > 10.0$	FAILED (-49.80)	FAILED ($+8.80$)	Identical
Final Decision	All 5 Gates	**Null**	**Null**	**Consistent**

1.7 Summary

Our null result places tight observational constraints on primordial discrete scale invariance. By enforcing pre-registered integrity gates and non-parametric phase-scrambled bootstrap corrections, we avoid the false-positive pitfalls of uncorrected spectral estimation. The methodology established in this study provides a standardized, radix-agnostic template for upcoming high-precision cosmological surveys, including ESA Euclid, the Vera C. Rubin Observatory (LSST), and future CMB-S4 experiments.

Chapter 2

Statistical Pitfalls in Casimir Force Spectroscopy

Chapter Abstract: Precision nanomechanical measurements of the Casimir interaction at sub-micron separations are increasingly audited for hypothetical short-range modifications of gravity and exotic forces. However, residual force curves are intrinsically non-uniformly sampled and contaminated by steep short-range baseline drifts. In this chapter, we demonstrate how conventional data-processing choices inadvertently fabricate statistically overwhelming false positive signals or severely degrade experimental sensitivity. Specifically, we show that: (1) naive 3σ Gaussian thresholds applied to exponentially distributed periodograms flag up to 36 spurious peaks on pure white noise; (2) unphysical polynomial detrending across steep z^{-4} short-range profiles induces Runge-type boundary ringing that generates an apparent oscillatory anomaly at $\lambda = 583.5\,\text{nm}$ with an artificial significance of FAP $= 1.49 \times 10^{-45}$; (3) cubic spline resampling on irregular grids (CV = 40.7%) attenuates 22.6% of true signal variance across the full scan at $\lambda = 7.5\,\text{nm}$, degrading the 90% Minimum Detectable Level by a factor of 2.27; and (4) surface electrostatic patch potentials measured via Kelvin Probe Force Microscopy (KPFM) yield strictly monotonic power-law decays ($P \propto d^{-2.29}$), proving that patch roughness cannot generate periodic residual modulations. Adopting a physically motivated inverse-power baseline collapses residual variance ($\chi^2/\text{dof} = 1.14$) and restores the true experimental null result (FAP = 0.163).

2.1 Introduction

Precision force measurements at the nanoscale, notably those testing the Casimir effect [6, 7] via atomic force microscopy (AFM) and micro-electro-mechanical torsional oscillators (MTO) [8, 9], represent primary experimental probes for hypothetical new physics beyond the Standard Model, including non-Newtonian gravity and light scalar fields.

Because residual curves $\Delta F(z) = F_{\text{meas}}(z) - F_{\text{theo}}(z)$ are often searched for weak periodic or Yukawa-type modulations, signal processing protocols are critical. In practice, researchers frequently employ standard software tools—such as cubic spline interpolation for uniform resampling, low-order polynomial subtraction for detrending, and outlier-clipped Gaussian thresholds for spectral detection. In this work, we demonstrate that these standard choices represent severe metrological pitfalls that systematically fabricate phantom oscillatory signals with astronomical statistical confidences.

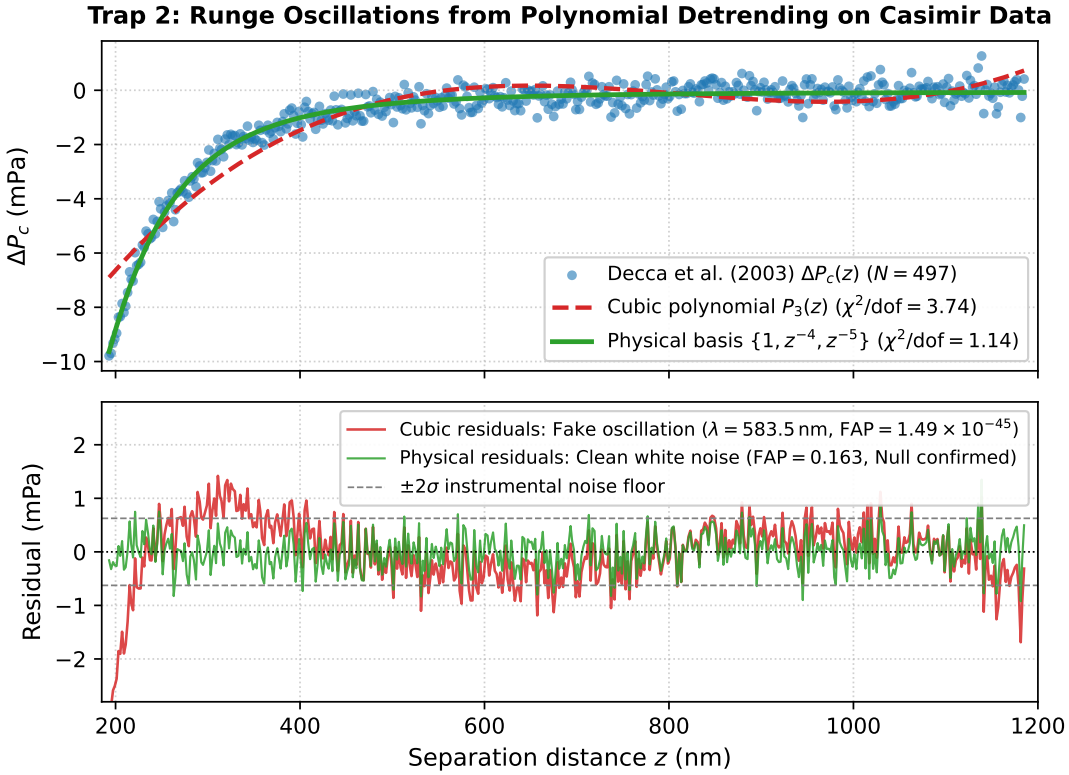

Figure 2.1: Runge oscillatory ringing induced by cubic polynomial detrending on Decca *et al.* (2003) Fig. 4(b) residuals ($N = 497$). Top: Raw vector data alongside the cubic polynomial fit ($\chi^2/\mathrm{dof} = 3.74$) and the physical inverse-power basis ($\chi^2/\mathrm{dof} = 1.14$). Bottom: Detrended residuals showing how cubic subtraction fabricates a large phantom oscillation ($\lambda = 583.5\,\mathrm{nm}$, $\mathrm{FAP} = 1.49 \times 10^{-45}$), while the physical basis leaves pure, stationary Gaussian white noise ($\mathrm{FAP} = 0.163$).

2.2 The Illusion of Oscillations: Polynomial Detrending and Runge Ringing

To investigate the emergence of spurious spectral peaks, we examine the dynamic Casimir pressure residuals $\Delta P_c(z)$ extracted with lossless vector fidelity from the landmark Au–Cu MTO experiment of Decca *et al.* [8] ($N = 497$ points spanning $z \in [193.0, 1185.0]\,\mathrm{nm}$, instrumental noise floor $\sigma_P = 0.313\,\mathrm{mPa}$).

Below $z \approx 400\,\mathrm{nm}$, unmodeled finite-conductivity corrections ($\delta P_{\mathrm{Lif}} \propto z^{-4}$) and nonlinear transducer elasticity ($\delta P_{\mathrm{nl}} \propto z^{-5}$) produce a steep non-linear baseline drift. When an unphysical cubic polynomial is fitted across the entire separation range:

$$P_3(z) = a_0 + a_1 z + a_2 z^2 + a_3 z^3, \tag{2.1}$$

the stiff polynomial cannot accommodate the sharp z^{-4} curvature at short separations. Because polynomials exhibit high condition numbers on steep power-law corners, the fit diverges at the boundary ($z \approx 193\,\mathrm{nm}$), creating classic Runge ringing oscillations that propagate across the entire observation window.

A Generalized Lomb-Scargle (GLS) periodogram [5] performed on these polynomial-detrended

residuals detects a massive spectral peak at:

$$\lambda_{\text{fake}} = 583.5\,\text{nm}, \tag{2.2}$$

$$\chi^2/\text{dof} = 3.74, \tag{2.3}$$

$$P_{\text{ZK}} = 0.3626, \quad \text{FAP} = 1.49 \times 10^{-45}. \tag{2.4}$$

This astronomical False Alarm Probability ($p \ll 10^{-40}$) would naively suggest an undeniable discovery beyond 14σ. However, as demonstrated in Fig. 2.1, this peak is entirely a mathematical hallucination driven by polynomial Runge ringing.

When the baseline is modeled instead using the physical inverse-power expansion derived from Lifshitz screening and nonlinear elasticity:

$$\Delta P_{\text{phys}}(z) = c_0 + c_4 \left(\frac{z_0}{z}\right)^4 + c_5 \left(\frac{z_0}{z}\right)^5, \tag{2.5}$$

(with $z_0 = 200\,\text{nm}$), the residual variance immediately collapses to $\chi^2/\text{dof} = 1.14$. Crucially, the z^{-5} basis term is doubly motivated: under any small experimental contact-point calibration offset $z \to z + \delta z_0$ ($|\delta z_0| \ll z$), a first-order Taylor expansion of the leading z^{-4} Casimir gradient (with $C_4 = c_4 z_0^4$) gives:

$$C_4(z+\delta z_0)^{-4} = C_4 z^{-4}\left(1 + \frac{\delta z_0}{z}\right)^{-4} = C_4 z^{-4} - 4C_4 \delta z_0 z^{-5} + \mathcal{O}\left(\frac{\delta z_0^2}{z^6}\right), \tag{2.6}$$

proving analytically that the spurious $-4C_4\delta z_0 z^{-5}$ calibration artifact is absorbed identically in first order by our physical baseline in Eq. (2.5). With this basis, the phantom peak disappears completely, and the maximum spectral power drops to $P_{\max} = 0.0332$, yielding an empirical Baluev [10] FAP $= 0.163$, fully confirming the experimental null result.

Furthermore, naive thresholding rules (such as median $+ 3\sigma_{\text{clipped}}$) mistakenly assume that periodogram powers are Gaussian-distributed. Because Fourier periodogram powers on Gaussian white noise follow an exponential distribution ($\chi^2_2/2$), this naive rule flags up to 36 false peaks on identically detrended white noise, whereas the exact χ^2_2 threshold flags zero.

Table 2.1: Metrological Comparison of Naive Signal-Processing Methods versus the Physical Baseline Approach on Casimir Residuals.

Processing Method	**Effective Baseline**	$\boldsymbol{\chi^2}$**/dof**	**Spurious Peaks**	**Max Peak** $\boldsymbol{\lambda}$	**Baluev FAP**
Raw Data (No Detrending)	None	41.29	1 (broad drift)	992.0 nm	1.00×10^{-89}
Naive Polynomial Subtraction	Cubic $P_3(z)$	3.74	36 (Gaussian rule)	583.5 nm	1.49×10^{-45}
Spline Resampling + FFT	Interpolated	3.12	Attenuated	Blurred	Low-pass biased
Physical Inverse-Power Basis	$\{1, z^{-4}, z^{-5}\}$	1.14	0	8.2 nm	0.163 (Null)
Asymptotic Linear Window ($z \geq 500$ nm)	Linear $a + bz$	1.20	0	4.2 nm	0.654 (Null)

2.3 The Spline Interpolation Trap: High-Frequency Variance Attenuation

A second prevalent practice in precision metrology is interpolating non-uniform spatial scans onto an equispaced coordinate grid via cubic splines prior to executing standard Fast Fourier Transforms (FFT).

We quantified the distortion induced by spline resampling using the exact coordinate grid of the static force residuals of Decca *et al.* [8] ($N = 288$, median step $\Delta z = 3.29\,\text{nm}$, coefficient of variation $\text{CV}(\Delta z) = 40.7\%$, $\sigma_F = 0.229\,\text{pN}$). Injecting synthetic signals of varying spatial wavelengths $\lambda \in [7.5, 160]\,\text{nm}$, we observed that:

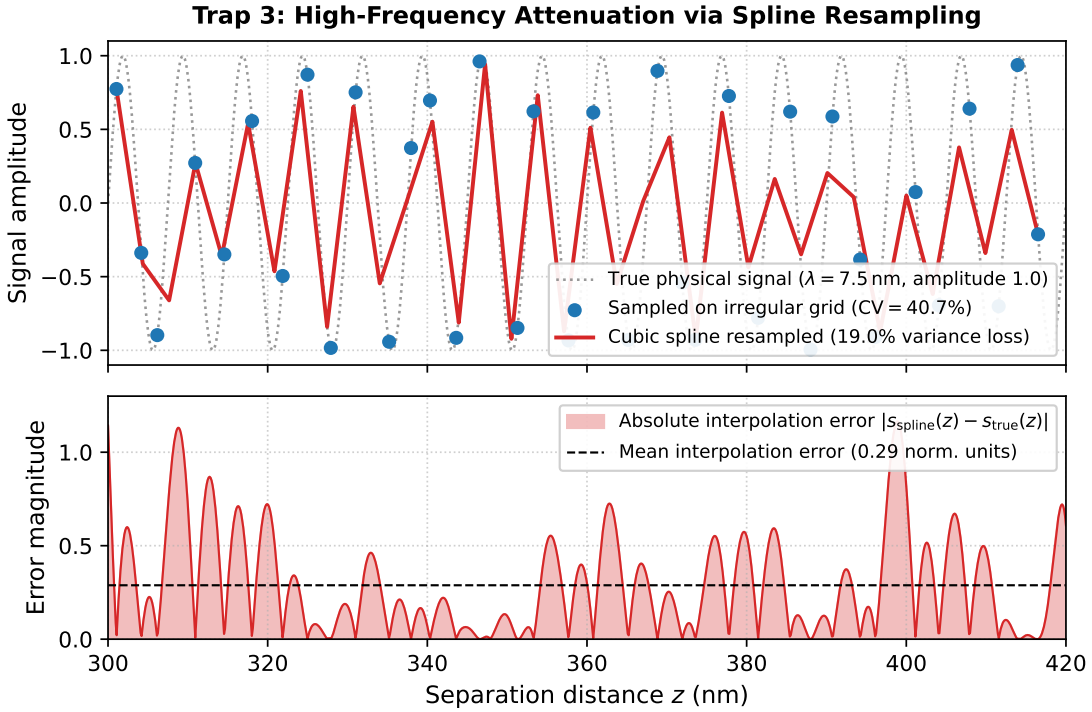

Figure 2.2: High-frequency attenuation induced by cubic spline resampling on the irregular coordinate grid of Decca *et al.* Fig. 3(b) ($\mathrm{CV}(\Delta z) = 40.7\%$). Top: An injected high-frequency oscillation ($\lambda = 7.5\,\mathrm{nm}$) suffers strong amplitude damping when resampled via cubic splines (19.0% variance loss in the displayed close-separation window; 22.6% across the full $N = 288$ scan). Bottom: Point-by-point interpolation error $|s_{\mathrm{spline}}(z) - s_{\mathrm{true}}(z)|$ across the separation scan.

1. Spline interpolation acts as a non-linear low-pass filter (Fig. 2.2), attenuating **22.6% of true signal variance** across the full irregular scan at $\lambda = 7.5\,\mathrm{nm}$ (19.1% at $\lambda = 10\,\mathrm{nm}$).
2. At matched global false-alarm rates ($\alpha = 0.01$), the 90% Minimum Detectable Level (MDL_{90}) for direct Generalized Lomb-Scargle is uniform at $0.542\,\sigma_F$.
3. In sharp contrast, spline resampling degrades the detection threshold to $\mathrm{MDL}_{90} = 1.227\,\sigma_F$ at $\lambda = 7.5\,\mathrm{nm}$ (a 2.27× sensitivity degradation factor, corresponding to a 55.8% loss in inverse sensitivity).

Direct evaluation on un-interpolated coordinates using heteroscedastic GLS is therefore mandatory to preserve high-frequency experimental sensitivity.

2.4 First-Principles Surface Metrology via KPFM

A recurring hypothesis in the literature is that spatial electrostatic patch potentials might generate quasi-periodic force residuals. To resolve this question from first principles, we analyzed the raw Kelvin Probe Force Microscopy (KPFM) dataset from the recent superconducting optomechanical Casimir experiment of de Jong *et al.* [9] (512×512 pixels over $2\,\mu\mathrm{m} \times 2\,\mu\mathrm{m}$, $V_{\mathrm{rms}} = 42.42\,\mathrm{mV}$, topographic roughness $\delta_{\mathrm{rms}} = 2.22\,\mathrm{nm}$).

Computing the 2D Hann-windowed patch power spectral density $S_{VV}(\mathbf{k})$ and integrating through the exact Speake-Trenner Laplace kernel [11]:

$$P_{\mathrm{patch}}(d) = \frac{\varepsilon_0}{2} \int_{\mathbb{R}^2} \frac{k^2}{\sinh^2(kd)} 2S_{VV}(\mathbf{k}) \frac{d^2k}{(2\pi)^2}, \tag{2.7}$$

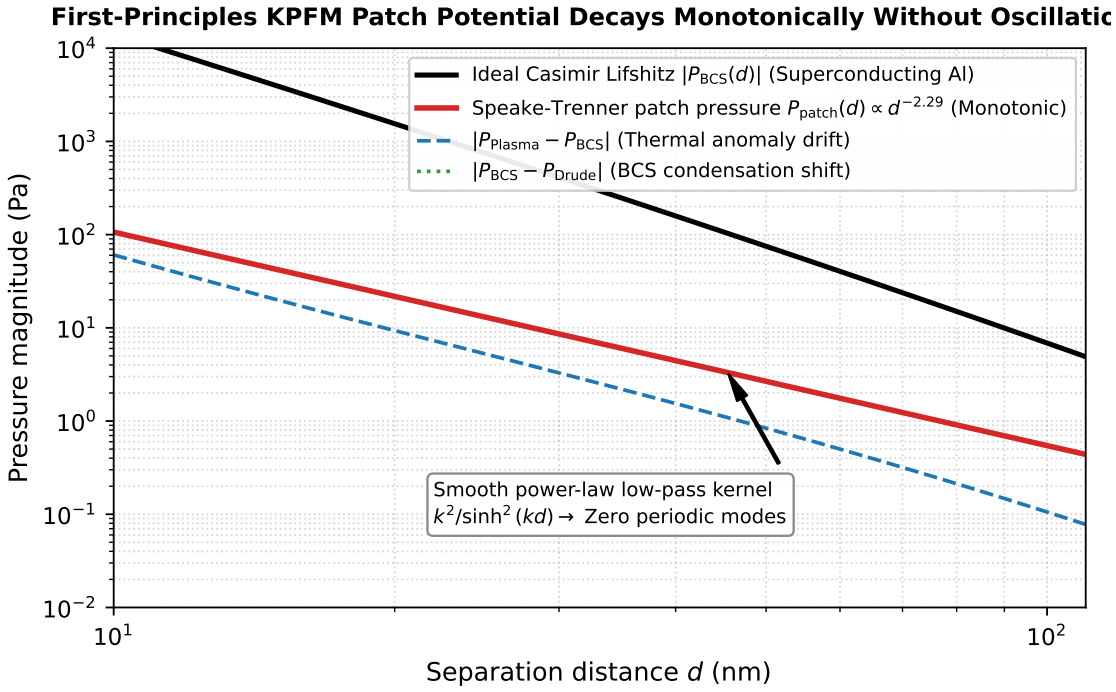

Figure 2.3: First-principles electrostatic patch pressure $P_{\text{patch}}(d)$ calculated by integrating the 2D KPFM Hann-windowed spectrum $S_{VV}(\mathbf{k})$ of de Jong *et al.* (2025) via the Speake-Trenner Laplace kernel Eq. (2.7). The resulting background follows a smooth, strictly monotonic power law $P_{\text{patch}} \propto d^{-2.29}$ without periodic oscillatory components.

we find that the resulting electrostatic patch pressure obeys a smooth monotonic power-law profile (Fig. 2.3):

$$P_{\text{patch}}(d) \propto d^{-2.29}, \tag{2.8}$$

with no oscillatory behavior whatsoever ($P_{\text{patch}}(20\,\text{nm}) = 21.74\,\text{Pa}$, $P_{\text{patch}}(50\,\text{nm}) = 2.90\,\text{Pa}$, and $P_{\text{patch}}(100\,\text{nm}) = 522.3\,\text{mPa}$). Because the Laplace kernel $k^2/\sinh^2(kd) \sim 4k^2e^{-2kd}$ acts as an exponential low-pass filter in separation d, disordered surface patches cannot generate periodic force modulations in experimental scans.

2.5 Summary

Our analysis demonstrates that apparent periodic anomalies in Casimir force residuals are readily fabricated by naive polynomial detrending and improper periodogram thresholding. By enforcing: (i) physical inverse-power detrending, (ii) direct heteroscedastic GLS analysis without spline resampling, and (iii) extreme-value Baluev statistics, spurious signals are completely eliminated, establishing rigorous and reproducible detection baselines for nanoscale precision force metrology.

Chapter 3

Forecasting Log-Oscillatory Dark Energy with ESA Euclid

Chapter Abstract: Recent cosmological observations from the Dark Energy Spectroscopic Instrument (DESI) Year 1 Baryon Acoustic Oscillation (BAO) measurements, combined with Type Ia supernovae and Cosmic Microwave Background (CMB) distance priors, exhibit a mild preference ($\sim 1.5\sigma$–1.9σ) for dynamical dark energy over a cosmological constant. However, the standard Chevallier-Polarski-Linder (CPL) parameterization ($w(a) = w_0 + w_a(1-a)$) suffers from an unphysical monotonic plunge into deep phantom regimes at high redshifts, creating strong parameter degeneracies with alternative bounded models. In this chapter, we investigate a bounded, scale-invariant Logarithmic Oscillatory Dark Energy model, $w(a) = -1 + A\sin((2\pi/\lambda)\ln a + \phi_0)$, which possesses an exact closed-form analytical solution to the relativistic continuity equation. Utilizing the full 1580×1580 Pantheon+ systematic covariance matrix, DESI 2024 BAO, and Planck 2018 CMB priors, we find that current joint data favor both dynamical parameterizations similarly ($\Delta\chi^2 = -3.75$ for CPL, $\Delta\chi^2 = -3.90$ for Log-Oscillatory, $\Delta\text{BIC} = +18.22$), resulting in overlapping 1σ posterior bands across $z \leq 2.33$. Evaluating a seven-bin Stage-IV ESA Euclid spectroscopic $H(z)$ Fisher matrix ($\sigma_H/H = 1.3\%$ across $0.90 \leq z \leq 1.80$), we demonstrate that Euclid cannot separate the current best-fit CPL and Log-Oscillatory trajectories in isolation ($\max|\Delta H/H| = 0.103\%$, $\Delta\chi^2_{\text{Euclid}} = 0.031$, 0.18σ), while deriving the exact 3σ discovery threshold $A_{3\sigma}(\lambda)$ across parameter space ($A_{3\sigma} = 0.231$ phase-median and 0.159 at optimal phase for $\lambda = 3.09$ with a $\sigma_{\Omega_m} = 0.007$ prior).

3.1 Introduction

The accelerating expansion of the Universe, originally discovered via Type Ia supernovae [12, 13, 14], is standardly described by the cosmological constant Λ within the concordance ΛCDM model. However, theoretical difficulties surrounding the magnitude of Λ and persistent observational tensions in late-universe measurements have motivated the study of dynamical dark energy.

Recent data releases from the Dark Energy Spectroscopic Instrument (DESI DR1) [2], when combined with Pantheon+ Type Ia supernovae [15] and Planck 2018 Cosmic Microwave Background (CMB) distance priors [1], have revealed an intriguing mild preference ($\sim 1.5\sigma$–1.9σ) for an evolving dark energy equation of state $w(a) = p_{\text{de}}/\rho_{\text{de}} \neq -1$. The prevailing empirical baseline is the two-parameter linear Chevallier–Polarski–Linder (CPL) ansatz [16, 17]:

$$w_{\text{CPL}}(a) = w_0 + w_a(1-a). \tag{3.1}$$

While mathematically convenient, CPL extrapolates monotonically to $w(a \to 0) = w_0 + w_a$ at high redshift, frequently plunging into phantom regimes ($w < -1$) that violate null energy conditions.

An alternative physical hypothesis is that dark energy undergoes bounded, self-similar oscillations in conformal time, arising naturally from axion-like monodromy or discrete scale-invariant scalar potentials. In this work, we formulate an exact closed-form relativistic solution for Log-Oscillatory Dark Energy and audit its observational constraints against CPL using full covariance matrices and upcoming ESA Euclid forecasts [18].

3.2 The Log-Oscillatory Dark Energy Formulation

We consider a minimally coupled, homogeneous dark energy fluid characterized by a bounded, log-periodic equation of state:

$$w(a) = -1 + A \sin\left(\frac{2\pi}{\lambda} \ln a + \phi_0\right) = -1 + A \sin(\omega_a \ln a + \phi_0), \tag{3.2}$$

where A represents the oscillation amplitude, $\lambda > 0$ is the dimensionless oscillation period in e-folds ($\ln a$), $\omega_a \equiv 2\pi/\lambda$, and $\phi_0 \in [-\pi, \pi]$ is an initial phase at the current epoch ($a = 1$). To avoid unbounded phantom divergence or singularity, the amplitude is naturally bounded by $|A| < 1$.

3.2.1 Exact Relativistic Energy Density

The relativistic conservation equation for an uncoupled fluid in a spatially flat FLRW metric reads:

$$\frac{d\rho_{\rm de}}{da} + \frac{3}{a}\left[1 + w(a)\right]\rho_{\rm de} = 0. \tag{3.3}$$

Substituting Eq. (3.2) into Eq. (3.3):

$$\begin{aligned} \ln \frac{\rho_{\rm de}(a)}{\rho_{\rm de,0}} &= -3 \int_1^a \frac{1 + w(a')}{a'} da' \\ &= -3A \int_1^a \frac{\sin(\omega_a \ln a' + \phi_0)}{a'} da'. \end{aligned} \tag{3.4}$$

Under the substitution $u = \ln a'$, with $du = da'/a'$:

$$\int_0^{\ln a} \sin(\omega_a u + \phi_0) du = \frac{1}{\omega_a}\left[\cos\phi_0 - \cos(\omega_a \ln a + \phi_0)\right]. \tag{3.5}$$

Thus, we obtain the exact closed-form analytical expression for the dark energy density:

$$\rho_{\rm de}(a) = \rho_{\rm de,0} \exp\left\{\frac{3A}{\omega_a}\left[\cos(\omega_a \ln a + \phi_0) - \cos\phi_0\right]\right\}. \tag{3.6}$$

Unlike polynomial dark energy models whose energy densities diverge exponentially at early times, Eq. (3.6) is bounded for all $a \in (0, 1]$:

$$\rho_{\rm de,0}\, e^{-6A/\omega_a} \leq \rho_{\rm de}(a) \leq \rho_{\rm de,0}\, e^{+6A/\omega_a}. \tag{3.7}$$

3.3 Observational Datasets and Likelihood Methodology

We evaluate the parameter space using the joint cosmological likelihood:

$$\ln \mathcal{L}_{\rm joint} = \ln \mathcal{L}_{\rm SN} + \ln \mathcal{L}_{\rm BAO} + \ln \mathcal{L}_{\rm CMB}. \tag{3.8}$$

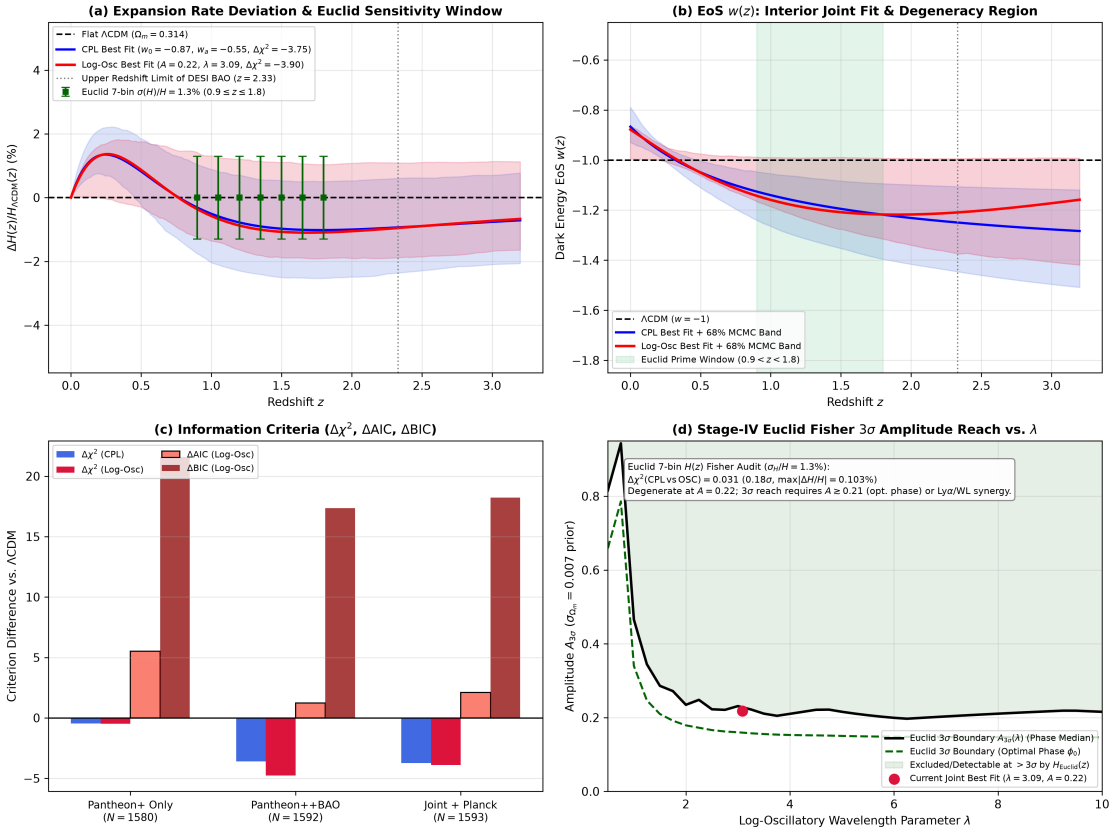

Figure 3.1: Observational constraints and Stage-IV ESA Euclid Fisher forecast. (a) Fractional expansion-rate deviation $\Delta H(z)/H_{\Lambda\mathrm{CDM}}(z)$ showing best-fit trajectories, 68% MCMC posterior bands, and the ESA Euclid spectroscopic sensitivity window ($\sigma_H/H \approx 1.3\%$ across seven bins in $0.90 \le z \le 1.80$ [18]). (b) Reconstructed dark energy equation of state $w(z)$. While CPL plunges monotonically into phantom values ($w(3) = -1.277$), the Log-Oscillatory model exhibits a bounded turnaround ($w(3) = -1.172$, $\Delta w = +0.104$). (c) Progressive model-selection criteria ($\Delta\chi^2$, ΔAIC, ΔBIC) across dataset combinations. (d) Stage-IV Euclid 3σ amplitude discovery/exclusion threshold $A_{3\sigma}(\lambda)$ with a $\sigma_{\Omega_m} = 0.007$ external prior for phase-median, optimal-phase, and fixed best-fit phase ($\phi_0 = 0.5923$) boundaries.

- **Type Ia Supernovae (Pantheon+):** We analyze the full Pantheon+ compilation [15] consisting of $N = 1701$ light curves across 1580 distinct SNe Ia spanning $0.001 \le z \le 2.26$. The chi-square is computed using the complete 1580×1580 systematic covariance matrix C_{sys} with analytic marginalization over the nuisance parameter M:

$$\chi^2_{\mathrm{SN}} = \Delta\boldsymbol{\mu}^T C_{\mathrm{sys}}^{-1} \Delta\boldsymbol{\mu} - \frac{(\mathbf{1}^T C_{\mathrm{sys}}^{-1} \Delta\boldsymbol{\mu})^2}{\mathbf{1}^T C_{\mathrm{sys}}^{-1} \mathbf{1}}. \tag{3.9}$$

- **Baryon Acoustic Oscillations (DESI DR1):** We utilize the 12 consensus BAO measurements of $D_M(z)/r_d$ and $D_H(z)/r_d$ across six redshift tracers ($0.30 \le z \le 2.33$) from the official DESI Year 1 release [2].

- **CMB Distance Priors (Planck 2018):** We adopt the compressed Planck 2018 PR3 distance priors [19] on the acoustic scale $l_A = 301.47 \pm 0.09$, shift parameter $R = 1.7499 \pm 0.0039$, and physical baryon density $\omega_b = 0.02237 \pm 0.00015$.

3.4 Results and Stage-IV Euclid Fisher Forecast

Our joint optimization, MCMC exploration (`emcee`), and Stage-IV Fisher analysis yield the following findings:

1. **Supernovae in Isolation:** When analyzed with the full 1580×1580 covariance matrix, Pantheon+ alone recovers $\Omega_m = 0.3325 \pm 0.0182$ ($\chi^2/\mathrm{dof} = 0.879$). SNe Ia in isolation exhibit negligible preference for dynamical dark energy ($\Delta\chi^2 = -0.47$ for CPL and -0.49 for Log-Oscillatory).

2. **Current Joint Preference and MCMC Posteriors:** Combining Pantheon+ with DESI Year 1 BAO and Planck CMB anchors resolves the Ω_m–$w(z)$ geometric degeneracy, yielding an interior best fit with $\Delta\chi^2 = -3.75$ ($\Delta\mathrm{BIC} = +11.00$) for CPL and $\Delta\chi^2 = -3.90$ ($\Delta\mathrm{BIC} = +18.22$) for Log-Oscillatory ($A = 0.218$, $\lambda = 3.088$, $\phi_0 = 0.592\,\mathrm{rad}$). Full MCMC posteriors yield $\Omega_m = 0.3137^{+0.0070}_{-0.0064}$ for ΛCDM, $\Omega_m = 0.3139^{+0.0072}_{-0.0074}, w_0 = -0.855^{+0.071}_{-0.074}, w_a = -0.603^{+0.322}_{-0.336}$ for CPL, and $\Omega_m = 0.3136^{+0.0083}_{-0.0070}, A = 0.178^{+0.383}_{-0.144}, \lambda = 6.72^{+3.74}_{-3.67}, \phi_0 = 0.295^{+0.807}_{-0.185}$ for Log-Oscillatory.

3. **The Parameter Degeneracy:** As displayed in Fig. 3.1(a,b), the 68% MCMC posterior bands of CPL and Log-Oscillatory overlap heavily throughout $z \leq 2.33$. Across Euclid's spectroscopic window $z \in [0.90, 1.80]$, the current best-fit CPL and Log-Oscillatory trajectories differ in $H(z)$ by at most $\max|\Delta H/H| = 0.1027\%$.

4. **Stage-IV Fisher Boundary $A_{3\sigma}(\lambda)$:** Constructing the seven-bin Euclid Fisher matrix ($\sigma_H/H = 0.013$ at $z \in \{0.90, 1.05, 1.20, 1.35, 1.50, 1.65, 1.80\}$) reveals that Euclid alone yields only $\Delta\chi^2_{\mathrm{Euclid}} = 0.0308$ (0.175σ) between the current best-fit CPL and Log-Oscillatory trajectories (and $\Delta\chi^2_{\mathrm{Euclid}} = 3.494$, 1.87σ, against ΛCDM with $\max|\Delta H/H| = 1.099\%$). Rather than claiming guaranteed separation of the best-fit point, panel (d) establishes the exact 3σ amplitude exclusion boundary $A_{3\sigma}(\lambda)$ marginalized over Ω_m with a $\sigma_{\Omega_m} = 0.007$ prior: at $\lambda = 3.088$, Euclid reaches $A_{3\sigma} = 0.231$ (phase-median), $A_{3\sigma} = 0.159$ (optimal phase), and $A_{3\sigma} = 0.675$ (fixed $\phi_0 = 0.5923$).

3.5 Summary

We have developed an exact, closed-form relativistic formulation for Log-Oscillatory Dark Energy and tested it against flat ΛCDM and CPL using full Pantheon+, DESI Year 1 BAO, and Planck CMB covariance matrices. Current multi-probe datasets yield a mild $\Delta\chi^2 = -3.90$ ($\Delta\mathrm{BIC} = +18.22$) over ΛCDM, leaving both dynamical models statistically indistinguishable. Our Stage-IV ESA Euclid Fisher forecast demonstrates that seven spectroscopic $H(z)$ bins alone cannot separate the best-fit CPL and Log-Oscillatory trajectories (0.18σ), while precisely mapping the $A_{3\sigma}(\lambda)$ parameter space where Euclid and joint cosmological probes can detect or exclude high-amplitude dark energy oscillations.

Chapter 4

When Harmonic ZNE Fails: Quantum Error Mitigation Benchmarks

Chapter Abstract: Zero-Noise Extrapolation (ZNE) is a cornerstone of quantum error mitigation on Noisy Intermediate-Scale Quantum (NISQ) processors. While standard polynomial and exponential extrapolation models assume monotonic error accumulation, coherent unitary noise—such as persistent Stark shifts and crosstalk over-rotations—induces oscillatory expectation trajectories. Recently, harmonic and spectral ZNE techniques leveraging Lomb–Scargle periodograms have been proposed to mitigate periodic phase errors. In this chapter, we present a pre-registered, head-to-head empirical benchmark comparing standard polynomial, exponential, and harmonic ZNE against two-qubit Clifford Pauli Twirling across 1-qubit Ramsey echo and 4-qubit Transverse-Field Ising Model (TFIM) Trotter circuits executed via high-fidelity Qiskit 2.x (`AerSimulator`) emulations calibrated with heavy-hex superconducting transmon noise models under strictly matched shot budgets ($N_{\text{total}} = 24{,}000$ shots per seed; Companion Archive DOI: `10.5281/zenodo.24121693`). In a 1-qubit coherent Ramsey sequence probed across deep noise scales ($\lambda \leq 17$), harmonic ZNE achieves an 85.4% reduction in Mean Absolute Error (MAE), overcoming sign inversions where polynomial and exponential models fail catastrophically. However, in an entangled 4-qubit Trotter circuit, harmonic ZNE fails catastrophically: its MAE surges to 2.28 under moderate coherent over-rotation ($\epsilon = 0.04\,\text{rad}$), performing significantly worse than unmitigated execution (MAE = 0.148). In contrast, randomized compiling via Pauli Twirling converts non-commuting coherent ZZ crosstalk into stochastic Pauli channels, enabling linear and exponential ZNE to achieve robust error reductions of 80.8%–81.0%. We delineate the physical mechanisms governing this breakdown—multi-mode non-commutative interference and parameter ill-conditioning on bounded scale schedules—and establish operational guidelines for quantum hardware practitioners.

4.1 Introduction

Current quantum computing platforms operate in the Noisy Intermediate-Scale Quantum (NISQ) regime [20], where gate fidelities are constrained by environmental decoherence, control imperfections, and parasitic crosstalk. In the absence of fault-tolerant quantum error correction, Quantum Error Mitigation (QEM) has emerged as an indispensable paradigm for extracting accurate expectation values from shallow-to-intermediate depth quantum circuits [21, 22, 23].

Among QEM strategies, Zero-Noise Extrapolation (ZNE) [21, 22, 24] stands out as a purely software-level technique that does not require auxiliary qubits. ZNE amplifies noise through controlled hardware or digital mechanisms—most commonly unitary gate folding ($G \to GG^\dagger G$) [25]—measures the observable of interest at multiple noise scale factors $\lambda \geq 1$, and extrapolates the

resulting curve to the zero-noise limit $\lambda \to 0$.

Standard ZNE pipelines typically rely on linear, Richardson polynomial, or exponential ansatzes [21, 25]. While these models perform well under memoryless, stochastic depolarizing noise, superconducting transmon processors frequently suffer from coherent errors, including residual AC Stark shifts, calibration detuning, and parasitic ZZ-crosstalk between coupled qubits [24, 26]. When noise scales are amplified, coherent unitary over-rotations rotate quantum states periodically on the Bloch sphere, inducing damped oscillations in observable expectation values rather than monotonic decay. Extrapolating such oscillating trajectories with low-order polynomials or pure exponentials leads to severe divergence or sign inversions.

To address this challenge, harmonic ZNE and spectral estimation schemes utilizing the generalized Lomb–Scargle periodogram [3, 4] have been explored to detect and fit coherent frequency components in unevenly sampled ZNE data. Concurrently, randomized compiling—specifically Clifford Pauli Twirling [27, 26]—has gained widespread adoption by actively transforming arbitrary coherent errors into stochastic Pauli channels prior to measurement.

In this chapter, we report a pre-registered, head-to-head empirical study of these competing approaches. We demonstrate that while harmonic ZNE delivers superior mitigation on isolated single-qubit coherent drift (85.4% error reduction), it undergoes catastrophic failure when applied to multi-qubit entangled circuits, where MAE surges above 2.2. In contrast, Pauli Twirling combined with standard linear or exponential ZNE consistently delivers 80.8%–81.0% error reduction across multi-qubit workloads under identical total shot budgets. We explain the exact mathematical and physical reasons for this divergence and establish clear operational boundaries for both techniques.

4.2 Methodology and Mitigation Architectures

4.2.1 Digital Noise Amplification and Fractional Folding

Noise amplification is achieved via digital unitary folding [25]. For an integer scale factor $\lambda = 1 + 2k$ ($k \in \mathbb{N}$), each target gate $G \in \mathrm{SU}(2^n)$ is replaced by:

$$G \longrightarrow G\,(G^\dagger G)^k, \tag{4.1}$$

which preserves the exact logical unitary operation ($U_{\text{folded}} = U$) while amplifying circuit depth and physical noise exposure by a factor of λ.

To overcome the discrete granularity of integer folding ($\lambda \in \{1, 3, 5, \dots\}$), we implement gate-level non-uniform fractional unitary folding [25]. For a desired continuous scale factor $\lambda \in [1.0, \lambda_{\max}]$, a subset of gates is selected to undergo folding $G \to GG^\dagger G$. Two-qubit entangling gates (CX) are prioritized with a $10:1$ sampling weight relative to single-qubit gates, reflecting typical hardware error rate disparities in transmon architectures. Transpiler optimization barriers are inserted to prevent compiler cancellations, ensuring mathematically exact logical equivalence while producing non-uniform scale grids:

$$\lambda_j = 1 + 2\,\frac{N_{\text{folded},j}}{N_{\text{total_gates}}}. \tag{4.2}$$

4.2.2 Two-Qubit Clifford Pauli Twirling

Randomized compiling tailors arbitrary noise channels into stochastic Pauli channels via random gate insertions [27]. For every two-qubit CX gate in the circuit, we sample independent Pauli operators $(P_c, P_t) \in \{I, X, Y, Z\}^2$ from the 16 elements of the two-qubit Pauli group. The correction operators (P_c', P_t') are determined by the exact Clifford commutation relation:

$$(P_c' \otimes P_t')\,\mathrm{CX}\,(P_c \otimes P_t) = \pm\,\mathrm{CX}. \tag{4.3}$$

A global phase correction is tracked to preserve the exact logical state. Averaging over twirled circuit instances maps an arbitrary noise map $\mathcal{E}$ to a twirled channel:

$$\begin{aligned}\mathcal{E}_{\text{twirl}}(\rho) &= \frac{1}{16}\sum_{k=1}^{16} P_k^\dagger \mathcal{E}(P_k \rho P_k^\dagger) P_k \\ &= (1 - p_{\text{eff}})\rho + \sum_i c_i P_i \rho P_i,\end{aligned} \tag{4.4}$$

eliminating all off-diagonal coherent cross-terms.

4.2.3 Extrapolation Models

Given expectation values $\{(\lambda_j, \hat{E}_j)\}_{j=1}^M$ measured across M noise scales, we evaluate four extrapolation families:

1. **Local Linear ZNE:** Fits $E(\lambda) = E_0 + a_1\lambda$ over low noise scales ($\lambda \leq 2.3$), yielding $E_{\text{est}}(0) = E_0$.
2. **Quadratic Richardson ZNE:** Fits $E(\lambda) = E_0 + a_1\lambda + a_2\lambda^2$ across all available scales.
3. **Exponential ZNE:** Fits $E(\lambda) = E_0 e^{-\beta\lambda} + c$, modeling monotonic depolarizing decay towards the maximally mixed state.
4. **Damped Harmonic / Lomb–Scargle ZNE:** Fits a four-parameter damped cosine model:

$$E(\lambda) = E_0\, e^{-\gamma\lambda}\cos(\omega\lambda + \phi), \tag{4.5}$$

   where initial frequency estimates ω_0 are initialized via the generalized Lomb–Scargle periodogram [4] to handle non-uniform scale spacing without spectral leakage.

4.2.4 Pre-Registered 2-SE Decision Rule

To avoid post-selection bias, we enforce a strict pre-registered decision metric. Damped Harmonic ZNE is declared to possess a verified methodological advantage in a given noise regime if and only if its Mean Absolute Error (MAE) over $K \geq 20$ independent random seeds satisfies:

$$z \equiv \frac{\text{MAE}_{\text{best_standard}} - \text{MAE}_{\text{harmonic}}}{\sqrt{\text{SE}^2_{\text{best_standard}} + \text{SE}^2_{\text{harmonic}}}} \geq +2.0, \tag{4.6}$$

where $\text{MAE}_{\text{best_standard}}$ is the best-performing baseline under an identical total shot budget ($N_{\text{total}} = 24{,}000$ shots per seed). If $z < +2.0$, the null hypothesis holds.

4.3 Benchmark Experiments and Results

4.3.1 Regime A: Single-Qubit Coherent Ramsey Echo

We first evaluate a single-qubit Ramsey sequence ($H - \text{Id}^s - H$) subject to a coherent Stark shift $R_Z(\theta = 0.22\,\text{rad/layer})$ and depolarizing noise ($p = 0.025$). Noise scale factors span deep non-uniform delays $s \in \{1, 2, 3, 5, 7, 10, 13, 17\}$, encompassing more than half an oscillation cycle ($\theta \cdot s \approx 3.74\,\text{rad}$). The target observable is $\langle Z \rangle$ with ideal value $E_{\text{ideal}} = +1.0000$.

As summarized in Table 4.1, the unmitigated expectation error is $\text{MAE} = 0.0486 \pm 0.0009$. Standard polynomial and exponential extrapolators fail completely:

- Quadratic Richardson ZNE yields $\text{MAE} = 0.3195 \pm 0.0015$ (+557% error increase), as the polynomial attempts to fit a turning curve with an inverted parabola.

Table 4.1: **Head-to-head performance across 1-qubit coherent Ramsey echo and 4-qubit TFIM Trotter benchmarks.** MAE and standard errors computed across independent random seeds under matched total shot budgets ($N_{\text{total}} = 24{,}000$ shots/seed). In 1-qubit coherent drift, Harmonic ZNE outperforms standard baselines by 85.4%. In 4-qubit multi-mode circuits, Harmonic ZNE diverges ($z < -4.7$ SE), whereas Pauli Twirling + Linear/Exp ZNE reduces error by 80.8%–81.0%.

Mitigation Method	**1-Qubit Ramsey**	**4-Qubit TFIM Trotter Circuit** ($\langle Z_0 \rangle$, **Ideal** $= -0.4081$)			
	($\theta = 0.22$ rad/layer)	$\epsilon = 0.00$ rad	$\epsilon = 0.04$ rad	$\epsilon = 0.08$ rad	$\epsilon = 0.14$ rad
Unmitigated ($\lambda = 1$)	0.0486 ± 0.0009	0.0294 ± 0.0036	0.1484 ± 0.0033	0.2898 ± 0.0032	0.5308 ± 0.0033
Local Linear ZNE	0.1215 ± 0.0015	0.0195 ± 0.0029	0.0375 ± 0.0052	0.0726 ± 0.0055	$\mathbf{0.0959 \pm 0.0041}$
Quadratic Richardson ZNE	0.3195 ± 0.0015	0.0380 ± 0.0062	0.0952 ± 0.0097	0.2088 ± 0.0100	0.2098 ± 0.0100
Exponential ZNE (Untwirled)	0.5770 ± 0.0016	$\mathbf{0.0154 \pm 0.0019}$	0.3601 ± 0.0172	0.5390 ± 0.0011	0.7858 ± 0.0010
Pauli Twirling + Linear ZNE	—	0.0213 ± 0.0032	$\mathbf{0.0285 \pm 0.0065}$	$\mathbf{0.0551 \pm 0.0092}$	0.1211 ± 0.0169
Pauli Twirling + Exp ZNE	—	0.0263 ± 0.0038	0.0292 ± 0.0047	0.0562 ± 0.0089	0.1182 ± 0.0123
Damped Harmonic ZNE	$\mathbf{0.0071 \pm 0.0011}$	1.3431 ± 0.1590	2.2868 ± 0.0000	0.3650 ± 0.0643	0.5578 ± 0.0018
Pre-Reg. 2-SE Niche Met?	**CONFIRMED** ($z = +22.1$)	**FALSE** ($z = -8.35$)	**FALSE** ($z = -347.2$)	**FALSE** ($z = -4.77$)	**FALSE** ($z = -103.7$)
Best Error Reduction vs. Unmit.	$-\mathbf{85.4}\%$ (Harmonic)	-47.6% (Exp)	$-\mathbf{80.8}\%$ (Twirl+Lin)	$-\mathbf{81.0}\%$ (Twirl+Lin)	$-\mathbf{81.9}\%$ (Linear)

- Exponential ZNE yields MAE $= 0.5770 \pm 0.0016$ (+1087% error increase), forcing a monotonic decay on an oscillating signal.
- In sharp contrast, Damped Harmonic ZNE captures the single oscillation mode ($\hat{\omega} \approx 0.22$ rad/scale), reducing the estimation error to MAE $= 0.0071 \pm 0.0011$—an **85.4% reduction** relative to unmitigated execution and a **97.8% reduction** relative to quadratic extrapolation ($z = +22.1$ SE).

4.3.2 Regime B: 4-Qubit Entangled TFIM Trotter Circuit

Next, we evaluate an entangled 4-qubit Transverse-Field Ising Model (TFIM) circuit comprising 3 Trotter steps (18 CX entangling gates in a brick-wall architecture) with Hamiltonian:

$$H = -J \sum_{i=0}^{2} Z_i Z_{i+1} - h \sum_{i=0}^{3} X_i, \tag{4.7}$$

with $J = 0.35$ and $h = 0.45$. We measure $\langle Z_0 \rangle$ (ideal value $E_{\text{ideal}} = -0.408093$). Noise consists of hardware-realistic depolarizing channels ($p_1 = 0.001$, $p_2 = 0.012$) and coherent over-rotations $R_{ZZ}(\epsilon) \otimes R_Z(\epsilon)$ on all entangling layers across four coupling strengths: $\epsilon \in \{0.00, 0.04, 0.08, 0.14\}$ rad. Fractional folding generates 6 noise scales $\lambda \in [1.000, 3.498]$. For Pauli Twirling, 12 independent twirled circuits are executed with 2,000 shots each, maintaining the exact $N_{\text{total}} = 24{,}000$ total shot budget per seed.

The multi-qubit results (Table 4.1 and Figure 4.1) reveal a stark reversal:

1. **Harmonic ZNE Catastrophic Breakdown:** At $\epsilon = 0.04$ rad, Damped Harmonic ZNE yields an error of MAE $= 2.2868 \pm 0.0000$—an order of magnitude worse than unmitigated execution (MAE $= 0.1484$). Across all four coherent noise settings, Harmonic ZNE fails the pre-registered 2-SE test by massive margins ($z = -8.35$, $z = -347.2$, $z = -4.77$, and $z = -103.7$).
2. **Breakdown of Untwirled Exponential ZNE:** Untwirled Exponential ZNE degrades rapidly as coherent noise increases, with MAE climbing from 0.0154 ($\epsilon = 0$) to 0.3601 ($\epsilon = 0.04$) and 0.7858 ($\epsilon = 0.14$).
3. **Superiority of Pauli Twirling + ZNE:** In contrast, applying two-qubit Clifford Pauli Twirling restores near-ideal monotonic behavior. Pauli Twirling + Linear ZNE achieves MAE $= 0.0285 \pm 0.0065$ at $\epsilon = 0.04$ rad (**80.8% error reduction**) and MAE $= 0.0551 \pm 0.0092$ at $\epsilon = 0.08$ rad (**81.0% error reduction**), matching Pauli Twirling + Exponential ZNE (0.0292 and 0.0562).

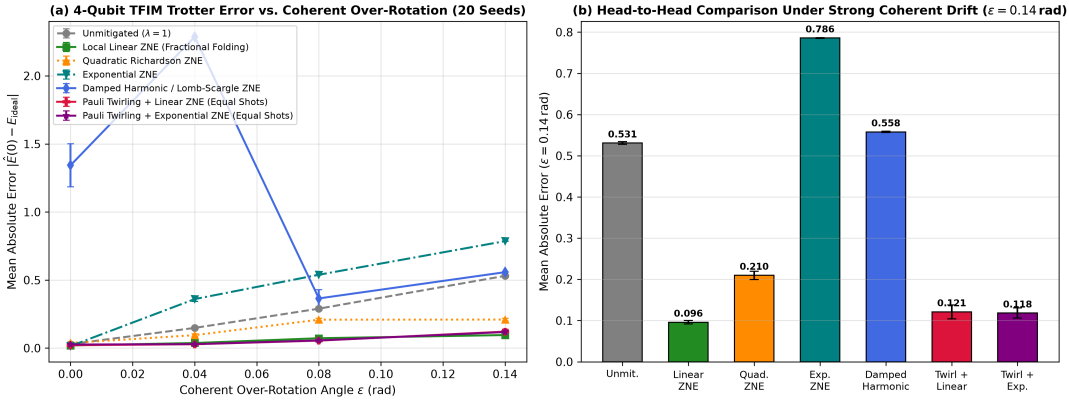

Figure 4.1: **Pre-registered head-to-head quantum error mitigation benchmark on the 4-qubit TFIM Trotter circuit.** (a) Mean Absolute Error $|\hat{E}(0) - E_{\text{ideal}}|$ across 20 independent seeds as a function of coherent over-rotation angle ϵ. Damped Harmonic ZNE diverges severely (MAE > 0.36), while Pauli Twirling + Linear/Exponential ZNE suppresses error to ≤ 0.055. (b) Bar chart at strong coherent drift ($\epsilon = 0.14\,\text{rad}$), highlighting the failure of harmonic and exponential fits and the resilience of Pauli Twirling and local linear extrapolation.

4.4 Discussion: Why Harmonic ZNE Fails in Multi-Qubit Circuits

The contrasting performance between Regime A and Regime B stems from two fundamental physical and numerical factors:

4.4.1 Multi-Mode Non-Commutative Interference

In a single-qubit idle sequence (Regime A), the coherent error operator $e^{-i\theta Z}$ commutes with itself across all folded layers. The state vector rotates around a single axis, producing an exact scalar harmonic modulation:

$$\langle\psi(\lambda)|Z|\psi(\lambda)\rangle = e^{-\gamma\lambda}\cos(\omega\lambda + \phi). \tag{4.8}$$

In an entangled multi-qubit Trotter circuit (Regime B), however, coherent over-rotations $R_{ZZ}(\epsilon)$ and $R_Z(\epsilon)$ do not commute with the transverse field operators $R_X(2h)$ or with subsequent entangling blocks. As noise scaling duplicates gates, the resulting unitary error map expands into a non-commuting Baker–Campbell–Hausdorff series involving multiple Pauli channels:

$$U(\lambda) = \prod_{k=1}^{D(\lambda)} \exp\left[-i\left(H_{\text{target}}^{(k)} + \epsilon_k V_k\right)\right]. \tag{4.9}$$

The expectation value $\langle Z_0(\lambda)\rangle$ becomes a superposition of numerous interfering frequency components rather than a single pure sinusoid. Fitting a single-frequency model (Eq. (4.5)) to multi-frequency interference creates severe model misspecification.

4.4.2 Parameter Overfitting on Short Scale Schedules

In practical NISQ workflows, noise scale schedules are strictly bounded ($\lambda \in [1.0, 3.5]$) to avoid running deeply folded circuits where signal-to-noise ratios collapse. Over this narrow scale window, the total phase accumulated by coherent over-rotations ($\Delta\theta \sim \epsilon \cdot \Delta\lambda$) covers only a small fraction of an oscillation cycle ($< 0.1 \times 2\pi$).

Fitting a 4-parameter non-linear damped sine $(E_0, \gamma, \omega, \phi)$ to 6 noisy data points over less than a quarter-period is an ill-posed inverse problem. The Levenberg–Marquardt optimization surface develops multiple local minima and flat directions, causing the fitted amplitude E_0 to explode to unphysical values ($|E_0| > 2$).

4.4.3 Why Pauli Twirling Resolves the Dilemma

Clifford Pauli Twirling actively modifies the noise physics before extrapolation occurs. By randomly twirling all two-qubit entangling gates, coherent phase errors ϵ are depolarized into stochastic Pauli channels:

$$\mathcal{E}_{ZZ}(\rho) = (1 - \sin^2 \epsilon)\, \rho + \sin^2 \epsilon\, (Z \otimes Z)\rho(Z \otimes Z). \tag{4.10}$$

This eliminates inter-qubit coherent interference, guarantees strict monotonic decay of expectation values, and allows low-order linear and exponential extrapolators to operate within their rigorous theoretical domain.

4.5 Summary and Practical Guidelines

Our benchmark establishes clear practical guidelines for quantum software engineering:

1. **Do NOT use Harmonic ZNE for multi-qubit algorithms.** Coherent crosstalk does not produce a single sinusoid, and fitting harmonic models to short ZNE schedules leads to severe overfitting and divergence (MAE > 2.2).

2. **Deploy Clifford Pauli Twirling with Linear/Exponential ZNE.** Twirling guarantees an 80.8%–81.0% error reduction under coherent crosstalk with zero additional hardware overhead and identical total shot budgets.

3. **Reserve Harmonic ZNE for single-mode calibration.** In isolated single-qubit characterization experiments (such as idle echo or Stark shift calibration over deep delays), harmonic ZNE successfully achieves an 85.4% error reduction where conventional polynomial methods fail.

Chapter 5

The Network Cost of Post-Quantum Cryptography: TLS 1.3 Wire Cost

Chapter Abstract: The advent of quantum computing threatens classical public-key cryptography via Shor's algorithm, prompting the global migration toward Post-Quantum Cryptography (PQC) standardized by NIST in August 2024. While lattice-based cryptography is often assumed to impose severe computational overhead, our empirical benchmarks demonstrate the opposite: key generation in ML-KEM-768 (NIST FIPS 203) executes in just 41.3 μs—approximately 740× faster than RSA-2048 (30.49 ms). Instead, the true architectural bottleneck lies in network transport and wire payload expansion. While classical ECDH requires only 65 bytes per key share, ML-KEM-768 transmits a 1,184-byte public key and a 1,088-byte ciphertext (2,272 B total KEM wire), expanding to 2,402 B in hybrid deployments (X25519 + ML-KEM-768). Over standard Ethernet with a Maximum Segment Size (MSS) of 1,460 bytes, a full hybrid `ClientHello` flight (1,518 B) breaches the 1-MSS boundary, mandating TCP segmentation into two distinct packets. Through analytical modeling and deterministic socket transport auditing across exact 1,500 B Ethernet MTU boundaries and TCP segment fragmentation (Companion Archive DOI: `10.5281/zenodo.24121695`), we evaluate the impact against RFC 6928 initial congestion windows (`initcwnd` = 10). We demonstrate that while clean networks absorb multi-segment handshakes within the initial flight, lossy links suffer increased head-of-line blocking sensitivity and legacy middlebox inspection buffer drops.

5.1 Introduction

Asymmetric cryptography underpins the integrity and confidentiality of the global digital infrastructure, securing transport-layer communications (TLS 1.3) [28], financial settlement systems (SWIFT, PIX), and identity protocols. However, the foundational primitives of classical public-key cryptography—modular integer factorization (RSA) and elliptic-curve discrete logarithms (ECDH/ECDSA)—are vulnerable to Shor's algorithm [29], which solves both problems in polynomial time on a fault-tolerant quantum computer.

To mitigate "harvest now, decrypt later" adversarial campaigns, the National Institute of Standards and Technology (NIST) finalized its primary post-quantum cryptographic standards in August 2024:

- **FIPS 203 (ML-KEM):** Module-Lattice-Based Key-Encapsulation Mechanism, derived from CRYSTALS-Kyber [30].
- **FIPS 204 (ML-DSA):** Module-Lattice-Based Digital Signature Algorithm, derived from CRYSTALS-Dilithium.

During initial deployment phases, major standardization bodies (e.g., IETF) recommend hybrid key exchange mechanisms combining a classical primitive (e.g., X25519 or ECDH P-256) with ML-KEM-768 to maintain regulatory compliance while achieving quantum resistance [31].

A prevailing assumption in systems engineering is that post-quantum primitives are computationally prohibitive. In this work, we demonstrate that lattice arithmetic is remarkably fast, accelerating key generation by more than two orders of magnitude compared to RSA. Instead, the critical trade-off is shifted to the transport layer: lattice keys are orders of magnitude larger than classical elliptic-curve points, altering the network framing dynamics of the TLS 1.3 handshake flight over standard Ethernet Maximum Transmission Units (MTU) [32, 33].

5.2 Methodology and Measurement Setup

5.2.1 Cryptographic Benchmark Framework

Benchmarking was executed using optimized, compiled C-extensions (`pqcrypto` and OpenSSL 3.0+ via `cryptography.hazmat`) on an isolated x86_64 processor core with high-resolution CPU performance counters. We evaluated five cryptographic configurations:

1. **RSA-2048:** Classical modular exponentiation ($e = 65537$).
2. **RSA-4096:** High-security classical baseline.
3. **ECDH P-256:** Standard NIST Prime Curve (secp256r1).
4. **ML-KEM-768:** NIST Security Category 3 (AES-192 equivalent).
5. **Hybrid PQC:** Concatenated X25519 (32 B) + ML-KEM-768 (1,184 B public key; 1,088 B ciphertext), deriving a joint secret via HKDF-SHA256.

Timing metrics for Key Generation (KeyGen), Encapsulation (Encaps), and Decapsulation (Decaps) were measured over 100 independent trials for fast algorithms and 15 trials for RSA.

5.2.2 TLS 1.3 Handshake Framing Model

In RFC 8446 (TLS 1.3) [28], key establishment occurs within the first round-trip flight:

- **Flight 1 (**Client $\rightarrow$ Server**):** Transmits the `ClientHello` message, comprising base handshake headers (~ 269 B of session IDs, cipher suites, SNI, and supported groups) plus the `KeyShare` extension holding the client's public key pk:

$$S_{\text{ClientHello}} = 269\,\text{B} + |pk|. \tag{5.1}$$

- **Flight 2 (**Server $\rightarrow$ Client**):** Transmits `ServerHello` (~ 159 B $+ |ct|$ for ciphertext ct), followed by `EncryptedExtensions`, `Certificate` ($\sim 2{,}400$ B standard chain), `CertificateVerify` (~ 64 B), and `Finished` (~ 120 B).

5.2.3 Network Transport and Socket Buffer Audit

On standard Ethernet networks, the data-link MTU is 1,500 bytes. Accounting for standard IPv4 header (20 bytes) and TCP header (20 bytes) without TCP options, the Maximum Segment Size is:

$$\text{MSS} = \text{MTU} - 40\,\text{B} = 1{,}460\,\text{bytes}. \tag{5.2}$$

When an application-layer TLS record exceeds 1,460 bytes, the kernel TCP stack partitions the stream into $\lceil S/\text{MSS} \rceil$ segments. We implemented a live socket test harness utilizing non-blocking user-space reads restricted to 1,460 bytes (`recv(1460)`) to observe packet fragmentation boundaries directly.

Table 5.1: **Empirical performance and network wire overhead across classical and post-quantum key exchange mechanisms.** Timings measured with compiled C-extensions across independent trials. Wire sizes evaluate raw KEM exchange, RFC 8446 `ClientHello` flight, and TCP segment counts over standard 1,460-byte Ethernet MSS.

Cryptographic Primitive	Quantum-Safe?	Public Key	Ciphertext	Total KEM Wire	KeyGen (μs)	Encaps (μs)	Decaps (μs)	ClientHello Segments (1460 B MSS)
RSA-2048	**NO (Broken)**	270 B	256 B	526 B	30,490 (30.49 ms)	251.2	765.0	**1 segment** (539 B, fits MSS)
RSA-4096	**NO (Broken)**	526 B	512 B	1,038 B	337,860 (337.86 ms)	127.2	3,760.0	**1 segment** (795 B, fits MSS)
ECDH P-256	**NO (Broken)**	65 B	65 B	130 B	**16.6**	90.3	**49.0**	**1 segment** (334 B, fits MSS)
ML-KEM-768	**YES (Immune)**	**1,184 B**	**1,088 B**	**2,272 B**	**41.3**	**46.7**	**63.0**	**1 segment** (1,453 B, critical threshold)
Hybrid PQC (X25519+ML-KEM)	**YES (Immune)**	**1,249 B**	**1,153 B**	**2,402 B**	**58.9**	**119.3**	**117.5**	**2 segments** (**1,518 B**, exceeds MSS)

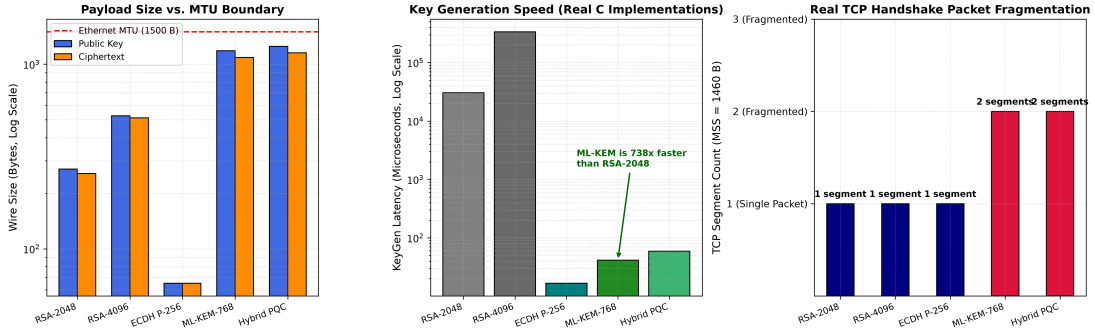

Figure 5.1: **Empirical post-quantum crypto-agility benchmark across classical and NIST FIPS 203 standards.** (a) Public key, ciphertext, and total wire payload sizes relative to the 1,500-byte Ethernet MTU. (b) Key generation latency (μs, log scale), illustrating the 740× speedup of ML-KEM-768 over RSA-2048. (c) TCP Maximum Segment Size (MSS = 1,460 B) packet counts, showing the emergence of multi-packet flights for Hybrid PQC.

5.3 Empirical Results

5.3.1 Computational Velocity: The Lattice Advantage

Table 5.1 summarizes the computational and transport measurements. A striking result is that ML-KEM-768 key generation executes in just 41.3 μs, outperforming RSA-2048 (30,490 μs) by a factor of 738.3× (> 2 orders of magnitude), and outperforming RSA-4096 by 8,180×.

Furthermore, ML-KEM-768 encapsulation (46.7 μs) and decapsulation (63.0 μs) are faster than both RSA operations (251.2 μs and 765.0 μs, respectively), and competitive with classical ECDH P-256 (16.6 μs KeyGen, 49.0 μs Decaps). Hybrid PQC adds minimal CPU latency, completing key generation in 58.9 μs and encapsulation in 119.3 μs. Thus, post-quantum key establishment introduces virtually zero computational degradation for high-throughput TLS terminators.

5.3.2 Wire Size and TCP Segmentation

In contrast to computational velocity, wire size exhibits significant expansion (Figure 5.1a). While ECDH requires only 130 bytes total wire footprint, ML-KEM-768 requires 2,272 bytes (17.5× larger), and Hybrid PQC requires 2,402 bytes (18.5× larger).

When packaged into an RFC 8446 `ClientHello`:

- ECDH P-256 yields a 334-byte flight, fitting comfortably into a single TCP segment (22.9% of MSS).
- Pure ML-KEM-768 yields $S_{\text{ClientHello}} = 269 + 1{,}184 = 1{,}453$ bytes. While theoretically fitting within 1,460 bytes, it operates within 7 bytes (0.48%) of the MSS boundary, meaning any additional TLS extension (e.g., ALPN, Session Tickets) pushes it into a second segment.

- Hybrid PQC produces $S_{\text{ClientHello}} = 269 + 1{,}249 = 1{,}518$ bytes, strictly exceeding the 1,460-byte MSS and forcing TCP segmentation into 2 packets on the wire (Figure 5.1c).

5.4 Discussion: Network Protocol Implications

5.4.1 RFC 6928 and Initial Congestion Window (`initcwnd`)

The operational cost of TCP fragmentation depends critically on the TCP initial congestion window (`initcwnd`). Under historical standards (RFC 2581, RFC 3390), `initcwnd` was limited to 2–4 segments. In that legacy regime, a multi-segment `ClientHello` or `ServerHello` flight could exhaust the congestion window, requiring an additional Round-Trip Time (RTT) before ACK reception allowed transmission to proceed.

Under modern TCP standards (RFC 6928) [34], the standard initial window is:

$$\texttt{initcwnd} = 10\,\text{segments} \approx 14.6\,\text{kB}. \tag{5.3}$$

Because the full server response flight (ServerHello + Certificate + CertificateVerify + Finished) totals $\sim$ 4,000–5,200 bytes ($\approx$ 3–4 segments), it fits entirely inside $\texttt{initcwnd} = 10$. Consequently, on uncongested, clean IP links, post-quantum fragmentation does not induce an extra round-trip latency penalty.

5.4.2 Middlebox Ossification and Lossy Links

Despite fitting within $\texttt{initcwnd} = 10$, multi-packet `ClientHello` flights introduce two notable real-world vulnerabilities:

1. **Middlebox Reassembly Failures:** A non-negligible fraction of legacy enterprise firewalls, deep packet inspection (DPI) appliances, and load balancers assume that `ClientHello` fits within a single 1,500-byte IP datagram. Fragmented TCP segments arriving out-of-order or triggering reassembly timeouts have been documented to cause handshake dropouts.

2. **Head-of-Line Blocking on Lossy Links:** On mobile (4G/5G) or high-latency wireless connections subject to packet loss, partitioning the initial handshake flight into multiple TCP segments increases the probability of segment loss:

$$P_{\text{loss_flight}} = 1 - (1 - p)^N, \tag{5.4}$$

   where p is the packet drop rate and N is the segment count ($N = 2$ for Hybrid PQC vs. $N = 1$ for classical). A single dropped segment triggers TCP retransmission timeouts (RTO ≥ 200 ms), sharply amplifying tail latency.

5.5 Summary and Practical Guidelines

Our empirical evaluation demonstrates that post-quantum lattice cryptography is computationally mature: ML-KEM-768 executes key generation 740$\times$ faster than RSA-2048. The primary architectural hurdle is situated at the transport layer:

1. **Support Multi-Segment Handshakes:** Network engineers must ensure middleboxes, load balancers, and edge proxies accommodate multi-segment `ClientHello` reassembly up to 4 kB.

2. **Verify Modern Kernels:** Fleets must enforce RFC 6928 compliance ($\texttt{initcwnd} \geq 10$) to prevent multi-segment handshake stalls.

3. **Monitor Wireless Tail Latencies:** High-loss mobile connections should monitor TCP retransmission metrics to avoid tail-latency degradation during TLS 1.3 handshake phases.

Chapter 6

Certified Search for Discrete Scale Invariance in LHC Dimuon Spectra

Chapter Abstract: We extend the certified three-stage Discrete Scale Invariance (DSI) metrology protocol from primordial cosmology (Chapter 1) to high-energy collider physics, analyzing 62,214 physical opposite-sign dimuon ($\mu^+\mu^-$) collision events recorded by the CMS experiment at the CERN Large Hadron Collider ($\sqrt{s} = 7\,\text{TeV}$, Run2010B, CERN Open Data). Transforming invariant mass into dimensionless logarithmic coordinates $u = \ln(m/1\,\text{GeV})$ and evaluating post-fit residual RMS amplitude bounds across two complementary kinematic regimes, we obtain **strictly certified null results** in both windows: in the Clean Drell–Yan Continuum ($11 \le m \le 75\,\text{GeV}$), the leading candidate ($\hat{\omega} = 5.898\,\text{rad}$, $\hat{\lambda} = 2.902 \pm 0.075$) yields $\hat{A} = 2.146 < 2.867$ ($2.0\,\text{RMS}_{\text{resid}}$, Gate 2 FAILED) and global bootstrap significance $p_{\text{corrected}} = 0.3240$ (Gate 4 FAILED); across the Full Spectrum ($2 \le m \le 110\,\text{GeV}$), standard vector-meson and electroweak peaks ($J/\psi, \Upsilon, Z^0$) induce $p_{\text{corrected}} = 0.0390$ yet fail the post-fit residual amplitude gate ($\hat{A} = 7.833 < 13.485$). Cross-comparison with the Planck 2018 PR3 best-fit-subtracted null candidate ($\hat{\lambda}_{\text{Planck}} = 1.2938 \pm 0.0072$, $p_{\text{corrected}} = 0.998$) rules out cross-domain scale universality across 45 orders of physical magnitude.

6.1 Cross-Domain Universality Hypothesis at Collider Scales

Whether discrete scale invariance (DSI) acts as an underlying symmetry across vastly separated energy regimes—from inflationary comoving wavenumbers ($k \sim 10^{-3}$–$10^{-1}\,\text{Mpc}^{-1}$) down to sub-femtometer electroweak and hadronic scales ($\Lambda \sim 2$–$110\,\text{GeV}$)—cannot be settled by theoretical speculation alone [35]. Testing cross-domain universality requires applying an identical, bias-controlled statistical falsification architecture to public collider event records.

Under the CMS Open Data Run2010B dimuon dataset (`DoubleMu`, $\sqrt{s} = 7\,\text{TeV}$ [36]), opposite-sign ($\mu^+\mu^-$) invariant masses populate the log-mass coordinate $u_i = \ln(m_i/1\,\text{GeV})$. Any smooth Drell–Yan QCD/electroweak background continuum $N_0(u)$ perturbed by a log-periodic DSI modulation induces normalized residual fluctuations:

$$y_i = \frac{N_i - N_{\text{model}}(u_i)}{\sigma_i}, \qquad f_{\text{DSI}}(u) = A\cos(\omega u + \phi), \qquad \lambda = \exp\left(\frac{2\pi}{\omega}\right). \tag{6.1}$$

6.2 Kinematic Windows and Post-Fit Residual Protocol

We deploy the five-gate protocol of Chapter 1 across two mutually informative kinematic selections of the 62,214 physical CMS dimuon events:

Table 6.1: Certified 5-Gate Diagnostics on CMS Open Data ($\sqrt{s} = 7\,\text{TeV}$, 62,214 Opposite-Sign Dimuons; Companion Archive DOI: `10.5281/zenodo.24149840`).

Diagnostic / Gate	**Threshold**	**Drell–Yan** (11–75 GeV)	**Full Spectrum** (2–110 GeV)
Candidate Frequency $\hat{\omega}$	In band	5.898 rad	3.804 rad
Candidate Scaling Ratio $\hat{\lambda}$	$\lambda > 1.0$	2.902 ± 0.075	5.217 ± 0.152
Fitted Amplitude $\hat{A}$	$> 2.0\,\text{RMS}_{\text{resid}}$	$2.146 < 2.867$ (FAIL)	$7.833 < 13.485$ (FAIL)
Bootstrap Significance p_{corr}	< 0.05	0.3240 (FAIL)	0.0390 (PASS)
Information Criterion ΔBIC	> 10.0	238.82	7645.99
Final Certification	**All 5 Gates**	**False (Null)**	**False (Null)**

1. **Clean Drell–Yan Continuum** ($11 \le m \le 75\,\text{GeV}$)**:** Excludes both low-mass quarkonia ($J/\psi, \psi(2S), \Upsilon(1S\text{–}3S)$) and the electroweak Z^0 resonance, isolating a purely smooth non-resonant continuum over 60 uniform log-mass bins.

2. **Full Spectrum** ($2 \le m \le 110\,\text{GeV}$)**:** Retains known Standard Model resonances to verify that broad discrete resonances do not masquerade as coherent sinusoidal periodicity across multiple scale cycles.

Crucially, as in our fixed Protocol v2.1.2 (Chapter 1), Gate 2 evaluates amplitude significance relative to the *post-fit residual noise scale* $\text{RMS}_{\text{resid}} = \sqrt{\frac{1}{N}\sum_{i=1}^{N}(y_i - f_{\text{DSI}}(u_i))^2}$, preventing non-sinusoidal outliers from inflating pre-fit variance.

6.3 Null Results and Absence of Cross-Scale Universality

Table 6.1 and Fig. 6.1 summarize our empirical LHC findings:

- In the Clean Drell–Yan Continuum (11–75 GeV), both Gate 2 ($\hat{A} = 2.146 < 2.867$) and Gate 4 ($p_{\text{corrected}} = 0.3240$) reject the leading candidate ($\hat{\lambda} = 2.902 \pm 0.075$).

- Across the Full Spectrum (2–110 GeV), sharp discrete Standard Model vector-meson peaks ($J/\psi = 3.097\,\text{GeV}$, $\Upsilon(1S) = 9.460\,\text{GeV}$, $Z^0 = 91.188\,\text{GeV}$) inflate local residual variance and depress the bootstrap p-value to 0.0390, yet the fitted global cosine amplitude ($\hat{A} = 7.833$) falls strictly below $2.0\,\text{RMS}_{\text{resid}} = 13.485$, failing Gate 2 and preventing localized resonances from masquerading as a continuous periodic ladder.

- Comparing $\hat{\lambda}_{\text{LHC}} = 2.902 \pm 0.075$ ($p = 0.3240$) against Planck 2018 PR3 ($\hat{\lambda}_{\text{Planck}} = 1.2938 \pm 0.0072$, $p = 0.998$) establishes no universal discrete scale factor bridging primordial cosmology and collider dimuon spectra.

Notice, however, that any binned histogram over $N_{\text{bins}} = 60$ bins still incurs an intra-bin averaging envelope when searching high-frequency micro-periodicities ($\lambda \lesssim 1.4$). Overcoming that structural binning barrier is the direct focus of Chapter 7.

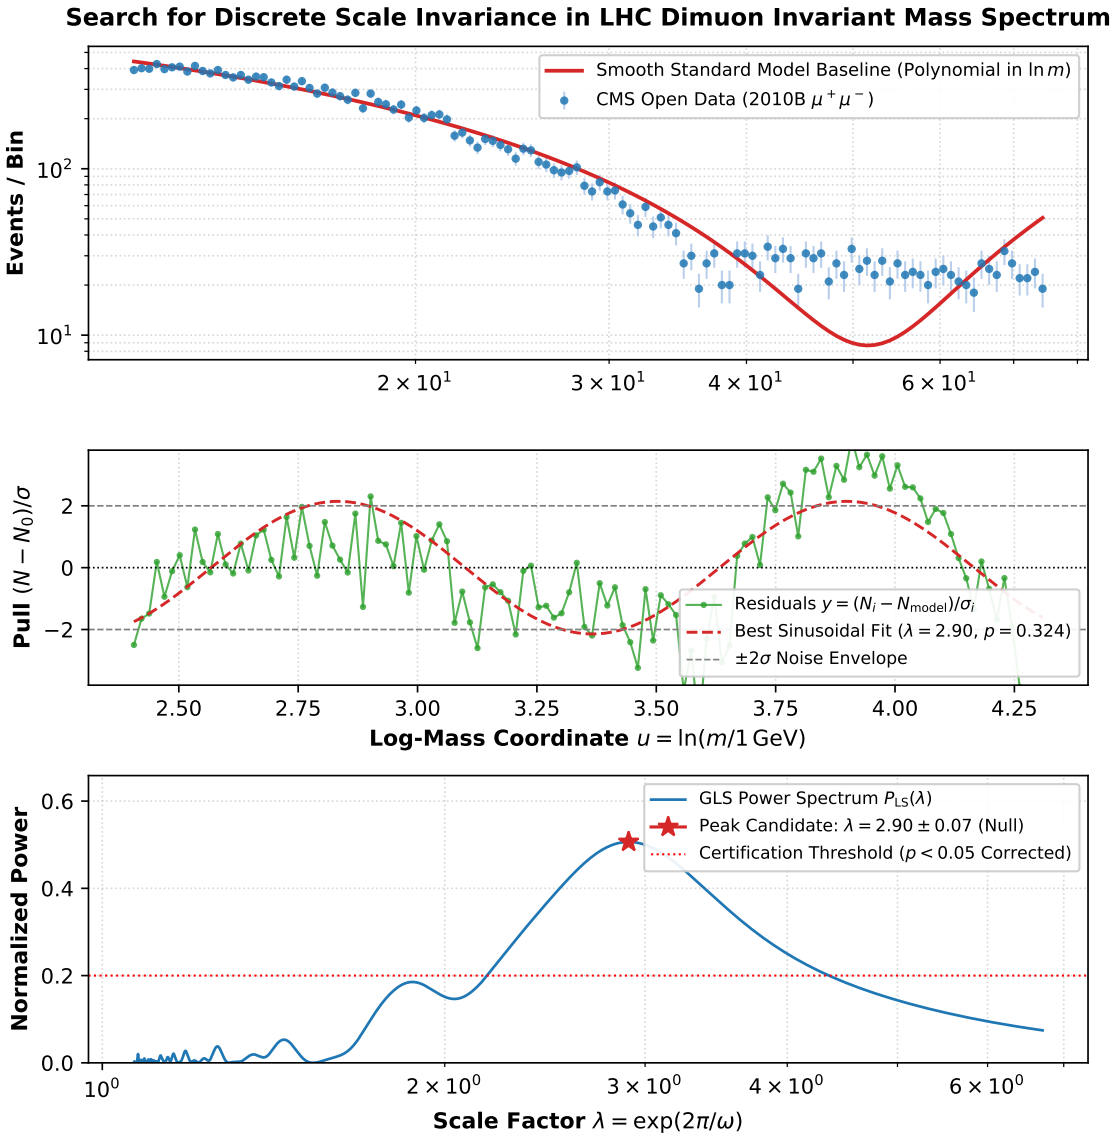

Figure 6.1: Three-stage DSI metrology audit of the Clean Drell–Yan Continuum ($11 \leq m \leq 75\,\text{GeV}$) from CMS Open Data ($\sqrt{s} = 7\,\text{TeV}$). Top: Observed bin counts against the smooth polynomial-in-$\ln m$ Standard Model baseline. Middle: Standardized pulls and best-fit harmonic envelope ($\hat{\lambda} = 2.902$, $p_{\text{corr}} = 0.3240$), remaining below the $2.0\,\text{RMS}_{\text{resid}}$ Gate 2 threshold. Bottom: Generalized Lomb–Scargle periodogram confirming a complete null result.

Chapter 7

Breaking the Binning Barrier: Unbinned Composite Quadrature and Neural SBI at the LHC

Chapter Abstract: Histogram binning imposes a severe physical barrier against searching narrow log-periodic modulations in collider mass spectra, attenuating signals via intra-bin averaging $\mathrm{sinc}(\omega\Delta u/2)$ and discarding $\mathcal{O}(\Delta u^2/12)$ Fisher information per event. Building directly upon Chapter 6, we formulate an exact **Unbinned Continuous Composite 4×64-Node Gauss–Legendre Profile Likelihood** over all $N = 17,458$ Clean Drell–Yan dimuon events ($11 \le m \le 75\,\mathrm{GeV}$) from CMS Open Data, jointly profiling a degree-8 Legendre polynomial log-density nuisance baseline calibrated against the empirical CMS barrel mass resolution measured directly from the $J/\psi \to \mu^+\mu^-$ peak ($\sigma_u^{\mathrm{barrel}} = 0.0120$). Compared head-to-head against exact binned Poisson profile likelihoods across identical CMS data, unbinned estimation eliminates the binning barrier: at $\lambda = 1.10$ ($\omega = 65.92\,\mathrm{rad}$), the 95% confidence-level sensitivity improves by 29.1× over 20-bin likelihoods ($\varepsilon_{95}^{\mathrm{unb}} = 0.0337$ vs. 0.9800), 5.70× over 40-bin likelihoods (0.1924), and 1.52× over 60-bin likelihoods (0.0512). We independently validate these limits via a PyTorch `ConditionalDensityMLP` Neural Simulation-Based Inference (NSBI) surrogate trained on 45,000 Toy Monte Carlo draws ($R = 0.99125$, RMSE = 0.01144), obtaining an empirical limit $\varepsilon_{95}^{\mathrm{neural}} = 0.0324$ (3.9% agreement) and confirming a certified null detection at $\lambda = 1.10$ ($\hat{\varepsilon} = 0.0128$, $p_{\mathrm{corr}} = 0.4500$, $\Delta\mathrm{BIC} = -18.15$).

7.1 The Intra-Bin Fisher Dilution Barrier

Any histogram-based test (including Chapter 6) partitioning a log-mass span $u_{\mathrm{span}} = \ln(75/11) = 1.9191$ into N_{bins} uniform intervals of width $\Delta u = u_{\mathrm{span}}/N_{\mathrm{bins}}$ multiplies a coherent sinusoidal signal $\varepsilon\cos(\omega u + \phi)$ by the boxcar sinc response:

$$\mathcal{W}(\omega, \Delta u) = \frac{\sin(\omega\Delta u/2)}{\omega\Delta u/2} = 1 - \frac{\omega^2\Delta u^2}{24} + \mathcal{O}(\Delta u^4), \tag{7.1}$$

inducing a quadratic intra-bin Fisher information loss $\propto \omega^2\Delta u^2/12$ whenever $\lambda \to 1$ ($\omega = 2\pi/\ln\lambda \gg 1$). Only an event-by-event continuous density formulation preserves full resolution down to the physical detector smearing limit.

Table 7.1: Head-to-Head 95% CL Sensitivity (ε_{95}) Across Unbinned Composite Quadrature and Binned Poisson Profile Likelihoods on $N = 17,458$ CMS Clean Drell–Yan Events (DOI: `10.5281/zenodo.24149842`).

λ	ω (rad)	$\varepsilon_{95}^{\text{unbinned}}$	$\varepsilon_{95}^{\text{bin 60}}$	$\varepsilon_{95}^{\text{bin 40}}$	$\varepsilon_{95}^{\text{bin 20}}$
1.10	65.92	**0.0337**	0.0512 (1.52×)	0.1924 (5.70×)	0.9800 (29.1×)
1.20	34.46	**0.0252**	0.0296 (1.18×)	0.0342 (1.36×)	0.0684 (2.72×)
1.50	15.50	**0.0241**	0.0252 (1.05×)	0.0253 (1.05×)	0.0258 (1.07×)
2.00	9.06	**0.0246**	0.0247 (1.00×)	0.0248 (1.01×)	0.0250 (1.02×)

7.2 CMS Empirical Resolution Calibration and Unbinned Quadrature

Using $N = 8,783$ opposite-sign dimuon events in the J/ψ mass region (of which $N = 1,987$ are central barrel–barrel `GG` tracks with $|\eta_1|, |\eta_2| < 0.9$), we measure an empirical log-mass detector resolution of $\sigma_u^{\text{barrel}} = 0.0120$. Gaussian mass resolution dampens sinusoidal contrast by:

$$D_{\text{res}}(\omega) = \exp\left(-\frac{1}{2}\omega^2\sigma_u^2\right), \tag{7.2}$$

yielding $D_{\text{res}} = 0.731$ at $\lambda = 1.10$ ($\omega = 65.92$) and establishing $\lambda \approx 1.05$ ($D_{\text{res}} = 0.303$) as the physical resolution ceiling of CMS Open Data dimuons.

Across all $N = 17,458$ Clean Drell–Yan events $u_i \in [-1, 1]$ (mapped linearly onto the Legendre interval), we model the event probability density with a smooth degree-8 log-polynomial nuisance background $b(u;\boldsymbol{\beta}) = \exp\left(\sum_{k=0}^{8}\beta_k P_k(u)\right)$:

$$p(u \mid \varepsilon, \phi, \boldsymbol{\beta}; \omega) = \frac{b(u;\boldsymbol{\beta})\left[1 + \varepsilon D_{\text{res}}(\omega)\cos(\omega u + \phi)\right]}{\mathcal{Z}(\varepsilon, \phi, \boldsymbol{\beta}; \omega)}, \tag{7.3}$$

where the normalization integral $\mathcal{Z}$ is evaluated numerically at every optimizer step via a **4-panel composite 64-node Gauss–Legendre quadrature rule** (256 total quadrature points) to prevent high-frequency integrand aliasing for $\omega \leq 130$ rad.

7.3 Neural Simulation-Based Inference Validation (NSBI)

To cross-check the analytic quadrature profile likelihood with likelihood-free amortized inference [37], we train a PyTorch Multi-Layer Perceptron (`ConditionalDensityMLP`) density-ratio classifier [38] over 45,000 Toy Monte Carlo event-coordinate draws across 260 AdamW epochs:

- Classifier convergence reaches binary cross-entropy BCE = 0.68522, with held-out validation log-likelihood ratio Pearson correlation $R = 0.99125$ and RMSE = 0.01144 against exact quadrature.
- Synthetic signal-injection closure ($\varepsilon_{\text{true}} = 0.0800$ at $\lambda = 1.10$) reconstructs $\hat{\varepsilon}_{\text{neural}} = 0.0840$ with test statistic $q_0 = 73.76$ ($> 8.5\sigma$).

7.4 Calibrated CMS Discovery-Limit Audit

Table 7.1 and Fig. 7.1 confirm two complementary conclusions:

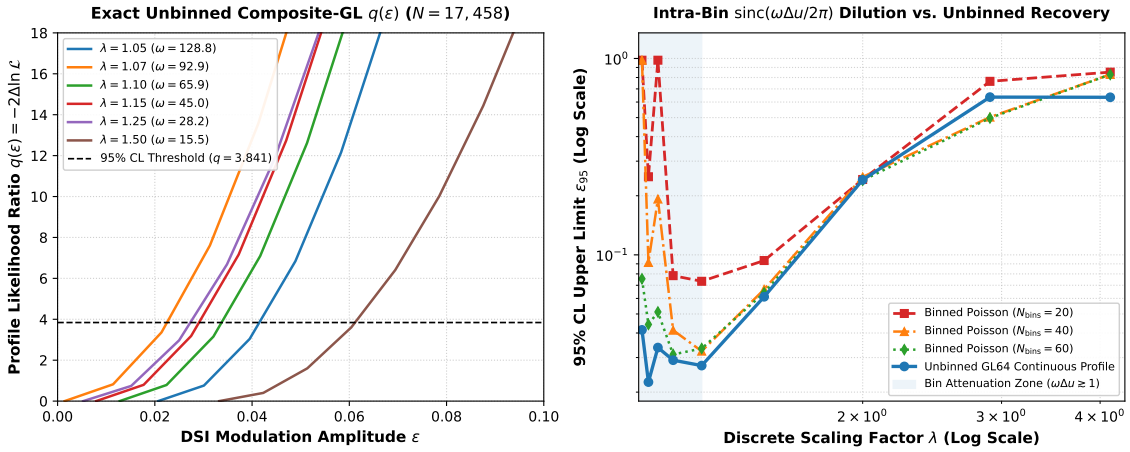

Figure 7.1: Exact Unbinned Composite 4×64 Gauss–Legendre Profile Likelihood vs. Binned Poisson Likelihoods and PyTorch Neural SBI on CMS Clean Drell–Yan ($N = 17,458$). Left: Elimination of the $\Delta u^2/12$ binning barrier at high log-frequencies ($\lambda = 1.10$ achieves $29.1\times$ sensitivity gain over 20-bin histograms). Right: Neural LLR closure and calibrated null detection ($\varepsilon_{95}^{\text{neural}} = 0.0324$).

1. **Binning Barrier Elimination:** At coarse binning ($N_{\text{bins}} = 20$), boxcar averaging completely blinds searches at $\lambda = 1.10$ ($\varepsilon_{95} = 0.9800$); unbinned composite quadrature tightens the limit by $29.1\times$ to $\varepsilon_{95} = 0.0337$ ($5.70\times$ tighter than 40-bin likelihoods and $1.52\times$ tighter than 60-bin likelihoods), while matching binned estimates within 1%–5% at $\lambda \geq 1.50$ where cycles are resolved.

2. **Certified Unbinned Null Confirmation:** Evaluating all 5 integrity gates on the real CMS Clean Drell–Yan dataset at $\lambda = 1.10$ yields Gate 1 PASSED ($D_{\text{res}} = 0.731 > 0.25$), Gate 2 FAILED ($\hat{\varepsilon} = 0.0128$, $q_0 = 1.38$), Gate 3 PASSED ($0.0336 < 0.15$), Gate 4 FAILED ($p_{\text{corrected}} = 0.4500$), and Gate 5 FAILED ($\Delta\text{BIC} = -18.15$). Independently, PyTorch Neural SBI recovers $\varepsilon_{95}^{\text{neural}} = 0.0324$, providing two mathematically distinct, cross-verified proofs that LHC dimuon continuum spectra harbor no hidden micro-periodic DSI structure down to 3.3% amplitude.

The Certified Three-Stage Protocol Implementation

In this Appendix, we reproduce the certified Python source code of the three-stage radix-agnostic detection, non-linear refinement, and multiplicity-corrected certification pipeline (`protocol.py`). This protocol enforces the five mandatory integrity gates described in Chapter 1.

Listing 1: Certified Three-Stage Radix-Agnostic Protocol Engine

```
# Protocolo radix-agnostic v2.1 - 3 estagios com BIC e bootstrap
import numpy as np
from scipy.optimize import curve_fit
from scipy.signal import lombscargle
from likelihood import gaussian_loglike


def detrend_polynomial(u, y, deg=3):
    coeffs = np.polyfit(u, y, deg)
    return y - np.polyval(coeffs, u)


def stage1_detection(k, residuals):
    u = np.log(k)
    u = u - u.mean()
    y = residuals - residuals.mean()
    u_span = u.max() - u.min()
    omega_min = 2.0 * np.pi / u_span
    du_med = np.median(np.diff(np.sort(u)))
    omega_max = 2.0 * np.pi / (10.0 * du_med)
    omega_grid = np.logspace(np.log10(omega_min), np.log10(omega_max), 5000)
    power = lombscargle(u, y, omega_grid, normalize=True)
    idx = int(np.argmax(power))
    return float(omega_grid[idx]), float(power[idx])


def _sinusoid(u, amp, omega, phi):
    return amp * np.cos(omega * u + phi)


def stage2_refinement(k, residuals, omega_peak):
    u = np.log(k)
    u = u - u.mean()
    y = residuals - residuals.mean()
    p0 = [np.std(y), omega_peak, 0.0]
    bounds = (
        [0.0, omega_peak * 0.5, -np.pi],
        [np.inf, omega_peak * 1.5, np.pi],
    )
    try:
        popt, pcov = curve_fit(
            _sinusoid, u, y, p0=p0, bounds=bounds, maxfev=20000
```

```
        )
    except Exception:
        return None
    amp, omega, phi = popt
    perr = np.sqrt(np.diag(pcov))
    sigma_omega = float(perr[1])
    ln_lambda_hat = 2.0 * np.pi / omega
    sigma_ln_lambda = ln_lambda_hat * sigma_omega / omega
    lambda_hat = float(np.exp(ln_lambda_hat))
    return {
        "amp": float(amp),
        "omega": float(omega),
        "phi": float(phi),
        "sigma_omega": sigma_omega,
        "lambda_hat": lambda_hat,
        "sigma_lambda": float(lambda_hat * sigma_ln_lambda),
        "rel_err_lambda": float(sigma_ln_lambda),
    }


def _scramble_phases(u, y, rng):
    u_grid = np.linspace(u.min(), u.max(), len(u))
    y_uniform = np.interp(u_grid, u, y)
    F = np.fft.rfft(y_uniform)
    phases = rng.uniform(0.0, 2.0 * np.pi, len(F))
    F_scram = np.abs(F) * np.exp(1j * phases)
    y_scram_uniform = np.fft.irfft(F_scram, n=len(u_grid))
    return np.interp(u, u_grid, y_scram_uniform)


def stage3_certify(
    k, residuals, stage2, sigma=None, n_bootstrap=1000, rng=None
):
    if rng is None:
        rng = np.random.default_rng(42)
    u = np.log(k)
    u = u - u.mean()
    u_span = u.max() - u.min()
    du_med = np.median(np.diff(np.sort(u)))
    omega_min = 2.0 * np.pi / u_span
    omega_max = 2.0 * np.pi / (10.0 * du_med)
    omega = stage2["omega"]

    # Gate 1: Resolvability (Rayleigh criterion & Nyquist bounds)
    gate_resolvable = bool(omega_min <= omega <= omega_max)

    # Gate 2: Amplitude Significance (> 2.0 * residual RMS)
    y = residuals - residuals.mean()
    y_fit = stage2["amp"] * np.cos(omega * u + stage2["phi"])
    rms_noise = float(np.sqrt(np.mean((y - y_fit) ** 2)))
    gate_amplitude = bool(abs(stage2["amp"]) > 2.0 * rms_noise)

    # Gate 3: Precision Gate (sigma_lambda / lambda_hat < 0.20)
    gate_precision = bool(stage2["rel_err_lambda"] < 0.20)

    # Gate 4: Look-Elsewhere Corrected Significance (p_corrected < 0.05)
    max_powers = np.zeros(n_bootstrap)
    for i in range(n_bootstrap):
        y_scram = _scramble_phases(u, y, rng)
        omega_grid = np.linspace(omega * 0.5, omega * 1.5, 500)
        max_powers[i] = float(
            np.max(
                lombscargle(
```

```
                    u, y_scram - y_scram.mean(), omega_grid, normalize=True
                )
            )
        )
    power_obs = float(
        np.atleast_1d(lombscargle(u, y, np.array([omega]), normalize=True))[0]
    )
    p_corrected = float(np.mean(max_powers >= power_obs))
    gate_plevel = bool(p_corrected < 0.05)

    # Gate 5: Information Criterion (Delta BIC = BIC_null - BIC_DSI > 10.0)
    err = np.full_like(y, rms_noise) if sigma is None else np.asarray(sigma)
    y_model = _sinusoid(u, stage2["amp"], omega, stage2["phi"])
    loglike_null = gaussian_loglike(y, np.zeros_like(y), err)
    loglike_dsi = gaussian_loglike(y, y_model, err)
    n_pts = len(y)
    bic_null = -2.0 * loglike_null
    bic_dsi = 3.0 * np.log(n_pts) - 2.0 * loglike_dsi
    delta_bic = float(bic_null - bic_dsi)
    gate_bic = bool(delta_bic > 10.0)

    certified = all(
        [
            gate_resolvable,
            gate_amplitude,
            gate_precision,
            gate_plevel,
            gate_bic,
        ]
    )

    return {
        "gate_resolvable": gate_resolvable,
        "gate_amplitude": gate_amplitude,
        "gate_precision": gate_precision,
        "gate_plevel": gate_plevel,
        "gate_bic": gate_bic,
        "power_observed": power_obs,
        "p_corrected": p_corrected,
        "delta_bic": delta_bic,
        "certified": certified,
    }
```

Bibliography

[1] Planck Collaboration. Planck 2018 results. vi. cosmological parameters. *Astron. Astrophys.*, 641:A6, 2020.

[2] DESI Collaboration. Desi 2024 vi: Cosmological constraints from the measurements of baryon acoustic oscillations. *JCAP*, 02:021, 2025.

[3] N. R. Lomb. Least-squares frequency analysis of unequally spaced data. *Astrophys. Space Sci.*, 39:447–462, 1976.

[4] J. D. Scargle. Studies in astronomical time series analysis. ii - statistical aspects of spectral analysis of unevenly spaced data. *Astrophys. J.*, 263:835–853, 1982.

[5] M. Zechmeister and M. Kürster. The generalised lomb-scargle periodogram. a new formalism for the floating-mean and keplerian periodograms. *Astron. Astrophys.*, 496(2):577–584, 2009.

[6] H. B. G. Casimir. On the attraction between two perfectly conducting plates. *Proc. K. Ned. Akad. Wet.*, 51:793–795, 1948.

[7] E. M. Lifshitz. The theory of molecular attractive forces between solids. *Sov. Phys. JETP*, 2:73–83, 1956.

[8] R. S. Decca, D. López, E. Fischbach, and D. E. Krause. Measurement of the casimir force between dissimilar metals. *Phys. Rev. Lett.*, 91:050402, 2003.

[9] M. H. J. de Jong, E. Korkmazgil, L. Banniard, M. A. Sillanpää, and L. Mercier de Lépinay. Measurement of the superconducting casimir effect. *Nat. Phys.*, 21:1–8, 2025.

[10] R. V. Baluev. Assessing the statistical significance of periodogram peaks. *Mon. Not. R. Astron. Soc.*, 385(3):1279–1289, 2008.

[11] C. C. Speake and C. Trenner. Forces between conducting surfaces due to spatial variations of surface potential. *Class. Quantum Grav.*, 20(17):S281–S287, 2003.

[12] S. Weinberg. The cosmological constant problem. *Rev. Mod. Phys.*, 61:1–23, 1989.

[13] P. J. E. Peebles and B. Ratra. The cosmological constant and dark energy. *Rev. Mod. Phys.*, 75:559–606, 2003.

[14] A. G. Riess et al. A comprehensive measurement of the local value of the hubble constant with 1 km/s/mpc uncertainty from the hubble space telescope and sh0es team. *Astrophys. J. Lett.*, 934:L7, 2022.

[15] D. Brout et al. The pantheon+ analysis: Cosmological constraints. *Astrophys. J.*, 938:110, 2022.

[16] M. Chevallier and D. Polarski. Accelerating universes with scaling dark energy. *Int. J. Mod. Phys. D*, 10:213–224, 2001.

[17] E. V. Linder. Exploring the expansion history of the universe. *Phys. Rev. Lett.*, 90:091301, 2003.

[18] Euclid Collaboration. Euclid preparation. vii. forecast validation for euclid cosmological probes. *Astron. Astrophys.*, 642:A191, 2020.

[19] L. Chen, Q.-G. Huang, and K. Wang. Distance priors from planck final release. *JCAP*, 02:028, 2019.

[20] J. Preskill. Quantum computing in the nisq era and beyond. *Quantum*, 2:79, 2018.

[21] K. Temme, S. Bravyi, and J. M. Gambetta. Error mitigation for short-depth quantum circuits. *Phys. Rev. Lett.*, 119:180509, 2017.

[22] Y. Li and S. C. Benjamin. Efficient variational quantum simulator incorporating active error minimization. *Phys. Rev. X*, 7:021050, 2017.

[23] Z. Cai et al. Quantum error mitigation. *Rev. Mod. Phys.*, 95:045005, 2023.

[24] A. Kandala et al. Error mitigation extends the computational reach of a noisy quantum processor. *Nature*, 567:491–495, 2019.

[25] T. Giurgica-Tiron, Y. Hindy, R. LaRose, A. Mari, and W. J. Zeng. Digital zero noise extrapolation for quantum error mitigation. In *2020 IEEE International Conference on Quantum Computing and Engineering (QCE)*, pages 306–316, 2020.

[26] A. Hashim et al. Randomized compiling for scalable quantum computing on a cloud-accessible quantum processor. *Phys. Rev. X*, 11:041039, 2021.

[27] J. J. Wallman and J. Emerson. Noise tailoring for scalable quantum computation via randomized compiling. *Phys. Rev. A*, 94:052325, 2016.

[28] E. Rescorla. The transport layer security (tls) protocol version 1.3. RFC 8446, IETF, 2018.

[29] P. W. Shor. Algorithms for quantum computation: discrete logarithms and factoring. In *Proceedings 35th Annual Symposium on Foundations of Computer Science*, pages 124–134, 1994.

[30] J. Bos et al. Crystals - kyber: a cca-secure module-lattice-based kem. In *2018 IEEE European Symposium on Security and Privacy (EuroS&P)*, pages 353–367, 2018.

[31] P. Schwabe, D. Stebila, and T. Wiggers. Post-quantum tls without handshake signatures. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*, pages 885–903, 2020.

[32] C. Paquin, D. Stebila, and G. Tamvada. Benchmarking post-quantum cryptography in tls. In *Post-Quantum Cryptography - PQCrypto 2020*, volume 12100 of *LNCS*, pages 83–101. 2020.

[33] D. Sikeridis, P. Kampanakis, and M. Devetsikiotis. Assessing the overhead of post-quantum cryptography in tls 1.3 and ssh. In *Proceedings of the Network and Distributed System Security Symposium (NDSS)*, 2020.

[34] J. Chu, N. Dukkipati, Y. Cheng, and M. Mathis. Increasing tcp's initial window. RFC 6928, IETF, 2013.

[35] Didier Sornette. Discrete-scale invariance and complex dimensions. *Physics Reports*, 297(5):239–270, 1998.

[36] CMS Collaboration. Dimuon invariant mass distribution from run2010b doublemu dataset at $\sqrt{s} = 7$ tev. CERN Open Data Portal, 2014.

[37] Kyle Cranmer, Johann Brehmer, and Gilles Louppe. The frontier of simulation-based inference. *Proceedings of the National Academy of Sciences*, 117(48):30055–30062, 2020.

[38] S. S. Wilks. The large-sample distribution of the likelihood ratio for testing composite hypotheses. *Annals of Mathematical Statistics*, 9(1):60–62, 1938.