Matrix Node Theory / Evans Quantum Energy Field

Open-Data Testing and Model Verification

A reproducible multi-domain assessment of a conservative phase–reserve model

Jordan Ryan Evans

Independent research

AI-assisted mathematical analysis, computation and editing

Research report OV-1.0

Execution snapshot: 6–7 September 2026

What this edition contains. Executed analyses of published DESI, FIRAS and CMS numerical products; XENON limit translations; clock, muon and neutrino comparisons; conservative-model stress tests; synthetic recovery controls; and downloadable source code, inputs and result records.

What it does not claim. A completed theory of everything, a new independently predicted physical constant, an independently confirmed EQEF signal, or a raw GWOSC / 100,000-event CMS analysis performed in this workspace. Those larger raw-file pipelines are supplied and synthetic-fixture tested, but acquisition failed here. This report is not independent peer review or an endorsement by any experiment.

Abstract

The potential-to-actuation interpretation of Matrix Node Theory is tested through an explicitly conservative compact-phase network coupled to reserve oscillators. The microscopic coefficients are held fixed. A separate implementation verifies conservation, common-phase symmetry, positivity of an energy partition, time reversal, clock reparameterization and the full coupled spectrum across four graph structures and 36 nonlinear preparations. Fifty directional spectral checks recover the conditional cubic-lattice relation $b_{\mathrm{axis}} - b_{\mathrm{body}} = -1/36$. This is a model consequence, not an identified physical detection.

The observational work uses three published numerical products: 13 DESI DR2 baryon-acoustic-oscillation measurements with their supplied covariance, 43 COBE/FIRAS spectral-residual channels with an approximate published covariance, and 26 CMS Drell–Yan normalized-ratio measurements. The fixed supplied cosmological comparator gives $\chi^2 = 28.50623$ for 13 measurements. A separately fitted standard flat-ΛCDM comparator gives $\chi^2 = 10.28401$ for nominally 11 degrees of freedom, with a refitted-bootstrap tail probability 0.500500. The improvement belongs to that standard comparator, not to a microscopic MNT prediction.

Separate FIRAS residual fits give $y = (-0.227 \pm 3.896) \times 10^{-6}$ and $\mu = (-13.112 \pm 36.962) \times 10^{-6}$ within the stated table-level approximation. Their uncertainties do not replace the collaboration's full systematic treatment. Twenty-four of 26 CMS marginal comparisons lie within two conditional standard deviations of the publication's NNLO FEWZ reference; an approximately 3.463σ local discrepancy is retained. Full inter-bin covariance is unavailable for that comparison, so no formal global CMS probability is assigned.

Additional work reconstructs 20 XENONnT limit rows, translates ten heavy-mediator limits into one stated contact-interaction convention, evaluates seven published-summary interfaces and a GW170817 propagation interval, and produces 120 conventional neutrino-propagation forecasts. The record contains 1684 successful software and mathematical assertions: 738 new and 946 retained assertions rerun. The new statistical calibration comprises 1,000 refitted DESI simulations and 4,000 FIRAS coverage trials. Neither these counts nor the number of numerical rows is an experimental accuracy percentage.

The raw detector code verifies file provenance, handles data-quality and rounding checks, and passes synthetic end-to-end tests. No real large raw-detector file was acquired in this workspace. The resulting publication distinguishes what has been established for a finite mathematical model, what agrees conditionally with known physics, and what remains to be mapped to an independently testable physical observable.

Reading the evidence

This is a testing report, not an attempt to make every row say "confirmed." A proposition about a specified Hamiltonian is proved from its assumptions. A numerical test checks an implementation to a declared tolerance. A published-data analysis tests a declared prediction or reference model against an external measurement product. A synthetic experiment tests the behavior of an estimator when its generating model is known. These are different activities.

A compact theory can be scientifically valuable before all of its physical identifications are established. The most useful report at that stage is one which makes its interfaces explicit. Where an MNT-specific source spectrum, mass, charge, interaction or cosmological history is not specified, a conventional reference is used and named as such. Agreement of that reference with observation is not promoted to a new discovery by changing its label.

The core model and the accompanying calculations are stated in ordinary Hamiltonian, statistical and relativistic notation. "Reserve" denotes explicit energy-storage degrees of freedom. "Actuation" denotes a resolved configuration or a detector registration. Identifying the reserve with EQEF is the physical hypothesis; the reserve construction itself is a particular mathematical realization, not a unique consequence of the terminology.

All headline numbers below are read from executable output files. The companion CSV ledger exposes individual assertions, including assertions that correctly recognize the limitations of a historical formula. The publication bibliography credits external scientific work. Development provenance, the inherited model snapshot and code hashes are recorded separately, rather than being offered as empirical evidence.

Contents

1. Scope, hypotheses and the meaning of a successful test

1.1 Three questions that should not be conflated

The investigation has three distinct questions. First, does the declared node–reserve system possess a consistent finite dynamics? Second, are the numerical algorithms correct and sufficiently accurate? Third, does an independently specified physical interpretation produce successful predictions for external measurements? The first two are necessary for the third but do not imply it.

For this edition, the answer to the first question is established for the stated finite Hamiltonian and its assumptions. The answer to the second is supported by the executed assertion ledger, analytic identities, independent numerical formulations and deliberately incorrect controls. The third question is only partly accessible: a number of conventional effective descriptions can be compared with measured data, but the finite model does not yet uniquely supply every particle, interaction and detector map.

The null statement "no additional observable deviation" is often allowed by a proposed extension. Its compatibility with data is a useful constraint, but it cannot favor that extension over another model which makes exactly the same prediction. For two models M_0, M_1 with identical likelihoods on the measured observables,

$$p(d \mid M_0) = p(d \mid M_1), \tag{1.1}$$

the data provide no likelihood preference between them. This simple equality prevents a conventional gravity or clock agreement from being recast as unique evidence for EQEF.

1.2 The hypothesis-to-observation contract

A testable microscopic proposal requires more than an evolution equation. Let ϑ collect its universal coefficients, z_0 its preparation and ν_d the experimentally constrained nuisance quantities for dataset d. A generic detector prediction is

$$\mu_{dj}(\vartheta, \nu_d) = \int_{\mathcal{T}_d} \mathrm{d}t \int \mathrm{d}\xi \, R_{dj}(\xi; \nu_d) \, F_d[z(t; \vartheta, z_0), \xi] + b_{dj}(\nu_d). \tag{1.2}$$

Here F_d maps the microscopic state into a physical emitted quantity, R_{dj} is the instrument response and b_{dj} is the background. Each term must carry a specified dimension and a reproducible numerical meaning. A detector response is not a free function selected to repair a theory discrepancy.

For the finite-model tests, z and its equations are fully given. For the DESI, FIRAS and CMS comparisons, the particular physical reference is stated before its result. For a full MNT recoil likelihood or collider production rate, the missing F_d is not filled with an arbitrary observed spectrum. This is a definition of the boundary of the test, not a presumption that such a map can never be developed.

1.3 Result classes used in the report

Class	Meaning
Mathematical verification	A theorem, identity or implementation check within stated assumptions.
Synthetic inference	Recovery or rejection using deliberately generated observations.
Published-product test	A calculation against an experiment's public vector, spectrum, covariance or limit table.
Published-summary test	A comparison with a reported estimate and uncertainty, without reconstructing the original likelihood.
Raw reconstruction	A calculation from actual event or strain files; not executed on the large files in this edition.
Forecast	An output for a declared future or hypothetical configuration, not an observation.
Independent physical prediction	An observable fixed without using its target measurement, then tested with an adequate physical likelihood. No new instance is established here.

The last class is intentionally demanding. A local content hash helps preserve a calculation but does not make a previously known measurement unknown, establish an independent timestamp, or prove that an observable map was derived without measurement input.

2. The frozen conservative system

Evidence class: Mathematical construction and numerical verification..

2.1 Variables, dimensions and Hamiltonian

Choose a finite connected graph with oriented incidence matrix B. The orientation is an accounting convention: a column has $+1$ at one end and -1 at the other. A compact phase θ_i and its canonical momentum p_i live on each vertex. Each edge has reserve coordinates x_{ea} and momenta P_{ea}. The relative phase is $\delta = B^T\theta$.

The tested Hamiltonian is

$$H = \sum_i \frac{p_i^2}{2I} + K\sum_e (1-\cos\delta_e) + \sum_{e,a}\left[\frac{P_{ea}^2}{2m_a} + \frac{m_a\Omega_a^2}{2}\left(x_{ea} - \frac{c_a \sin\delta_e}{m_a\Omega_a^2}\right)^2\right]. \tag{2.1}$$

Take $I, m_a, K, \Omega_a^2 > 0$ and real c_a. The phase is dimensionless. If x is assigned a length dimension, m_a is an inertial mass and c_a a force; alternatively all quantities may be expressed in one explicit system of model units. The numerical tests use model units throughout. They do not identify an arbitrary simulation second with a measured physical second.

The shifted square contains the reserve–node interaction and its counterterm. Expanding it gives a free reserve, a bilinear interaction $-c_a x_{ea}\sin\delta_e$ and a positive term $c_a^2\sin^2\delta_e/(2m_a\Omega_a^2)$.

The counterterm is part of the model, not a discrepancy-dependent fit. Conservative oscillator reservoirs and their reduced memory equations are established mathematical constructions [1, 2, 3]; here they implement the proposed potential/actuation interpretation.

The common numerical coefficients are fixed at

$$I = K = m_a = 1, \qquad \boldsymbol{\Omega} = (0.65, 1.15, 2.4), \qquad \boldsymbol{c} = (0.22, 0.13, 0.07). \tag{2.2}$$

There is no separate coefficient for an electron, a muon, a galaxy or a gravitational-wave event. Conversely, using identical coefficients on every edge does not by itself prove that the eigenmodes have those physical identities.

2.2 Equations of motion

Define

$$r_{ea} = x_{ea} - \frac{c_a \sin\delta_e}{m_a \Omega_a^2}, \qquad f_e = K \sin\delta_e - \cos\delta_e \sum_a c_a r_{ea}. \tag{2.3}$$

The canonical evolution is

$$\dot\theta_i = p_i / I, \qquad \dot p_i = -(Bf)_i, \tag{2.4}$$

$$\dot x_{ea} = P_{ea}/m_a, \qquad \dot P_{ea} = -m_a \Omega_a^2 r_{ea}. \tag{2.5}$$

Dots denote differentiation with respect to the declared evolution parameter, interpreted as local proper time only when a physical clock map has been given. The coupling is reciprocal: a reserve force changes the nodes, while the node phase drives the reserve. Removing either feedback term changes the model and invalidates its conservation proof.

Proposition 2.1 (Conservation and a nonnegative partition)**.** *For a solution of* (2.5)*, the total H is constant. In the stated partition,*

$$H_A = \sum_i \frac{p_i^2}{2I} + K \sum_e (1 - \cos\delta_e), \qquad H_R = H - H_A, \tag{2.6}$$

both H_A and H_R are nonnegative. Changes in one are balanced by the other.

Proof. Writing all canonical coordinates as z, the equations are $\dot z = J\nabla H$ with $J^T = -J$. Consequently $\dot H = (\nabla H)^T J \nabla H = 0$. Each kinetic or shifted-square contribution is nonnegative, as is $1 - \cos\delta_e$. Thus $\dot H_A + \dot H_R = 0$. A different allocation of interaction energy changes the two separate labels, not the conserved total. □

The interpretation of a smaller H_A is less resolved energy in this partition, not destruction. An observed detector deficit can be identified with reserve transfer only after ordinary acceptance, backgrounds and instrumental losses have been accounted for.

Proposition 2.2 (Common-phase charge and existence)**.** *The common-phase transformation $\theta \mapsto \theta + \alpha\mathbf{1}$ leaves H unchanged. Its canonical charge $Q = \sum_i p_i$ is conserved. A finite initial energy gives a global solution on the finite graph.*

Proof. Since $B^T\mathbf{1} = 0$, δ is invariant, and $\dot{Q} = -\mathbf{1}^T Bf = 0$, the elementary form of the Noether relation [4]. Finite H bounds p, P and the shifted reserve coordinates r. Bounded sine functions then bound x. The compact phases have smooth bounded velocities on every finite interval. The smooth vector field cannot develop a finite-time escape from these bounds. Standard continuation of the local solution therefore supplies a global solution. □

This proof is a finite-system result. It does not automatically establish an infinite-volume continuum limit, relativistic causality, a chiral quantum field theory or Einstein's dynamical equation.

3. What the reserve implies: memory, return and fluctuation

3.1 Eliminating the reserve exactly

The reserve equation is a forced oscillator,

$$\ddot{x}_{ea} + \Omega_a^2 x_{ea} = \frac{c_a}{m_a} \sin\delta_e. \tag{3.1}$$

Set $u_e(\tau) = \sin\delta_e(\tau)$. Integration with the initial conditions gives

$$\begin{aligned} x_{ea}(\tau) = x_{ea}(0)\cos\Omega_a\tau + \frac{P_{ea}(0)}{m_a\Omega_a}\sin\Omega_a\tau \\ + \frac{c_a}{m_a\Omega_a}\int_0^\tau \sin[\Omega_a(\tau - s)]u_e(s)\,\mathrm{d}s. \end{aligned} \tag{3.2}$$

Integrating the last term by parts and subtracting $c_a u_e(\tau)/(m_a\Omega_a^2)$ yields

$$\sum_a c_a r_{ea}(\tau) = \xi_e(\tau) - \int_0^\tau \Gamma(\tau - s)\dot{u}_e(s)\,\mathrm{d}s, \tag{3.3}$$

where

$$\Gamma(v) = \sum_a \frac{c_a^2}{m_a\Omega_a^2}\cos\Omega_a v, \tag{3.4}$$

$$\xi_e(\tau) = \sum_a c_a\left[r_{ea}(0)\cos\Omega_a\tau + \frac{P_{ea}(0)}{m_a\Omega_a}\sin\Omega_a\tau\right]. \tag{3.5}$$

Substitution into the node equation produces a history-dependent force with the same coefficients as the full conservative system. There is no independent damping constant inserted after an observational mismatch. The initial reserve state is retained rather than silently set to a convenient fitted noise term.

The exact memory equation and the direct coordinate evolution are two representations of the same model. Agreement between them is a useful implementation test because an omitted initial-slip or counterterm contribution produces a calculable disagreement. It is not two independent experimental confirmations.

3.2 Finite reserves do not guarantee irreversible drag

A finite sum of cosines can return stored energy. The system can show recurrences and coherent exchange rather than monotone damping. Approximating the memory kernel by a local friction coefficient requires an additional spectral or time-scale limit. Taking such a limit is a new assumption which must be named and tested. It cannot be used interchangeably with the finite three-frequency reserve when evaluating a frequency-dependent observation.

This distinction supplies a concrete implication of the EQEF interpretation: a measured dissipative response, a reactive frequency shift and any return signal must be described consistently by the same response model. One cannot assign a new independent "loss" parameter to each of these measurements without changing the hypothesis.

3.3 A conditional fluctuation relation

Prepare the shifted reserve coordinates and momenta in a classical canonical ensemble at temperature T, with independent edges. Their variances are

$$\mathbb{E}[r_{ea}(0)^2] = \frac{k_B T}{m_a \Omega_a^2}, \qquad \mathbb{E}[P_{ea}(0)^2] = m_a k_B T. \tag{3.6}$$

Using the explicit form of ξ and the cosine subtraction identity gives

$$\mathbb{E}[\xi_e(\tau)\xi_{e'}(s)] = \delta_{ee'} k_B T \Gamma(\tau - s). \tag{3.7}$$

Thus response and fluctuations are linked under that preparation, consistent with the established fluctuation–dissipation framework [5]. The thermal ensemble is an additional physical assumption; (3.7) is not a derivation of quantum Born probabilities. Nor does a classical thermal reservoir establish a cosmological dark-energy equation of state.

4. Propagation with the reserve retained

4.1 Full normal modes rather than a node-only approximation

Linearize about uniform phase and the associated reserve equilibrium. For a graph-Laplacian eigenvalue λ, one coupled block has stiffness and mass matrices

$$S(\lambda) = \begin{pmatrix} \lambda(K + A_0) & -\sqrt{\lambda}\boldsymbol{c}^T \\ -\sqrt{\lambda}\boldsymbol{c} & \mathrm{diag}(m_a\Omega_a^2) \end{pmatrix}, \quad M = \mathrm{diag}(I, m_1, m_2, m_3), \tag{4.1}$$

with $A_0 = \sum_a c_a^2/(m_a\Omega_a^2)$. The normal modes satisfy

$$Sv = \omega^2 M v. \tag{4.2}$$

Eliminating the reserve amplitudes gives, away from a bare pole,

$$K\lambda - I\omega^2 - \lambda \sum_a \frac{c_a^2}{m_a\Omega_a^2}\frac{\omega^2}{\Omega_a^2 - \omega^2} = 0. \tag{4.3}$$

The apparent pole belongs to the eliminated-coordinate representation. The full coupled matrix is the appropriate object near hybridization; a singular denominator in a reduced expression is not by itself a physical instability.

4.2 A parameter-canceling directional consequence

For a simple cubic lattice of spacing a,

$$\lambda(\mathbf{k}) = 4\sum_{j=1}^{3}\sin^2(k_j a/2) = (ka)^2 - \frac{(ka)^4}{12}\sum_j \hat{k}_j^4 + O((ka)^6). \tag{4.4}$$

Let $A_1 = \sum_a c_a^2/(m_a\Omega_a^4)$. Expanding (4.3) at frequencies below every reserve gap gives

$$\omega^2 = \frac{K}{I}\lambda - \frac{KA_1}{I^2}\lambda^2 + O(\lambda^3). \tag{4.5}$$

Consequently, with $c_* = a\sqrt{K/I}$,

$$\frac{\omega}{c_* k} = 1 + b_{\hat{\mathbf{k}}}(ka)^2 + O((ka)^4), \tag{4.6}$$

$$b_{\hat{\mathbf{k}}} = -\frac{1}{24}\sum_j \hat{k}_j^4 - \frac{A_1}{2I}. \tag{4.7}$$

An axis has $\sum_j \hat{k}_j^4 = 1$ and a body diagonal has $1/3$. Their difference therefore obeys

$$\boxed{b_{\text{axis}} - b_{\text{body}} = -\frac{1}{36}.} \tag{4.8}$$

This prediction cancels the isotropic reserve contribution. It survives changes in the three reserve parameters only while the identical-edge, gapped-reserve and cubic-lattice assumptions remain valid. Anisotropic couplings, gapless reserves or a different graph require a new derivation. The relation is not claimed to be unique among all lattice models.

The coefficient difference is a potentially clean laboratory target for a physically identified realization. It is not yet a bound on photons or tensor gravitational waves: the tested branch is a scalar collective mode and its dimensional lattice spacing is not independently fixed. Comparing (4.8) directly with a catalog graviton-mass limit would mix distinct dispersion hypotheses and observables.

5. Nonlinear stress campaign and independent formulations

Evidence class: Executed finite-model tests; no empirical parameter fitting..

5.1 Preparations and checks

Four six-vertex graphs are used: a ring, a path, a star and a ring with two chords. Each is evaluated at phase-amplitude scales 0.05, 0.35 and 1.2, with three fixed-seed random preparations at each scale. Every edge uses (2.2). This changes the initial-value problem and graph structure, not a particle-specific coupling.

There are 36 trajectories, each evolved for 20 model-time units at 201 recorded times. DOP853 is used with relative tolerance 10^{-10} and absolute tolerance 10^{-12}. Twelve preparations also compare direct proper-time evolution with a nonuniform reference-clock parameterization.

The finite-difference Hamiltonian derivative check is performed along a separately drawn state-space direction; it is not simply a re-evaluation of the same symbolic derivative.

The tests include total energy, the common-phase charge, a nonnegative partition, reversal of all canonical momenta, a common-phase shift, reversal of every edge orientation together with its reserve coordinates, the equilibrium Hessian, graph-mode decomposition and 50 random spatial directions in the coupled spectrum. They check distinct algebraic and numerical failure modes. Their combined count does not correspond to a number of distinct physical laws discovered.

5.2 Recorded worst discrepancies

The maximum recorded energy discrepancy is

$$\max_{\tau} \frac{|H(\tau) - H(0)|}{\max(1, |H(0)|)} = 7.895 \times 10^{-10}. \tag{5.1}$$

The denominator is the explicit normalization used by the code; for small initial energies this is not a relative fractional error. The largest common-charge drift is 1.244×10^{-14}. The maximum time-reversal state-coordinate difference is 6.914×10^{-9} and the maximum nonuniform-clock difference is 8.719×10^{-11}. These errors are numerical, in model units, and must not be advertised as experimental measurement accuracies.

The trajectories span more than one linearized preparation. Nevertheless, they do not sample all possible graphs or initial conditions, and they do not prove numerical convergence for arbitrary stiffness. Analytic energy bounds supply the general finite-model statement; these computations supply executable checks of the stated implementation.

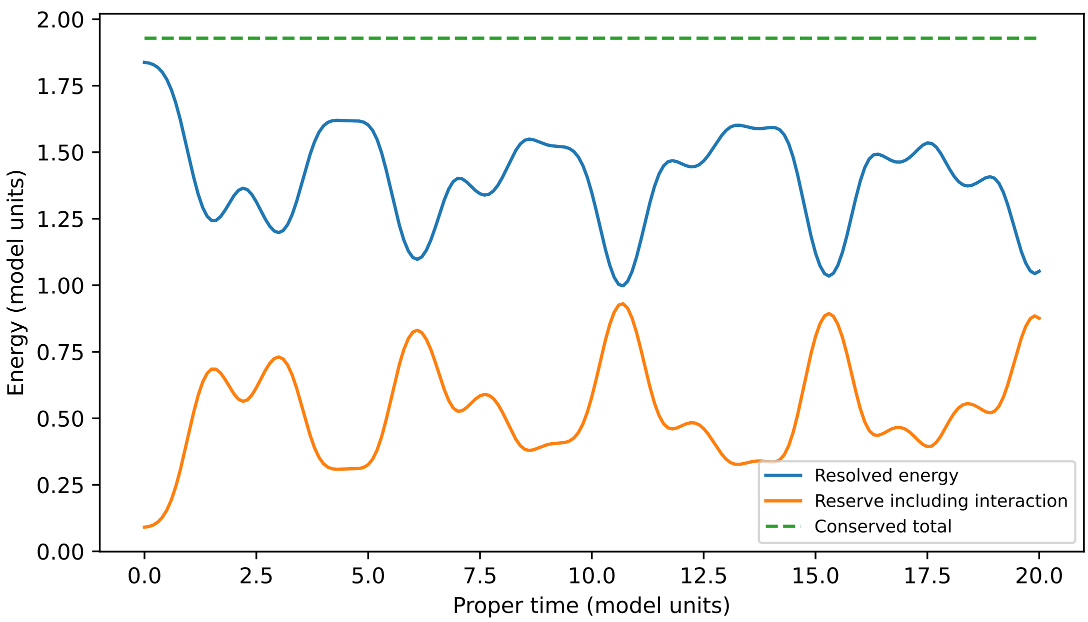

Figure 1: One representative conservative trajectory. The reserve includes its interaction contribution; neither partition is treated as energy destroyed. This is a model-unit calculation, not measured particle energy.

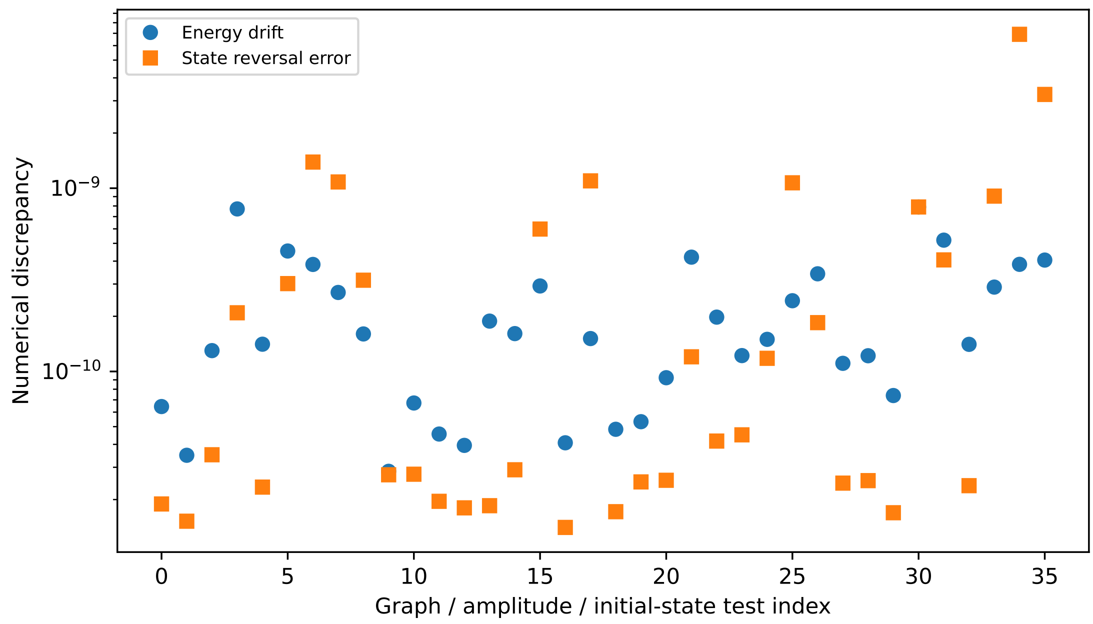

Figure 2: Numerical errors across 36 preparations. The two plotted quantities use different physical normalizations and are not percentages of theoretical accuracy.

6. Time, conversion and what an observation identifies

6.1 A rate is not a rest mass

Once a timelike clock comparison has been specified, write $g(t) = \mathrm{d}\tau/\mathrm{d}t > 0$. A resolved event-rate model can take the form

$$\lambda_{\mathrm{obs}}(t) = \epsilon_d(t)\, g(t)\, \eta(t)\, r_P(t), \tag{6.1}$$

where r_P is a predicted available conversion rate per local time, η is a conversion efficiency and ϵ_d is detector efficiency. Expected counts are the integral of this rate plus background. If any factor is unknown, the observation may constrain only their product.

Equation (6.1) is not a rule that every energy or particle mass receives a multiplicative time-dilation correction. Relativistic energy measurements are observer-dependent contractions of four-momentum; invariant mass and locally defined dimensionless couplings are different quantities. A photon has no timelike rest clock, so its propagation parameter must not be identified with a massive particle's proper time. Standard relativistic conventions are retained where used [6, 7].

Proposition 6.1 (Unidentified absolute conversion)**.** *If only λ_{obs} is measured and the normalization of r_P is unconstrained, absolute η cannot be determined from* (6.1).

Proof. For any admissible positive a, the transformation $\eta \mapsto a\eta$, $r_P \mapsto r_P/a$ preserves the product. The likelihood is constant along this direction. A numerical optimizer selecting one point on that flat family does not turn it into an identified physical measurement. □

The same issue applies to an unknown detector normalization. Independent source or preparation calibration is therefore a scientific input, not an optional cosmetic correction.

This is the reason the report does not estimate a universal EQEF number by calling every unexplained discrepancy "potential."

6.2 Finite frames and continuous rates

For a separately assumed constant proper-time hazard Γ, survival obeys $\mathrm{d}S/\mathrm{d}\tau = -\Gamma S$, giving

$$S(\Delta\tau) = e^{-\Gamma\Delta\tau}, \qquad p(\Delta\tau) = 1 - e^{-\Gamma\Delta\tau}. \tag{6.2}$$

The semigroup relation $S(u + v) = S(u)S(v)$ ensures that subdividing a frame leaves the predicted survival unchanged. A fixed fraction assigned to every arbitrary coordinate-time step does not have this property. This reconciles a universal rate with interval-dependent observed fractions, but the hazard assumption does not follow uniquely from the finite oscillatory reserve. A coherent finite system can instead exhibit return and memory.

6.3 A transformed cosmological clock does not add a fluid

For a homogeneous expansion written in proper cosmic time, $H_\tau = a^{-1}\mathrm{d}a/\mathrm{d}\tau$. Under $\mathrm{d}\tau = g\mathrm{d}t$, $H_t = gH_\tau$. The proper age remains

$$\tau_0 = \int \frac{\mathrm{d}a}{aH_\tau(a)} = \int \frac{g(a)\,\mathrm{d}a}{aH_t(a)}. \tag{6.3}$$

Changing the time label consistently does not alter a BAO distance prediction or create new cosmological energy. A physical EQEF contribution requires a specified stress-energy and exchange law. In an expanding spacetime the local covariant relation is $\nabla_\mu T^{\mu\nu} = 0$ for the total system; it is not the finite-box assertion that an arbitrary comoving volume always contains a constant energy.

7. Statistical methods and controls

7.1 Correlated data and explicit nuisance parameters

Let $d \in \mathbb{R}^n$ be a published measurement vector with covariance C, and $m(\theta)$ the specified model prediction. For a Gaussian approximation with parameter-independent covariance,

$$\chi^2(\theta) = [d - m(\theta)]^T C^{-1}[d - m(\theta)]. \tag{7.1}$$

The matrix must be symmetric and positive definite on the retained data subspace. If an exact normalization constraint makes it singular, one must remove a redundant direction or use the correctly defined constrained likelihood. Treating all normalized bins as independent would be incorrect.

The DESI covariance supplied here is positive definite, with correlated pairs at the same redshift and no nonzero cross-redshift entries. The FIRAS covariance approximation is also positive definite. Both are explicitly checked. We calculate quadratic forms by Cholesky solves rather than relying on an unstable dense inverse. For a linear model $m = X\beta$, the generalized least-squares estimator is

$$\widehat{\beta} = (X^T C^{-1} X)^{-1} X^T C^{-1} d, \qquad V_\beta = (X^T C^{-1} X)^{-1}. \tag{7.2}$$

The FIRAS implementation scales design columns before solving so that a temperature derivative and a dimensionless amplitude are not numerically compared on incompatible magnitudes. A separately evaluated normal-equation form is used as a numerical cross-check, not as a second physical measurement.

A fitted nuisance parameter is not necessarily a retrofit. For example, an experiment's calibration uncertainty or an explicitly fitted reference cosmology has a legitimate statistical role. Retrofitting would occur if such fitting were concealed and the resulting value were advertised as a target-independent prediction of the microscopic theory. All fitted quantities in this report are named.

7.2 Tail probabilities, intervals and selection

A tail probability such as $\mathbb{P}(\chi^2_\nu \geq \chi^2_{\text{obs}})$ asks how unusual that statistic would be under its stated model. It is not the probability that a theory is true, and $p = 0.5$ does not mean "50 percent accurate." The usual $n - p$ degrees of freedom is exact for suitable linear Gaussian models and only approximate for nonlinear fits or preprocessed data. Likelihood-ratio asymptotics also have assumptions and boundary exceptions [8].

The report uses exact matrix arithmetic identities and numerical tolerances to validate implementation. It uses parametric simulation to examine one nonlinear goodness-of-fit calibration. It does not use millions of decimal places as a surrogate for experimentally supported precision. Numerical precision, estimator uncertainty, model approximation and systematic uncertainty remain separate quantities.

A collection of local residuals contains expected outliers. For m valid marginal tests, a Bonferroni bound controls the familywise event by comparing the minimum marginal probability with α/m, without requiring test independence. It cannot repair incorrect marginal error models. For the CMS comparison the covariance and full theory systematics are not reconstructed, so the tabulated local-screening statistic is not promoted to a discovery or a formal rejection of the publication's analysis.

7.3 Bootstrap and held-out data have different roles

In a parametric bootstrap, simulated data are drawn from the specified fitted reference distribution and the fit is repeated. This tests the sampling behavior of the algorithm under that reference. A held-out comparison instead omits part of the actual data from parameter estimation and predicts it from the retained part. Neither makes an already public dataset historically unknown to the analyst.

For a training subset T and held-out subset H, the conditional Gaussian residual covariance is

$$C_{H|T} = C_{HH} - C_{HT} C_{TT}^{-1} C_{TH}. \tag{7.3}$$

The conditional residual mean has the corresponding $C_{HT} C_{TT}^{-1} r_T$ term. Linearized parameter uncertainty adds

$$C_{\text{pred}} = C_{H|T} + J_{\text{eff}} V_\theta J_{\text{eff}}^T, \quad J_{\text{eff}} = J_H - C_{HT} C_{TT}^{-1} J_T. \tag{7.4}$$

The DESI implementation uses this general expression, even though its selected redshift blocks have vanishing off-block covariance. Overlapping training sets make the seven resulting predictive scores dependent; their probabilities must not be multiplied.

7.4 Deliberately wrong cases

The analysis includes wrong-model and corrupted-input controls. A matter-only expansion shape is compared with the same BAO vector. An injected nonuniversal channel remains detectable after the correct clock transformation in the retained suite. Raw-file fixtures include a corrupted mass, negative energy, invalid data-quality bits, nonfinite strain and a wrong publisher checksum. A proper test system must detect these failures rather than quietly substitute convenient data.

8. DESI DR2: the full supplied covariance test

Evidence class: Executed published compressed-data analysis and a separately fitted standard comparator..

8.1 Input product and predicted quantities

The calculation uses the public GC-combined DR2 BAO vector and covariance distributed with the Cobaya BAO data products [9, 10]. There are 13 numerical measurements at seven effective redshifts. These are compressed distance measurements, not raw galaxy spectra or a full shape likelihood. The archived input vector and covariance were reused and checked against the public table; this workspace did not download a new galaxy catalog.

For a flat standard cosmological comparator define

$$E(z)^2 = \Omega_m(1+z)^3 + \Omega_r(1+z)^4 + 1 - \Omega_m - \Omega_r, \qquad A = \frac{c}{H_0 r_d}. \tag{8.1}$$

The fixed radiation parameter is $\Omega_r = 9.2 \times 10^{-5}$. The three predicted distance ratios are

$$D_M(z)/r_d = A \int_0^z \frac{\mathrm{d}u}{E(u)}, \tag{8.2}$$

$$D_H(z)/r_d = A/E(z), \tag{8.3}$$

$$D_V(z)/r_d = A\left[\frac{z}{E(z)}\left(\int_0^z \frac{\mathrm{d}u}{E(u)}\right)^2\right]^{1/3}. \tag{8.4}$$

The integrals use 80-point Gauss–Legendre quadrature and are checked against 160 points. Changing quadrature resolution is a numerical convergence check, not a change in the physical model.

8.2 The fixed benchmark is left unchanged

The fixed comparator uses $H_0 = 67.4\,\mathrm{km\,s^{-1}\,Mpc^{-1}}$, $r_d = 147.1\,\mathrm{Mpc}$ and $\Omega_m = 0.315$. These are supplied reference inputs. Evaluating (7.1) gives

$$\chi^2_{\text{fixed}} = 28.50623, \qquad p_{\text{conditional}} = 0.007686. \tag{8.5}$$

The degrees of freedom used in this conditional diagnostic is 13 because no parameter is fitted in this evaluation. The result is a mismatch of this particular supplied background, not a rejection of all flat-ΛCDM models or of every possible MNT cosmology. In particular, the microscopic Hamiltonian has not independently selected these cosmological inputs.

Seven individual entries lie within one marginal standard deviation and ten within two. These counts ignore correlation and therefore have less statistical meaning than the full quadratic form. A small marginal count discrepancy and an appreciable correlated global discrepancy can coexist without any contradiction.

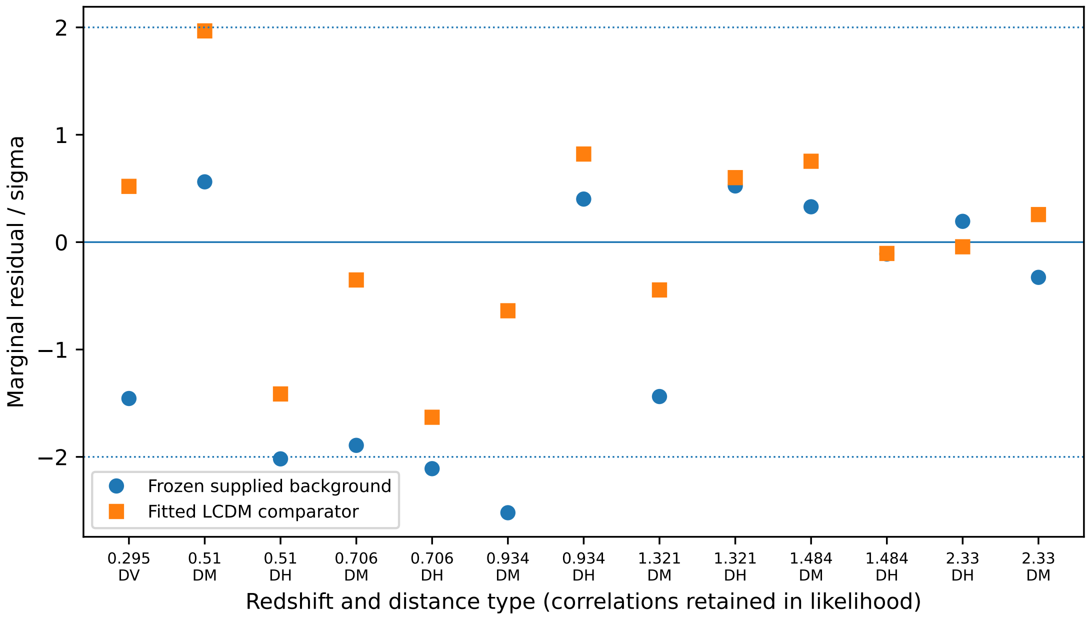

Figure 3: DESI marginal residuals for the supplied fixed comparator and the separately fitted flat-ΛCDM reference. The likelihood, unlike this visual display, uses the full covariance.

8.3 A standard reference fit, not a new MNT alignment

For fixed Ω_m, write the prediction as $Af(\Omega_m)$. The best amplitude is analytic:

$$\widehat{A}(\Omega_m) = \frac{f^T C^{-1} d}{f^T C^{-1} f}. \tag{8.6}$$

The remaining scalar minimization is performed over $0.05 \le \Omega_m \le 0.60$. The resulting fit is

$$\widehat{\Omega}_m = 0.297137 \pm 0.008575, \tag{8.7}$$

$$\widehat{A} = 29.521925 \pm 0.213072, \tag{8.8}$$

$$H_0 r_d = 10154.909\,\mathrm{km\,s^{-1}}, \tag{8.9}$$

$$\chi^2_{\text{fitted}} = 10.28401 \quad (\nu_{\text{nominal}} = 11). \tag{8.10}$$

The displayed parameter errors are local linearized errors from the two-parameter Jacobian. Their covariance is included in the JSON output. BAO alone determines a combination $H_0 r_d$ in this model; quoting separate precision measurements of H_0 and r_d from this two-parameter fit would introduce information not present in it.

The nominal tail probability is 0.505036. The substantial improvement over the fixed comparator is genuine as a standard reference fit. It is not credited to EQEF: no new microscopic calculation generated the fitted cosmology. A future MNT cosmological history would be evaluated against the same vector with its own declared parameter count and physical predictions.

8.4 Refitted bootstrap calibration

A thousand synthetic vectors are generated as

$$d^{(b)} = m(\widehat{\theta}) + L\epsilon^{(b)}, \qquad LL^T = C, \quad \epsilon^{(b)} \sim N(0, I). \tag{8.11}$$

Both reference parameters are re-estimated for each vector. The Monte Carlo tail estimate is

$$\widehat{p} = \frac{1 + \#\{\chi^2_b \geq \chi^2_{\text{obs}}\}}{1001} = 0.500500. \tag{8.12}$$

The binomial sampling uncertainty is approximately 0.016. This agrees with the nominal calibration at the resolution of the experiment. It verifies the nonlinear fitting and covariance machinery under its assumed data model. It does not test whether the assumed model is the actual microscopic origin of cosmic expansion.

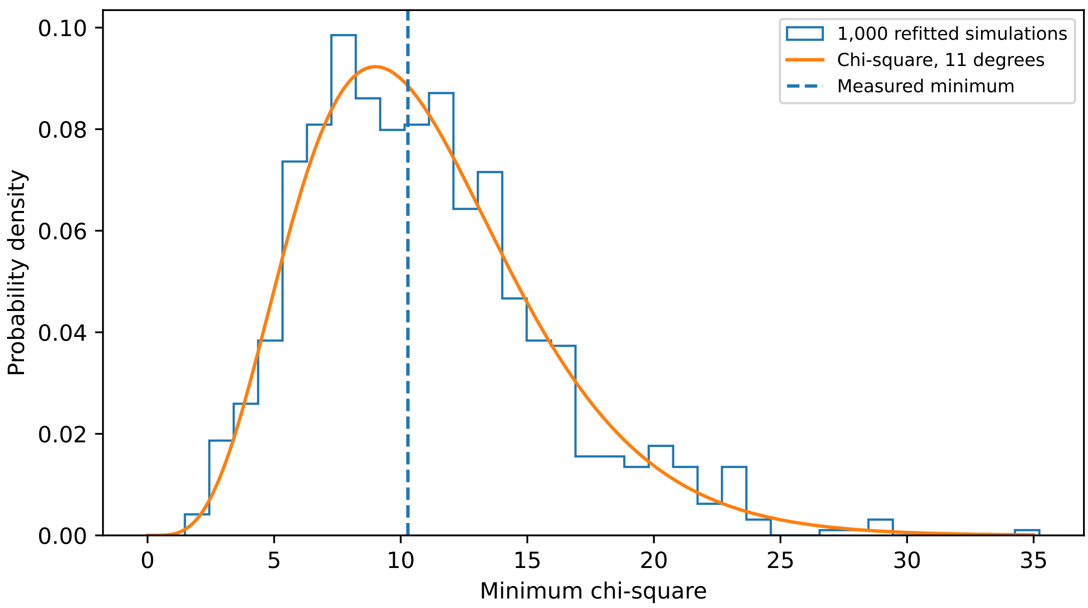

Figure 4: A thousand refitted synthetic realizations of the standard comparator. The vertical line is the statistic from the published BAO vector. These are statistical calibration trials, not additional observations.

8.5 Seven redshift-group holdouts

For each redshift the corresponding one- or two-component block is withheld. The standard comparator is fitted on the remaining measurements and used to predict the withheld block. The predictive Mahalanobis statistic uses (7.4). The lowest nominal held-out tail probability is approximately 0.10; none of these individual blocks gives a five-percent rejection within this approximation. The grouping was not a prospective experimental blind, and the seven training sets overlap.

Table 1: Leave-one-redshift-group-out comparator fits. Conditional predictive scores include linearized parameter uncertainty.

Omitted z	N_{train}	Ω_m	A	Q_{pred}	Nominal p
0.295	12	0.29446	29.4436	0.3723	0.5417
0.510	11	0.29232	29.3803	4.5799	0.1013
0.706	11	0.29928	29.6379	4.4215	0.1096
0.934	11	0.29957	29.5769	0.8891	0.6411
1.321	11	0.29863	29.5522	0.4396	0.8027
1.484	11	0.29748	29.5192	0.6788	0.7122
2.330	11	0.29887	29.5488	0.0984	0.9520

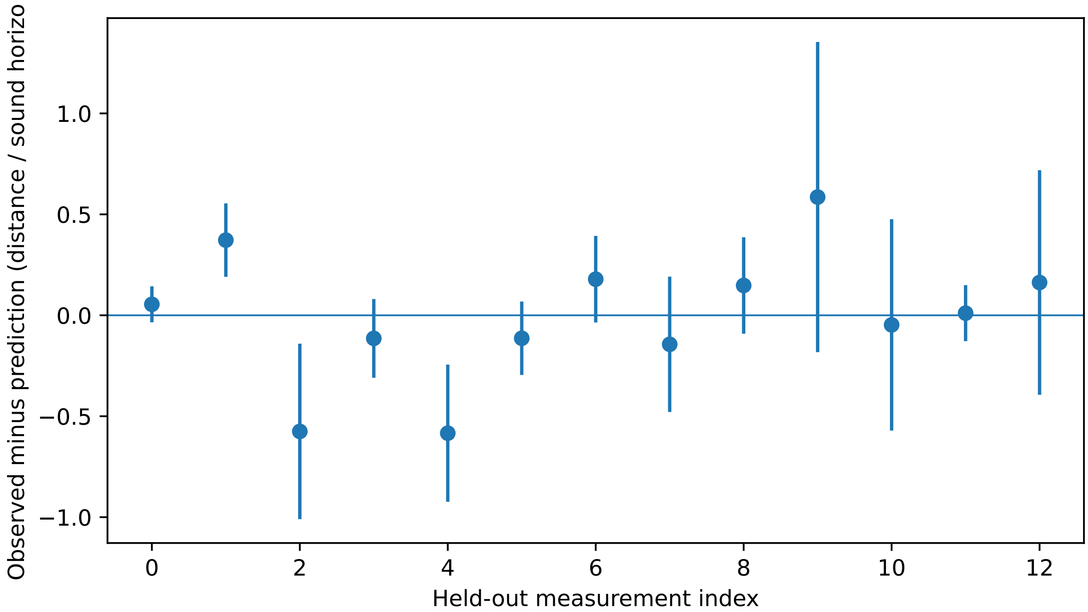

Figure 5: Held-out prediction scores by redshift. Parameter uncertainty is propagated using the local Jacobian. Overlapping training sets prevent treating these points as independent replications.

8.6 A negative control and the age calculation

A matter-only expansion shape with its amplitude fitted gives $\chi^2 \simeq 1420.18$ for the same vector. This demonstrates that the reference-fitting procedure does not make an arbitrary expansion history agree automatically. It does not select a unique dark-energy mechanism.

For the original supplied radiation-inclusive background, the proper-time age integral is

$$\tau_0 = \frac{1}{H_0}\int_0^1 \frac{\mathrm{d}a}{a\sqrt{\Omega_r a^{-4} + \Omega_m a^{-3} + \Omega_\Lambda}} = 13.7906925\,\text{Gyr}. \tag{8.13}$$

This is a dependent calculation from specified cosmological parameters. It is neither a separate age observation nor a target-independent prediction of the finite node model. The reference cosmological framework is conventional [11]; no age likelihood is added to the BAO test as if the same inputs were new data.

9. COBE/FIRAS: spectral residuals and conversion constraints

Evidence class: Executed 43-channel published-residual fits; conditional covariance approximation..

9.1 What is actually measured in the input

The NASA LAMBDA product supplies a wavenumber, a reconstructed monopole intensity, a residual, its uncertainty and a modeled Galactic spectrum at each channel. Its intensity column is explicitly formed from a 2.725 K blackbody plus the residual [12]. Fitting that constructed column and announcing a newly predicted CMB temperature would be circular. This report therefore fits the residual column, in kJy/sr, and uses temperature variation only as a nuisance spectral direction.

The covariance is constructed from the published approximation of Fixsen et al. [13],

$$C_{ij} = \sigma_i \sigma_j Q_{|i-j|}. \tag{9.1}$$

The 43 correlation-lag coefficients and the complete numerical residual input are distributed. This does not reconstruct every calibration, sky subtraction or instrument parameter of the collaboration analysis. Residual fitting may already have projected out modes, so nominal degrees of freedom and conditional errors here are diagnostics of the adopted approximation rather than replacement collaboration results.

9.2 Templates and their units

Using frequency $\nu = 100c\tilde{\nu}$ for wavenumber $\tilde{\nu}$ in cm^{-1}, define $x = h\nu/(k_B T_0)$ at $T_0 = 2.725$ K. The Planck spectrum and its temperature derivative are

$$B_\nu(T_0) = \frac{2h\nu^3}{c^2(e^x - 1)}, \qquad D_\nu = \frac{\partial B_\nu}{\partial T} = \frac{2h\nu^3}{c^2}\frac{xe^x}{T_0(e^x - 1)^2}. \tag{9.2}$$

The formula follows from thermal radiation theory [14]; it is not claimed as an MNT derivation. For table fitting, SI intensity is divided by 10^{-23} to obtain kJy/sr. The two distinct distortion templates are

$$Y_\nu = T_0 D_\nu [x \coth(x/2) - 4], \tag{9.3}$$

$$M_\nu = -T_0 D_\nu / x. \tag{9.4}$$

The chemical-potential template is accompanied by an independently fitted temperature direction. A different convention adding a multiple of D_ν gives the same fitted distortion subspace when that nuisance is free. The coefficients are fitted with either sign as local spectral perturbations; no interpretation of a negative best fit as a standalone negative-energy fluid is made.

Let G_ν be the supplied Galactic template. Four separate linear models are evaluated:

$$r_\nu = D_\nu \Delta T + g G_\nu, \tag{9.5}$$

$$r_\nu = D_\nu \Delta T + g G_\nu + y Y_\nu, \tag{9.6}$$

$$r_\nu = D_\nu \Delta T + g G_\nu + \mu M_\nu, \tag{9.7}$$

$$r_\nu = D_\nu \Delta T + g G_\nu + a_{\mathrm{grey}} B_\nu. \tag{9.8}$$

The y, μ and grey-amplitude components are not fitted jointly. Their amplitudes describe different conditional perturbations; they are not interchangeable estimates of a single universal EQEF constant.

9.3 Results and their interpretation

The separate results are

$$y = (-0.227 \pm 3.896) \times 10^{-6}, \tag{9.9}$$

$$\mu = (-13.112 \pm 36.962) \times 10^{-6}, \tag{9.10}$$

$$a_{\mathrm{grey}} = (11.863 \pm 49.458) \times 10^{-6}. \tag{9.11}$$

Each includes zero comfortably in this approximation. The baseline residual quadratic form is 49.7728; adding the y component changes it to 49.7694. A nuisance fit almost always reduces a quadratic form, so that decrease is not a discovery criterion by itself.

Table 2: Residual-template minima. Degrees of freedom and tail probabilities are conditional diagnostics of the adopted table-level approximation.

Model	χ^2	Nominal df	Nominal p
baseline	49.77284	41	0.16362
y	49.76945	40	0.13843
mu	49.64700	40	0.14106
grey	49.71530	40	0.13959

The full conditional parameter intervals are tabulated in the appendix. The published collaboration limits include a more extensive systematic analysis [13]. A smaller formal interval obtained from a compressed statistical table must not be called a stronger experiment. In particular, the nuisance ΔT here is not an absolute CMB temperature measurement at microkelvin accuracy.

9.4 Four thousand estimator-recovery trials

For the y-template design, synthetic data with fixed injected $\Delta T = 20\,\mu$K, $g = 0.015$ and $y = 10^{-5}$ are generated using the same covariance. The linear estimator is applied to 4,000 realizations. Its recovered z statistics have means near zero and widths near one. The measured 95-percent coverages are given below. They validate implementation and error propagation under the simulation assumptions; they do not validate those assumptions in nature.

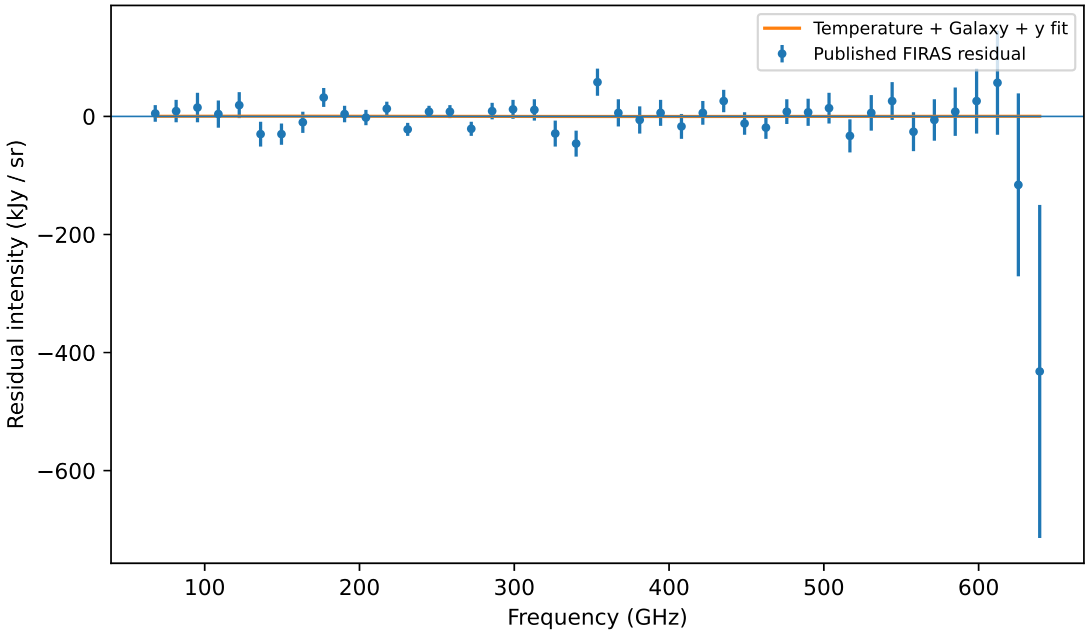

Figure 6: The published FIRAS spectral residuals and the fitted baseline / distortion model. Error bars are marginal; the fit uses the published approximate correlation structure. The curves are not observations of EQEF.

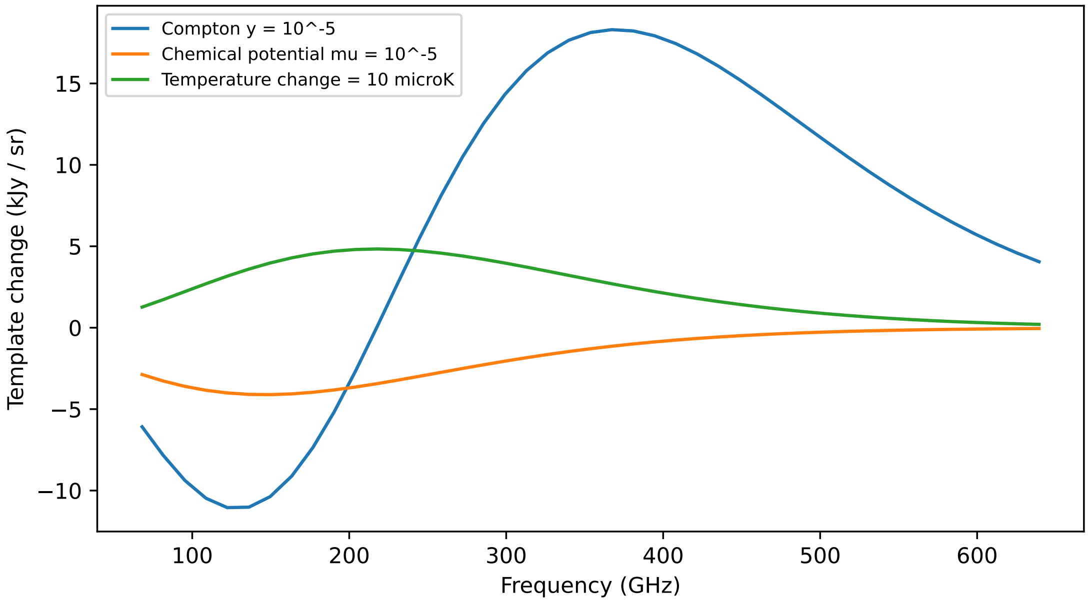

Figure 7: Unprojected physical spectral templates in intensity units, shown at the illustrative amplitudes in the legend. The separate nuisance-projected y–μ correlation is 0.824046; it is calculated in the likelihood metric, not read directly from these curves.

Table 3: 4,000 synthetic Gaussian repetitions of the same estimator, not 4,000 experimental observations.

Parameter	Mean z	SD of z	95% coverage
ΔT	0.0179	0.9787	95.275%
g	0.0246	0.9791	95.875%
y	-0.0426	0.9901	95.150%

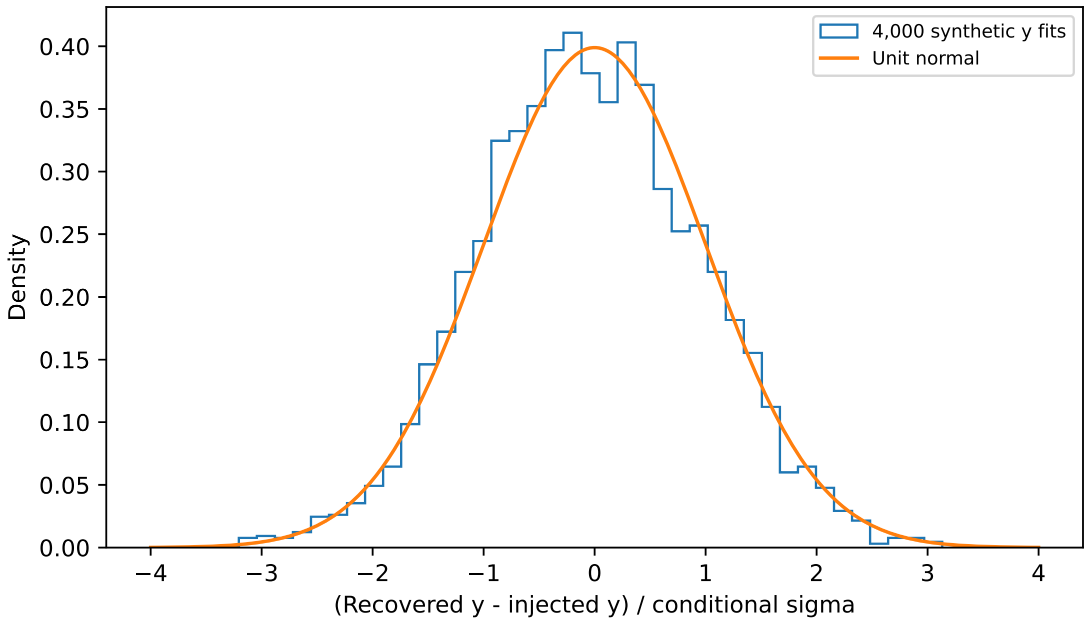

Figure 8: Synthetic coverage calibration of the conditional y estimator. The figure tests sampling behavior of the code with a known injected coefficient, not the existence of a cosmological reserve field.

9.5 What this says about EQEF

A theory predicting a concrete frequency-dependent conversion or energy-injection history can generate a spectral residual and be evaluated with (7.2) or a more complete likelihood. The current finite reserve model does not select the cosmological preparation, photon coupling or thermalization history needed to do that. Conservation of total node–reserve energy does not force the resolved photon spectrum to remain Planckian under arbitrary coupling.

A uniform grey factor is especially instructive. It can be degenerate with an absolute source amplitude or calibration. A small fitted grey coefficient is not an estimate of the fraction of all potential which actuates. To make that inference, the source spectrum and instrument transfer must be independently specified. The FIRAS analysis therefore supplies a well-defined constraint interface, not a measured absolute n_c or EQEF conversion constant.

10. CERN / CMS: 26 measured Drell–Yan ratios

Evidence class: Executed published-table comparison; no new event-level CMS analysis..

10.1 Observable and reference

The CMS product consists of 13 invariant-mass bins in each of the dimuon and dielectron channels for the 7 TeV Drell–Yan measurement. The sample corresponds to 35.9 pb^{-1}, and the tabulated ratio is normalized to the 60–120 GeV region [15]. We use the full-phase-space pre-final-state-radiation columns of Tables 8 and 9, not raw event counts. The comparison theory is the publication's Table 10 NNLO FEWZ result with its stated PDF choice and uncertainties. FEWZ was not executed in this workspace.

Write the measured quantity as

$$R_i = \frac{\sigma_i}{\sigma_{60\text{–}120}}, \qquad r_i = R_i/\Delta M_i. \tag{10.1}$$

The graph shows r_i per GeV so bins of unequal width can be compared. The statistical calculation uses the R_i values and their errors consistently; an exact linear change of units would leave a full-covariance likelihood unchanged.

10.2 Marginal uncertainty treatment

Let R_i^{ref} be the tabulated FEWZ prediction. The experimental uncertainty is the publication's statistical-plus-systematic quadrature error. A directional reference error is formed by adding the relevant PDF and "other" percentage errors in quadrature. The resulting diagnostic pull is

$$z_i = \frac{R_i^{\mathrm{obs}} - R_i^{\mathrm{ref}}}{\sqrt{\sigma_{i,\mathrm{exp}}^2 + \sigma_{i,\mathrm{ref,dir}}^2}}. \tag{10.2}$$

For an observed value above the prediction, the upward reference uncertainty is used, and conversely below. This is a transparent marginal approximation, not a complete nuisance-parameter likelihood. Cross-bin errors are correlated through common normalization, efficiencies, unfolding and theory inputs. Accordingly, the code records a diagonal sum of squared pulls but deliberately assigns it no formal global p value.

10.3 Results, including the local discrepancy

Eleven of 13 muon bins and all 13 electron bins lie within two conditional standard deviations of the stated reference, for 24 of 26 in total. The largest marginal discrepancy is the muon 96–106 GeV bin: its ratio is 0.041 ± 0.003, versus the published reference 0.052. It gives approximately -3.463 in (10.2) and is retained without adjustment.

A local residual is not by itself a new-physics result. The comparison does not reconstruct the full correlated experimental likelihood, and its reference values and uncertainties are rounded published outputs. A Bonferroni screen of the 26 marginal diagnostics is recorded in the machine-readable file, but no global discovery claim or rejection of the collaboration's full analysis follows. The scientifically useful next step for that discrepancy would be a full covariance and systematics treatment, not a fitted EQEF correction to that one bin.

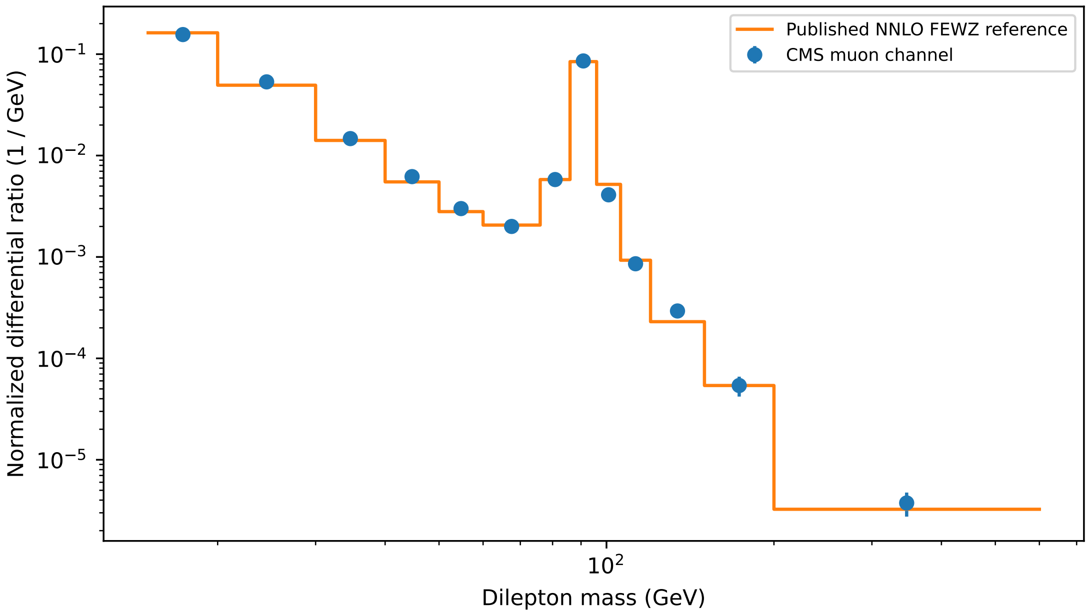

Figure 9: Published CMS dimuon normalized spectrum compared with the publication's FEWZ reference. This is a real published measurement-product comparison, not a reconstruction of 100,000 event records and not a microscopic MNT cross-section prediction.

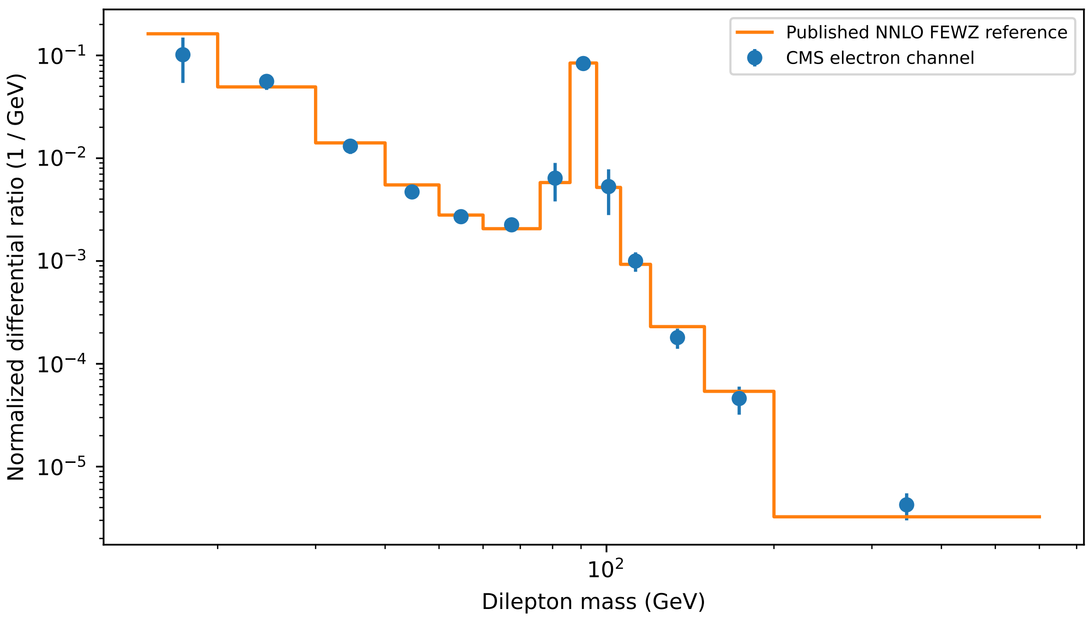

Figure 10: The separate dielectron channel, in the same convention. Agreement of two conventional reference comparisons is not two independent determinations of an EQEF parameter.

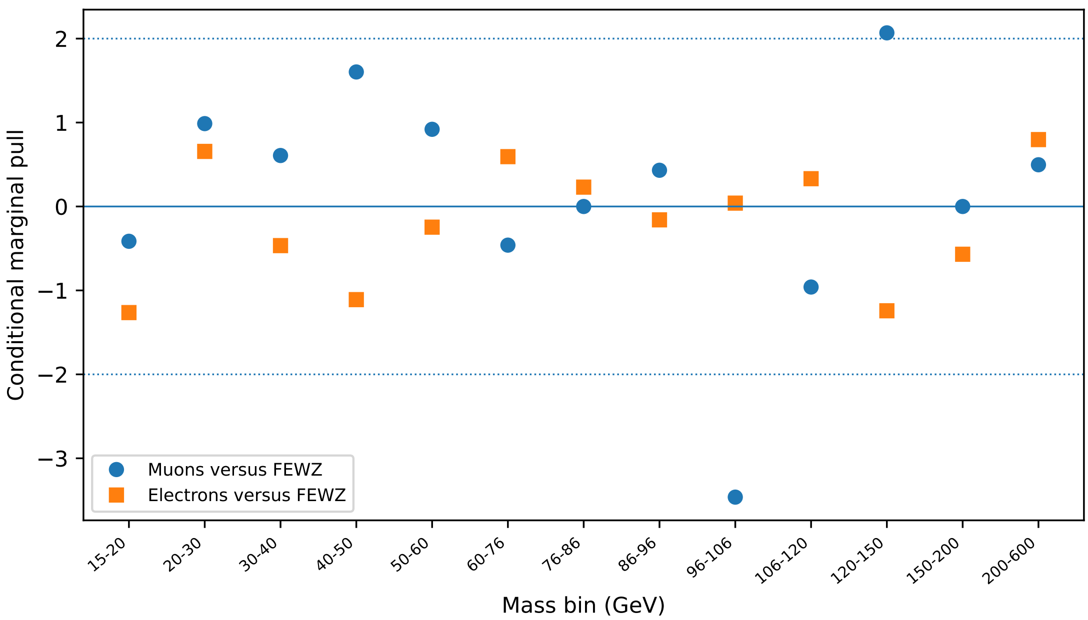

Figure 11: All 26 conditional marginal pulls. The largest local discrepancy is shown rather than removed. Correlation and shared systematic limitations prevent turning this graph into a universal accuracy percentage.

10.4 An exact identifiability result for a universal efficiency

Suppose the only change to every mass bin is a common multiplicative conversion efficiency η, with the same factor in the normalization window. Then

$$R_i' = \frac{\eta\sigma_i}{\eta\sigma_{60\text{–}120}} = R_i. \tag{10.3}$$

The code checks the cancellation across several trial efficiencies. Therefore this normalized spectrum cannot identify an absolute common efficiency. It can test a predicted mass-dependent shape change, or a species-dependent change after a physically justified response model, but it cannot determine a universal absolute conversion factor merely because the ratios agree.

This is a substantive inference result for the proposed validation strategy. A single universal microscopic law is a valuable requirement, yet a chosen observable can be insensitive to its overall scale. To test that scale requires an absolute production or decay observable with independently known source normalization, luminosity and response. Such quantities must be supplied by a physical model, not inferred from the same discrepancy being tested.

10.5 Cross-channel comparison

The muon and electron ratios are compared using the quadratic sum of their marginal errors. Twelve of 13 bins lie within two such standard deviations; the largest difference is about 2.267 standard deviations in the 120–150 GeV interval. This test uses the publication's limited inter-channel-correlation approximation and is not a microscopic derivation of lepton universality. Known charge assignments, weak currents and source spectra remain ingredients of the conventional calculation.

Frequency-dependent actuation may eventually generate a distinguishable line shape or rate. At present, the finite scalar model supplies no uniquely identified quark current, hadron parton distributions, leptonic matrix element or detector acceptance. Calling the published FEWZ curve an MNT output would conceal those missing physical identifications rather than solve them.

11. XENONnT: limits, dimensions and a declared interaction

Evidence class: Published-limit reconstruction and conditional operator translation..

11.1 The input is an exclusion product, not a list of recoil events

The retained XENONnT input contains ten heavy-mediator spin-independent limits and ten light-mediator limits at trial masses from 3 to 12 GeV. The collaboration publishes its light-dark-matter analysis and associated numerical products for external use [16, 17]. These two sets of rows describe different interaction assumptions applied to the experiment. They must not be counted as twenty independent observations of a particle.

The numerical audit reconstructs the power-constrained upper limit as

$$\sigma_{\mathrm{PC}}(m) = \max\left[\sigma_{\mathrm{unconstrained}}(m), \sigma_{\mathrm{sensitivity},-1}(m)\right] . \tag{11.1}$$

All twenty rows reproduce their archived limit values within the stated numerical tolerance. This establishes faithful manipulation of that product. It is not a new event-level likelihood evaluation, a reproduction of the collaboration's detector calibration, or the discovery of a dark-matter signal.

11.2 A scale translation that can actually be reproduced

For the heavy-mediator rows only, define a nucleon-level contact-interaction convention in natural units,

$$\sigma_N = \frac{\mu_N^2}{\pi\Lambda^4}, \qquad \mu_N = \frac{m_\chi m_N}{m_\chi + m_N}, \qquad m_N = 0.939\ \mathrm{GeV}. \tag{11.2}$$

This equation defines the normalization used in this translation; it is not a uniquely selected MNT operator. It assumes the stated isospin-conserving spin-independent heavy-mediator interpretation. Since a cross section in natural units has dimension GeV^{-2},

$$\sigma_N[\mathrm{cm}^2] = \frac{\mu_N[\mathrm{GeV}]^2}{\pi\Lambda[\mathrm{GeV}]^4}\left(1.973269804\times 10^{-14}\right)^2. \tag{11.3}$$

The resulting lower scale associated with a published upper cross-section limit is

$$\Lambda_{\min}(m_\chi) = \left[\frac{\mu_N^2(\hbar c)^2}{\pi\sigma_{N,\max}}\right]^{1/4} . \tag{11.4}$$

Each scale is substituted back into (11.2) to verify the original limit. Across these ten masses the translated scale runs from about 1.644 to 31.305 TeV. The full table is in Appendix A. This wide change is a mass-dependent experimental sensitivity, not evidence that the fundamental scale changes from particle to particle.

The light-mediator table carries a different convention, with tabulated units $\mathrm{cm}^2\mathrm{MeV}^4$. Its entries are not inserted into (11.4). Treating the two tables as one cross-section column would create an artificial numerical agreement through a units error.

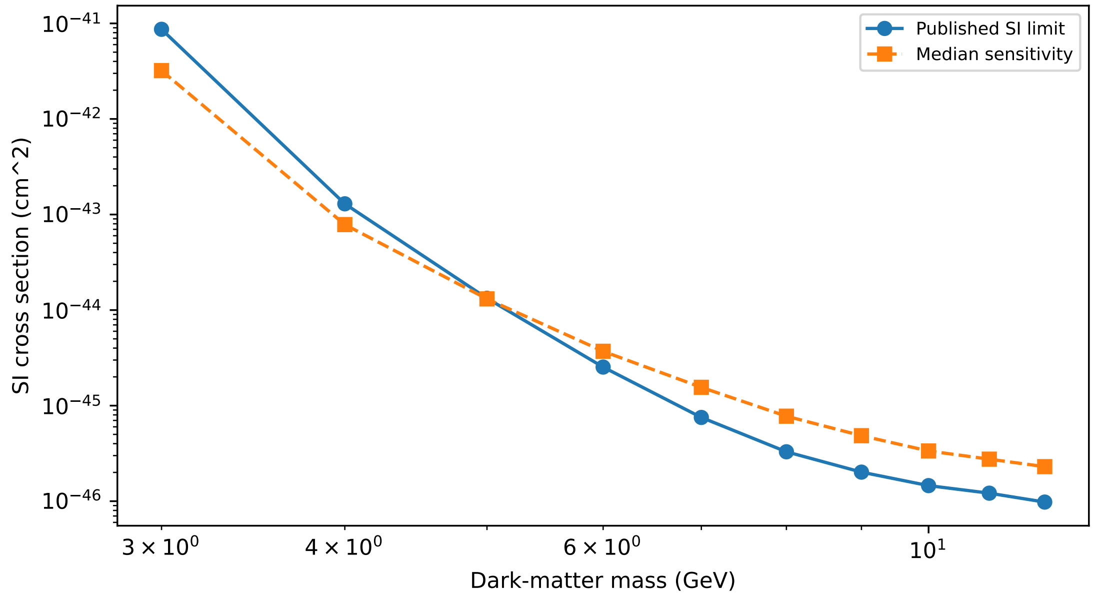

Figure 12: The retained XENONnT heavy-mediator spin-independent upper limits and median sensitivity. These are collaboration limit products, not a simulated MNT signal. The companion table translates the limits through one explicitly defined contact-interaction normalization.

11.3 What would turn this into an MNT recoil test

A physical hypothesis would have to predict a mass, an abundance or local density, a velocity distribution, and a differential interaction with the detector. In a simple elastic-scattering description, the source rate has the schematic form

$$\frac{\mathrm{d}R}{\mathrm{d}E_R} = \frac{\rho_\chi}{m_\chi m_T} \int_{v \geq v_{\min}} \mathrm{d}^3 v\, f_{\mathrm{lab}}(\mathbf{v})\, v \frac{\mathrm{d}\sigma_T}{\mathrm{d}E_R}. \tag{11.5}$$

A detector prediction then requires the response kernel and efficiency,

$$\lambda_b = \mathcal{E} \int \mathrm{d}E_R \frac{\mathrm{d}R}{\mathrm{d}E_R} \int_b \mathrm{d}x\, \epsilon(x, E_R)\, p(x \mid E_R) + \lambda_{b,\mathrm{background}}. \tag{11.6}$$

Here $\mathcal{E}$ is the exposure and x denotes measured detector variables. None of these terms can be replaced by the name "EQEF." In particular, energy which is stored in a reserve in the finite Hamiltonian is not automatically a gravitationally clustering particle population with a nuclear recoil interaction.

The useful result of this section is a precise interface: a future MNT prediction in the convention (11.2) can be compared with the table without changing the limit calculation. A different operator must be treated with the appropriate collaboration recasting procedure and response, rather than by rescaling the strongest point on this curve.

12. Clocks, muons and universal gravitational response

Evidence class: Published-summary comparisons with a retained conventional reference..

12.1 Why these are seven comparisons, not one combined likelihood

The seven summary inputs concern different observables and different modeling assumptions. Their uncertainties and systematics are not supplied as one joint covariance. We therefore report each standardized residual separately,

$$z = \frac{O_{\text{measured}} - O_{\text{reference}}}{\sqrt{\sigma^2_{\text{measured}} + \sigma^2_{\text{reference}}}}, \tag{12.1}$$

with the denominator adjusted to the specific supplied information. There is no pooled probability of MNT being correct. A result close to zero demonstrates compatibility of that stated reference and measurement at this level of approximation.

12.2 Optical clocks and gravitational redshift

For a small height difference in a weak approximately uniform gravitational field, the retained metric relation is

$$\frac{\Delta \nu}{\nu} \simeq \frac{g\Delta h}{c^2}. \tag{12.2}$$

With $g = 9.8\,\mathrm{m\,s^{-2}}$ and $\Delta h = 0.33\,\mathrm{m}$, this gives 3.5983×10^{-17}. The selected Chou et al. optical-clock result, $(4.1 \pm 1.6) \times 10^{-17}$, gives a standardized residual +0.314 [18]. These are externally supplied geometric conditions and an established weak-field law. This comparison checks the observation interface; it does not independently establish that the finite node graph generates the gravitational potential.

The Galileo redshift comparison uses the published violation parameter $(0.19 \pm 2.48) \times 10^{-5}$ from Delva et al. [19]. The conventional null gives $z = +0.0766$. It constrains an anomalous redshift contribution in that parameterization. It does not measure an absolute conversion efficiency, because the observation has already been reduced through an orbit and clock model.

12.3 Moving muon lifetimes

The retained lifetime comparison combines the MuLan rest lifetime with the rounded Lorentz factor used for the older moving-muon example [20, 21]:

$$\tau_{\text{lab}}^{\text{ref}} = \gamma\tau_0 = 29.33\,(2.1969803\,\mu\text{s}) = 64.437432199\,\mu\text{s}. \tag{12.3}$$

The selected moving positive-muon summary is $64.419 \pm 0.058\,\mu$s. Propagating the stated rest-lifetime error gives $z = -0.318$. The calculation does not reconstruct the historical beam momentum distribution, detector acceptance or correlated systematics. In particular, rounding γ is not an uncertainty model. The agreement is an illustrative consistency check, not a replacement for the original experiment's inference.

If an additional proper-time conversion mechanism were to alter a decay law, it would need to supply a rate contribution with units $\mathrm{s^{-1}}$ and the relevant state dependence. Merely comparing clocks does not determine that contribution. A coordinate-time exponential with a modified rate can represent the same proper-time decay after a change of clock variable, so the invariant quantity and measurement convention must be fixed before interpreting a residual.

12.4 Universal free fall

The MICROSCOPE comparison uses

$$\eta_{\text{Eotvos}} = (-1.5 \pm 2.3_{\text{stat}} \pm 1.5_{\text{syst}}) \times 10^{-15} \tag{12.4}$$

from the experiment's final result [22]. Combining the two stated errors in quadrature for a Gaussian diagnostic gives $z = -0.546$ relative to universal free fall. This convention is recorded; it does not claim that all systematic effects are intrinsically random Gaussian variables.

A universal metric coupling is compatible with this result. A reserve coupled differently to different matter compositions would require a separate calculation of the composition-dependent acceleration. The fact that a scalar finite model conserves energy is not enough to determine that acceleration. Conservation and universal coupling are distinct requirements.

12.5 The muon magnetic anomaly and a changing reference calculation

The comparison is frozen to the 2025 experimental world average and the specified 2025 Standard Model reference [23, 24]. In units of 10^{-11},

$$a_\mu^{\text{exp}} = 116592071.5 \pm 14.5, \tag{12.5}$$

$$a_\mu^{\text{SM,ref}} = 116592033 \pm 62. \tag{12.6}$$

Their difference is 38.5 ± 63.673, or $z = +0.605$. This reference uses a lattice-based treatment of the leading hadronic vacuum-polarization contribution. The choice matters: the theory compilation discusses tensions among hadronic inputs, and a historical discrepancy is not carried forward as an unquestioned current target.

No new MNT loop calculation was performed. An EQEF contribution to a_μ would require a specified interaction, a renormalization convention, and a calculation of the relevant vertex form factor. It cannot be obtained by setting an undefined conversion deficit equal to the measured minus predicted anomaly. That procedure estimates a residual by definition rather than predicting its cause.

12.6 GW170817 and electromagnetic propagation

The retained interval from the joint gravitational-wave and gamma-ray observation is

$$-3 \times 10^{-15} \leq \frac{c_g - c}{c} \leq 7 \times 10^{-16}, \tag{12.7}$$

under the publication's emission-time assumptions [25]. The conventional null lies inside it. The code assigns no Gaussian p value to this interval. This is not a fresh analysis of the detector strain or an independent measurement of the emission delay.

The scalar acoustic branch in Section 4 is not identified with the tensor gravitational wave in this observation. Such an identification would first require deriving the physical mode, its coupling to matter and the detector, and its propagation law on the relevant background. The interval is a valuable requirement for that construction, not a completed bridge on its own.

13. Neutrino mixing: measured summaries and forward spectra

Evidence class: Two published-summary benchmarks and 120 conventional propagation forecasts..

13.1 Observed oscillations and the input convention

The benchmark has normal ordering, with

$$(\theta_{12}, \theta_{13}, \theta_{23}) = (33^\circ, 8.5^\circ, 45^\circ), \quad (\Delta m^2_{21}, \Delta m^2_{31}) = (7.5 \times 10^{-5}, 2.5 \times 10^{-3})\,\mathrm{eV}^2. \quad (13.1)$$

These values were supplied to the conventional propagation example. They were not generated by the node eigenvalue problem. The Daya Bay summary used here reports $\sin^2 2\theta_{13} = 0.0851 \pm 0.0024$ and, for normal ordering, $\Delta m^2_{32} = (2.466 \pm 0.060) \times 10^{-3}\,\mathrm{eV}^2$ [26]. The benchmark comparisons give $z = -0.159$ and $+0.683$, respectively. The second uses $\Delta m^2_{32} = \Delta m^2_{31} - \Delta m^2_{21}$, not a silent substitution of one splitting for another.

13.2 A unitary three-flavor calculation

Let $s_{ij} = \sin\theta_{ij}$ and $c_{ij} = \cos\theta_{ij}$. In the standard convention,

$$U = \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta} & c_{23}c_{13} \end{pmatrix}. \quad (13.2)$$

The mass-squared matrix relevant for constant-density matter propagation is

$$M_f^2 = U\,\mathrm{diag}(0, \Delta m^2_{21}, \Delta m^2_{31})U^\dagger + \mathrm{diag}(A, 0, 0), \qquad A = 2\sqrt{2}G_F n_e E. \quad (13.3)$$

The matter term is the conventional coherent forward-scattering contribution [27, 28]. In the numerical unit convention,

$$A\,[\mathrm{eV}^2] = 1.526 \times 10^{-4}\, Y_e \rho\,[\mathrm{g\,cm^{-3}}]\,E\,[\mathrm{GeV}], \qquad Y_e = 0.5. \quad (13.4)$$

For antineutrinos, U is conjugated and A changes sign. A common mass-squared shift contributes only an overall phase and is removed. The evolution operator is

$$S(L,E) = \exp\left[-i\,2(1.26693268)\frac{L\,[\mathrm{km}]}{E\,[\mathrm{GeV}]}M_f^2[\mathrm{eV}^2]\right], \qquad P_{\mu\to\beta} = |S_{\beta\mu}|^2. \quad (13.5)$$

The factor two is essential: the familiar $1.267\Delta m^2 L/E$ appears inside a two-flavor probability's sine, whereas the amplitude phase contains twice that quantity.

The code evaluates both a matrix exponential and an independent Hermitian eigendecomposition. For the latter, $M_f^2 = V\,\mathrm{diag}(\lambda_j)V^\dagger$ gives

$$S = V\,\mathrm{diag}\big(e^{-i\kappa\lambda_j}\big)V^\dagger. \quad (13.6)$$

Hermiticity ensures unitarity, and hence $\sum_\beta P_{\mu\to\beta} = 1$. The two solvers agree to at most 4.99×10^{-15} across the stated grid, and probability conservation passes its numerical tolerance.

13.3 What was forecast

The grid contains five CP phases, neutrinos and antineutrinos, six energies and two densities: $5 \times 2 \times 6 \times 2 = 120$ points. The baseline is 1,300 km; densities are 0 and 2.8 g cm^{-3}. The complete CSV includes appearance, survival and tau-channel probabilities, solver differences and unitarity errors.

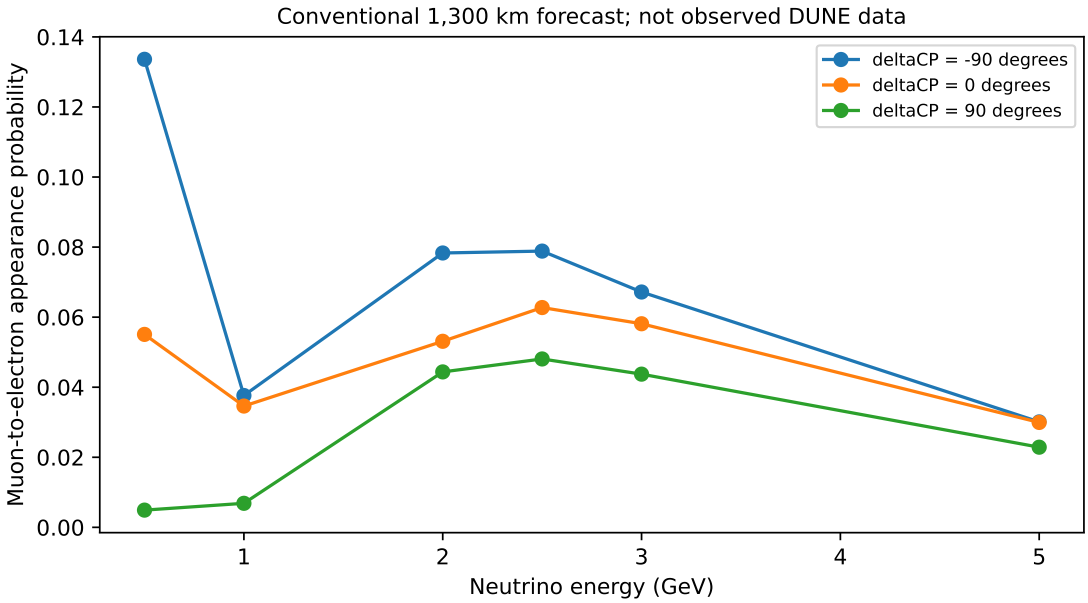

Figure 13: A subset of the conventional constant-density 1,300 km forecasts. CP phase, mass differences and mixing angles are inputs. Lines connect calculated grid points and do not represent observed DUNE far-detector event spectra.

A probability is not an event spectrum. A DUNE prediction would also require flux, cross sections, energy migration, efficiencies, exposure and backgrounds [29, 30]. Public prototype detector measurements and test-beam data are real data, but they are not the future far-detector oscillation sample. This edition does not label a probability forecast as an already measured DUNE validation.

For MNT, the required additional step is to determine which stable modes carry the appropriate weak flavor content, why their mixing matrix has this form, and which eigenvalue differences enter (13.3). Frequency-dependent actuation is compatible with posing that spectral problem, but it does not numerically solve it.

14. GWOSC raw-strain pipeline: supplied, fixture-tested, not run on real files

Evidence class: Executable local reconstruction method; synthetic verification only in this workspace..

14.1 Acquisition and provenance

GWOSC supplies the GW150914 Hanford and Livingston 32-second, 4,096 Hz strain records, together with an illustrative numerical-relativity waveform [31, 32]. The pipeline pins the two HDF5 filenames and the publisher's MD5 checksums, then records SHA-256 digests for any files actually processed. Files are downloaded atomically to a temporary suffix, subject to a size cap. A failed acquisition exits with an error and never substitutes simulated strain.

Both acquisition routes available in this workspace failed. The local failure log records the network-resolution error. Consequently, there are no measured raw-strain SNRs, time delays, waveform likelihoods or dispersion limits reported here. The following describes the supplied executable method and the synthetic controls which were actually run, not a completed real-detector analysis.

14.2 Conditioning and timing diagnostic

The loader reads strain, start time, sample spacing and data-quality metadata. It requires matching detector grids, finite data and valid available DATA flags. This initial check is not a full search-quality veto prescription. The method follows the distinction between raw strain, conditioning and inference emphasized in the collaboration's data-analysis guide [33].

The power spectral density is estimated from off-event intervals after excluding the vicinity of the known signal, using median-averaged Welch segments. A Tukey window controls spectral leakage. The whitened display is bandpassed from 35 to 350 Hz, and the cross-correlation is evaluated for delays inside a physical ±10 ms interval. If $H(t)$ is the Hanford series and $L(t)$ the Livingston series, the diagnostic compares $H(t)$ with $L(t-\ell)$ over the fixed event window. A positive ℓ therefore means that the Hanford pattern arrives later.

The largest absolute correlation is reported with its sign. Searching over lag and applying conditioning changes the null distribution, so the correlation coefficient is not called a detection significance. Off-event shifted-window maxima are saved as diagnostics. These overlapping windows are correlated and do not provide a calibrated independent false-alarm-rate measurement.

14.3 A fixed-template phase profile

The optional diagnostic applies a cubic Fourier-phase deformation,

$$\widetilde{h}_{\beta,t_0,\phi,A}(f) = A\widetilde{h}_0(f)\exp\left[\mathrm{i}\beta(f/f_*)^3 - 2\pi\mathrm{i}ft_0 + \mathrm{i}\phi\right], \qquad f_* = 150\text{ Hz}. \tag{14.1}$$

The template h_0 is the public illustrative waveform, which was scaled and adjusted using this known event. It is not a waveform predicted from the finite MNT Hamiltonian, nor an unseen signal selected before observation.

With a one-sided noise spectrum $S_n(f)$, use the standard noise-weighted inner product

$$(a|b) = 4\Re\int_{f_{\min}}^{f_{\max}} \frac{\widetilde{a}(f)\widetilde{b}(f)^*}{S_n(f)}\,\mathrm{d}f. \tag{14.2}$$

For fixed (β, t_0), define the complex matched product Z and template norm N_h. Maximizing over a common quadrature phase and nonnegative amplitude gives

$$\widehat{\phi} = \arg Z, \qquad \widehat{A} = |Z|/N_h, \qquad \rho^2_{\text{profile}} = |Z|^2/N_h. \tag{14.3}$$

The code profiles over $-2 \leq \beta \leq 2$ radians at f_* and a ± 0.02 s time-offset grid, separately for each detector. It saves the difference between the maximum ρ^2 and each profile value as a *conditional* $\Delta\chi^2$ diagnostic. A best fit at a grid boundary would not establish a finite interval. No automatic confidence-limit claim is made from these curves.

In a physical source analysis, masses, spins, orientation, distance, calibration, waveform systematics and PSD uncertainty must be treated consistently. Cubic propagation phase is also correlated with source phasing. Holding a known-event template fixed can make a conditional interval look much stronger than a physically marginalized one. The pipeline is therefore a transparent first reconstruction and sensitivity diagnostic, not a substitute for that full inference.

14.4 Connection to the lattice relation, under an extra physical hypothesis

A branch with $\omega = ck[1 + b(ka)^2 + \cdots]$ has, to leading order,

$$k(\omega) = \frac{\omega}{c} - b\frac{a^2\omega^3}{c^3} + \cdots . \tag{14.4}$$

With the Fourier convention of (14.1), propagation through a static distance D contributes the additional phase

$$\delta\varphi = b\frac{Da^2\omega^3}{c^3}, \qquad \beta = b\frac{Da^2(2\pi f_*)^3}{c^3}. \tag{14.5}$$

This is a conditional map, not an assertion that the scalar branch is a gravitational wave. Cosmological propagation also requires a redshift-dependent integral rather than an unqualified Euclidean distance. The sign and scale of a phase bound only become an MNT test after the physical branch, preferred-frame convention if present, lattice scale and source model are fixed.

14.5 End-to-end controls actually executed

The synthetic HDF5 fixture includes a deliberately shifted signal and opposite detector sign. Its injected Hanford-after-Livingston delay is 0.0068359375 s; the conditioning and lag search recover the same grid value, with correlation −0.989394. This tests timing convention and parser behavior, not the measured delay of GW150914.

A separate noiseless frequency-domain injection has $\beta = 0.35$, $t_0 = 0.004$ s and amplitude 2×10^{-20}. The profile recovers those grid values to floating-point precision. Because this injection is noiseless and on the search grid, exact recovery tests the algebra and implementation rather than frequentist uncertainty coverage. Rejection controls include an invalid checksum, nonfinite strain and bad data-quality flags.

15. The local CMS event-reconstruction method

Evidence class: Executable parser and invariant-mass reconstruction; synthetic fixture tests only for the large-sample pipeline..

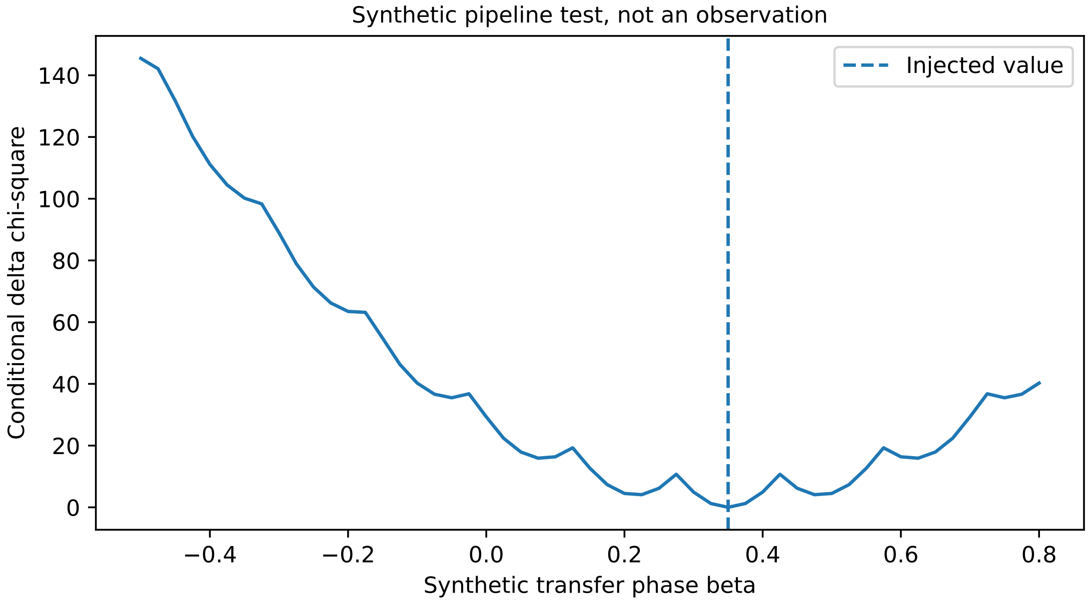

Figure 14: Synthetic fixed-template phase recovery with an injected cubic coefficient $\beta = 0.35$. The vertical scale is the conditional profile discrepancy. The input here is deliberately generated, not real GWOSC strain.

15.1 What a four-vector calculation can establish

The local runner accepts an actual CSV containing the two measured lepton four-vectors. In units with $c = 1$, the pair invariant mass is

$$M^2 = (E_1 + E_2)^2 - (p_{x1} + p_{x2})^2 - (p_{y1} + p_{y2})^2 - (p_{z1} + p_{z2})^2. \tag{15.1}$$

Recomputing this quantity can reveal a column-order, units, parsing or rounding error. It is not a prediction of the rest mass of a particle from the microscopic theory. The CMS educational dimuon release is genuine recorded data, but its documentation explicitly limits its suitability for a full CMS physics analysis [34].

The metadata-based downloader only accepts CSV URLs actually supplied by the CERN record. A missing CSV is reported as an acquisition failure; no plausible-looking filename is invented. A user may instead supply a CSV obtained from the official data product. The source path, checksum, number of records and validation status are written into the result.

15.2 Decimal precision and interval reconstruction

For a printed decimal value with last represented place 10^e, the rounding interval is taken as the central value plus or minus $10^e/2$. This assumes rounding to nearest at that printed precision; a source using truncation or exact integer coordinates needs a different convention. Interval sums are formed for the two energies and the three momentum components.

If the total energy is in $[E_-, E_+]$ and a momentum component is in $[p_-, p_+]$, its squared interval is

$$p^2 \in \left[\begin{cases} 0, & p_- \le 0 \le p_+, \\ \min(p_-^2, p_+^2), & \text{otherwise}, \end{cases} \quad \max(p_-^2, p_+^2) \right]. \tag{15.2}$$

Combining these bounds gives a conservative interval for (15.1). The lower and upper mass bounds are then compared with the rounding interval of a supplied mass column. This avoids declaring an otherwise consistent event wrong merely because its rounded four-vector components produce a slightly different central mass.

An event whose entire M^2 interval is negative is flagged as unphysical under the stated convention. A small negative central value with an interval extending above zero is distinguished from that case. Positive finite lepton energies and finite momentum components are required. Duplicate run/event identifiers are reported rather than automatically deleted, because several candidates can belong to one event and require an explicit selection policy.

15.3 Executed controls and physical limits

The fixture contains 1,200 generated two-lepton rows. All 1,200 pass the declared rounding-interval overlap test. A corrupted reported mass is detected, and an invalid negative-energy input is rejected. These checks contribute to the 23 raw-pipeline fixture assertions. They are not real collisions and are not counted as new external measurements.

A publication-grade collider signal test would additionally specify event selection, luminosity, efficiencies, backgrounds, detector response, theory cross sections, nuisance parameters and the statistical test. The measured Drell–Yan comparison in Section 10 supplies a useful published-data check now. The local event runner supplies an auditable entry point for real event records without pretending that reconstruction alone establishes an MNT resonance law.

16. Structural tests: frequency, gauge symmetry and gravity

16.1 A mode frequency is one property, not a complete particle identity

The physical proposal treats persistent particles as stable actuated patterns, with frequency and coherence affecting which patterns can form. This is a useful reason to calculate the spectrum of a single Hamiltonian rather than insert one adjustable potential per particle. The present coupled spectrum is a concrete implementation of that first step.

A spectral value alone does not determine charge, spin, statistics, chirality or a scattering amplitude. Even if an effective quantum description assigns $E = \hbar\omega$, the identification of E with a particle rest energy requires a rest-frame excitation with the relevant dispersion and conserved quantum numbers. The numerical value of $\hbar$ is not generated by replacing a classical frequency with an energy label. The atomic and mass identities retained in the package therefore keep their supplied-input status [35, 36].

This distinction also prevents an ambiguity in the phrase “observed, then actuated.” A detector registration is a physical interaction with a record-producing system. It is not, without an additional dynamical law, a claim that conscious observation produces the particle. The reciprocal threshold construction in the inherited model includes an apparatus-like degree of freedom and its backreaction. Whether its statistics reproduce a real quantum detector remains a separate test.

16.2 An exact scalar-link test

Consider the candidate link constructed only from differences of a single-valued node phase,

$$U_{ij} = e^{\mathrm{i}(\theta_i - \theta_j)}. \tag{16.1}$$

Around a closed oriented loop,

$$U_{12}U_{23}\cdots U_{n1} = \exp[\mathrm{i}\{(\theta_1 - \theta_2) + \cdots + (\theta_n - \theta_1)\}] = 1. \tag{16.2}$$

This is a pure-gradient connection. The stress suite verifies this identity for fifty random phase loops. Such a construction does not by itself produce independent plaquette curvature. Singular configurations, nontrivial topology or independent link variables would change the assumptions and must be stated explicitly.

For a non-Abelian gauge theory, the link variables and generators carry additional structure; in particular, the generators do not generally commute [37, 38, 39]. Merely grouping three scalar frequencies does not establish the eight-generator $\mathfrak{su}(3)$ algebra, and grouping two frequencies does not establish weak-isospin representations. A valid emergence calculation would derive the low-energy constrained degrees of freedom, their transformation law and effective action from the frozen microscopic system.

The established electroweak theory supplies a concrete target for the representation and interaction content [40, 28]. Chiral lattice matter has additional consistency requirements; the hypotheses and scope of the Nielsen–Ninomiya result are one reason these cannot be bypassed by naming an antisymmetric mode a fermion [41]. Equation (16.2) is a useful localized diagnosis, not a proof that all possible node models are incapable of gauge emergence.

16.3 Quantum composition and probability

The inherited quantum-channel examples are rerun. Their positivity and normalization checks are conditional on a Hilbert-space description, density matrices and specified completely positive maps. Those are standard mathematical structures [42, 43]. Verifying their implementation is valuable but does not derive their physical necessity from the classical reserve Hamiltonian.

A microscopic account of quantum measurement would also have to confront the observed correlation structure. The Bell and CHSH constraints distinguish classes of local hidden-variable models under their explicit assumptions [44, 45]; deterministic evolution alone is not a resolution. The report does not assign a probability law merely because phase dynamics can be complicated or chaotic. A preparation distribution, composition rule and measurement map are needed before comparing predicted outcome frequencies with a Bell experiment.

16.4 Propagation geometry versus autonomous gravity

A continuum wave equation can often be expressed using an effective metric. This establishes how one selected mode propagates on a medium or background; it does not automatically establish an autonomous Einstein equation for that background [46]. In the present construction, clock reparameterization and acoustic dispersion are calculated. A derivation of the physical gravitational field would additionally identify dynamical geometric variables, the constraint equations, matter coupling and the relevant continuum limit.

An assumed effective Einstein–Hilbert action can be varied consistently and compared with observations. Doing so is a conditional recovery of established gravity, not a derivation of that action from the finite six-node system. Thermodynamic approaches to gravity also require explicit physical assumptions [7, 47]. These distinctions preserve the genuine finite-model achievements without turning an analogy into a completed gravitational theory.

17. Twenty coverage areas and ten decisive requirements

17.1 Coverage is not a count of confirmations

The campaign covers twenty named areas at different levels. Some contain newly executed measurement-product analyses; others contain finite-model theorems, retained calculations, forecasts or an explicitly unexecuted raw-data pipeline. The detailed scope table is in Appendix C. A field with no specified physical prediction is not assigned an invented goodness-of-fit score.

The three main new measurement products contain 82 values: 13 BAO measurements, 43 spectral residuals and 26 collider ratios. Their correlations, common normalizations and reference-model dependencies matter. Counting them as 82 independent MNT predictions would erase the most important information about the analysis.

The complete execution ledger records 1684 successful software assertions. This includes checks that enforce an expected failure, verify a symmetry identity, reject a corrupted file, or establish consistency between two numerical methods. Their success means that the stated software criteria were met. It is not a posterior probability of the microscopic theory.

Table 4: Recorded assertions by executed suite. Repeated scalar checks are not independent physical experiments.

Suite	Origin	Passed	Total
desi	new implementation	15	15
firas	new implementation	20	20
cms	new implementation	13	13
stress	new implementation	376	376
constraints	new implementation	291	291
raw_fixture	new implementation	23	23
baseline	retained suite	172	172
retained_audit	retained suite	23	23
clock_conversion	retained suite	325	325
conservative_reconstruction	retained suite	277	277
threshold_response	retained suite	82	82
coupled_spectrum	retained suite	67	67

17.2 Ten requirements which an eventual physical completion must satisfy

Table 5: Ten make-or-break requirements, with the exact achievement of this edition.

No.	Requirement	Result and boundary
1	Closed, conservative dynamics	Explicit finite Hamiltonian, global finite-graph evolution and conserved quantities are established under stated assumptions. Nonlinear implementations pass their tolerances.
2	Stable and physically admissible states	The tested equilibrium spectra and energy partition are nonnegative. This does not prove stability of every nonlinear preparation or identify the observed matter spectrum.
3	Consistent clocks and observed relativity	Time reparameterization is verified; selected clock summaries are compatible with the retained metric reference. Emergence of that metric is not established.
4	One microscopic law across preparations	The same edge coefficients are used for every new graph and initial state. No particle-by-particle coefficients are fitted. The graph and physical species map remain separate model choices.
5	Quantum probability and correlations	Conditional quantum-channel calculations are verified. The probability rule and the observed quantum correlation structure are not uniquely derived from the finite classical model.
6	Correct gauge structure and chiral matter	The scalar pure-gradient loop identity is verified. Non-Abelian curvature, representations, anomaly consistency and observed chiral matter still require a constructive physical map.
7	Particle masses and interactions	No new fundamental mass or coupling is claimed. The CMS comparison uses the publication's conventional reference, not an MNT matrix element.
8	Universal gravitational response	Selected free-fall, redshift and propagation summaries are compatible with the conventional metric null. A microscopic Einstein limit and compact-object solution are not newly derived here.
9	Cosmological history and spectra	The fixed BAO comparator is strained; a fitted standard comparator is compatible. FIRAS residual shapes are fitted with declared nuisances. No microscopic cosmic history or primordial perturbation spectrum is selected.

No.	Requirement	Result and boundary
10	Distinctive, independent prediction	A directional lattice relation is derived conditionally. No new independent parameter-free MNT physical alignment has been observed. Large raw-detector analyses are not counted as completed.

These requirements are useful because they separate a failed particular assumption from a missing physical identification. A finite conservative model can pass Requirement 1 while Requirement 6 remains unestablished. Conversely, a physical branch that predicts a definite excluded dispersion cannot be protected by pointing to its energy conservation. Each claim must survive the test appropriate to it.

18. What would make the next prediction decisive

18.1 A complete prediction includes its measurement map

A comparison becomes distinctive when the model determines an observable which differs from the reference, with numerical choices fixed independently of the target. In the notation of (1.2), the required object is not only H but the complete chain

$$(H, \vartheta, \text{preparation}) \longrightarrow O_{\text{source}} \longrightarrow O_{\text{detector}} \longrightarrow \mathcal{L}(\text{data} \mid \vartheta). \tag{18.1}$$

The same microscopic ϑ must be used across domains, while independently measured apparatus properties remain legitimate nuisance inputs. “One law” does not mean that a dimensionful event rate has the same numerical value in every particle preparation, detector, time unit or environment.

The first physical test need not determine all Standard Model parameters. A dimensionless ratio, directional difference or functional relation can be especially useful when unknown scales cancel. But the cancellation has to follow from the equations, and the physical field to which it applies must be identified before measurement.

18.2 The directional relation as a concrete example

Equation (4.8) is a possible model-level target because the same reserve correction appears in both directions and cancels. Suppose a physical wave branch and the relevant local propagation directions were independently identified. After determining the long-wavelength speed and using a common wave-number convention, a measured difference of quadratic coefficients could test the value $-1/36$.

This would test the declared cubic, identical-edge, gapped-reserve realization. It would not establish every version of MNT, nor show uniqueness among all lattice models. A physical anisotropy search must also specify frame orientation, averaging, source behavior and instrumental directional effects. Those choices cannot be selected after finding a promising anisotropy in the data.

There is a sharp mathematical falsification within the model: changing the two directions while preserving the assumptions cannot change the analytic coefficient difference. The present fifty-direction verification checks that statement numerically. The remaining work for an experimental claim is the apparatus and scale identification, not another decimal-place match of the same identity.

18.3 Calibrated predictions can still be informative

A single independent calibration of a dimensional scale does not invalidate every subsequent prediction. It changes the evidentiary label. The calibration observation must not also be counted as a prediction, and uncertainty in that scale must be propagated to held-out observables.

The DESI analysis illustrates the separation. Its fitted A and Ω_m produce a legitimate standard-model goodness-of-fit and held-out residual assessment. They do not become a first-principles cosmological prediction by being inserted into a section titled EQEF. Likewise, a measured source flux or detector efficiency can calibrate an event-rate prediction, but an arbitrary source normalization inferred from the desired residual cannot establish an absolute conversion fraction.

18.4 Precommitment, blinding and hashes

The package records a frozen inherited model, an analysis plan, exact input hashes, code hashes and dated execution outputs. These records make later changes visible. A hash proves identity of bytes relative to a supplied digest; it does not prove scientific correctness, historical priority or that an analyst had never seen a public measurement.

The CMS, DESI and FIRAS products were already public and their broad behavior was known. These comparisons are therefore not described as historically blind discoveries. The holdouts and simulation controls have narrower, stated roles. A future genuinely prospective test should timestamp the complete physical prediction and evaluation rule before the relevant data are available, while keeping software-bug corrections and new model versions separate.

19. Reproduction, independent checking and safe publication

19.1 Reproduce the executed analysis without a paid AI service

The companion package contains the numerical tables actually used, executable Python scripts, frozen-input manifests and the complete assertion ledger. After installing the declared numerical dependencies in an ordinary local Python environment, the default command is

```
python -m pip install -r requirements.txt
python run_validation.py
```

The numerical stage then requires no network access or API key. Library installation may require an internet connection. The shipped environment used Python 3.13.5; exact numerical-library versions are recorded. Different BLAS implementations or library versions can change last floating-point digits, so comparisons use numerical tolerances rather than a claim of byte-identical output on every platform [48].

The report-generation workspace enforces a short per-command limit. The final execution was therefore performed in two explicit stages,

```
python run_validation.py --skip-prior
python run_validation.py --prior-only
```

Both stages completed. The new suites ran from 01:07:06 to 01:07:27 UTC on 7 September 2026; the retained suites ran from 01:07:32 to 01:07:53 UTC. The machine-readable stage records distinguish fresh execution from the runner's optional collection-only mode. Splitting the process changes neither its input data nor its scientific equations.

19.2 Run the larger raw-file pipelines locally

On a computer with working access to the official repositories, the optional commands are

```
python code/raw_detectors.py gw150914 --fetch
python code/raw_detectors.py cms --fetch
```

If CERN's record metadata does not supply a directly usable CSV, obtain an appropriate file from the official release and invoke

```
python code/raw_detectors.py cms --file /path/to/real_dimuon.csv
```

The program deliberately fails rather than inventing a URL or generating replacement collisions. It verifies the required columns before analysis. Raw results are written into a separate directory and cannot silently replace the published OV-1.0 results. A new real-data run should receive its own timestamp, provenance record and interpretation.

A successful local download is only the beginning of a detector analysis. Before using its outputs in a paper, inspect the quality flags, selection, source assumptions, frequency or mass range and systematic treatment described above. A real HDF5 file plus a numerical fit does not automatically provide a valid physical confidence interval.

19.3 A practical independent audit

An independent reviewer can check the work in three steps. First, compare the supplied data with the external products and verify the declared row selection and units. Second, reproduce one central calculation by a different method, such as solving the DESI covariance system independently or checking the FIRAS template dimensions. Third, trace each claimed physical inference back through its nuisance assumptions and model-to-observable map.

The code, tables and log files make this possible without accepting an AI-generated narrative on trust. A polished document is not itself a validation certificate, and an automated test count is not peer review. Conversely, the presence of a clearly recorded limitation is not evidence that a numerical calculation was never performed. Each executed result has its own inputs and output artifact.

19.4 Website wording and version control

The appropriate website heading is *Open-Data Testing and Reproducibility*. The subtitle can state: *Published measurement comparisons, conservative-model stress tests and executable* *detector pipelines.* The page should link the report, source package, scope table and assertion ledger together.

The page should also preserve the single essential qualification: the large raw GWOSC and CMS pipelines were fixture-tested but not run on real files in this edition. Readers should not have to discover that only after interpreting a headline as a new detector result. Future physical data analyses can be added as new versioned releases without rewriting the history of this one.

20. Conclusions

The new campaign verifies the specified conservative model across a broader set of nonlinear preparations and supplies three substantive published-data analyses. The finite-model checks establish conservation, symmetry, clock covariance and spectral consequences within their assumptions. The observation analyses establish a covariance-aware BAO comparison, a conditional CMB spectral-residual fit and a published collider-spectrum comparison with its local discrepancy retained.

The fixed cosmological comparator gives $\chi^2 = 28.50623$, while the separately fitted flat-ΛCDM reference gives $\chi^2 = 10.28401$. FIRAS's separate distortion fits remain compatible with zero within their table-level approximation. Twenty-four of 26 CMS marginal comparisons lie within two conditional standard deviations, with the largest discrepancy about 3.463σ. Selected clock, muon and neutrino summaries are compatible with their declared conventional references. These are concrete results, not a single global accuracy percentage.

The record contains 1684 passed assertions and 5,000 new statistical calibration replications. Those numbers describe the executed work, not independent experimental confirmations. The precise unresolved scientific task is to select a physical realization and observable map that turn the microscopic model into a distinctive prediction. The report supplies verified mathematical tools, real published constraints and a reproducible way to expose that prediction to failure when it is specified. It does not replace the missing physical identification with a fitted residual or a claim of completion.

A. Numerical inputs and results sufficient to reconstruct the comparisons

A.1 DESI vector, residuals and covariance

The following table lists every measurement used in the 13-component comparison. "Fixed" refers to the supplied background, while "fit" refers to the separately optimized standard flat-ΛCDM comparator. Marginal residuals are shown for inspection; the statistical calculation uses the full covariance, not the sum of squared entries in these marginal columns.

Table 6: All 13 DESI input values, predictions and marginal pulls. The likelihood uses the full supplied covariance.

z	Quantity	Observed	σ	Fixed	Fit	Fixed pull
0.295	D_V/r_d	7.94168	0.07609	8.05248	7.90216	-1.46
0.510	D_M/r_d	13.58758	0.16837	13.49304	13.25636	+0.56
0.510	D_H/r_d	21.86295	0.42887	22.72855	22.46930	-2.02
0.706	D_M/r_d	17.35069	0.17993	17.69123	17.41408	-1.89
0.706	D_H/r_d	19.45535	0.33387	20.15966	19.99984	-2.11
0.934	D_M/r_d	21.57564	0.16178	21.98331	21.67902	-2.52
0.934	D_H/r_d	17.64149	0.20104	17.56088	17.47655	+0.40
1.321	D_M/r_d	27.60086	0.32456	28.06750	27.74537	-1.44
1.321	D_H/r_d	14.17602	0.22455	14.05852	14.04114	+0.52
1.484	D_M/r_d	30.51190	0.76356	30.26046	29.93677	+0.33
1.484	D_H/r_d	12.81700	0.51801	12.87421	12.87161	-0.11
2.330	D_H/r_d	8.63155	0.10106	8.61198	8.63592	+0.19
2.330	D_M/r_d	38.98897	0.53168	39.16335	38.85248	-0.33

The symmetric covariance is specified by its upper-triangular nonzero elements. Unlisted elements are zero, and the lower triangle is obtained by symmetry. Indices follow the preceding measurement table. This compact representation reconstructs every element of the archived 13×13 matrix.

Table 7: Nonzero upper-triangular entries of the DESI covariance, in the row order of the input table. Mirror entries below the diagonal; all unlisted entries vanish.

i	j	z_i	z_j	C_{ij}
1	1	0.295	0.295	0.00578998687
2	2	0.510	0.510	0.0283473742
2	3	0.510	0.510	-0.0326062007
3	3	0.510	0.510	0.18392804
4	4	0.706	0.706	0.0323752442
4	5	0.706	0.706	-0.0237445646
5	5	0.706	0.706	0.111469198

Continued

i	j	z_i	z_j	C_{ij}
6	6	0.934	0.934	0.0261732816
6	7	0.934	0.934	-0.0112938006
7	7	0.934	0.934	0.0404183878
8	8	1.321	1.321	0.105336516
8	9	1.321	1.321	-0.0290308418
9	9	1.321	1.321	0.0504233092
10	10	1.484	1.484	0.583020277
10	11	1.484	1.484	-0.195215562
11	11	1.484	1.484	0.268336193
12	12	2.330	2.330	0.0102136194
12	13	2.330	2.330	-0.0231395216
13	13	2.330	2.330	0.282685779

A.2 FIRAS residual channels

The second intensity column in the NASA source is deliberately not treated as an independent absolute blackbody measurement. The table below preserves the residual, uncertainty and modeled Galactic contribution actually used. Wavenumber is converted to ordinary frequency using $\nu = c\,\widetilde{\nu}$, with the factor 100 converting cm^{-1} to m^{-1}. The likelihood operates on residual intensity in kJy/sr.

Table 8: FIRAS numerical residual input. Frequency is wavenumber; the three intensity columns are in kJy/sr. The separate reconstructed blackbody column is not fitted as independent temperature data.

$\tilde{\nu}$ (cm^{-1})	Residual	σ	Galaxy
2.27	5	14	4
2.72	9	19	3
3.18	15	25	-1
3.63	4	23	-1
4.08	19	22	3
4.54	-30	21	6
4.99	-30	18	8
5.45	-10	18	8
5.90	32	16	10
6.35	4	14	10
6.81	-2	13	12
7.26	13	12	20
7.71	-22	11	25
8.17	8	10	30
8.62	8	11	36

Continued

$\tilde{\nu}$ (cm^{-1})	Residual	σ	Galaxy
9.08	-21	12	41
9.53	9	14	46
9.98	12	16	57
10.44	11	18	65
10.89	-29	22	73
11.34	-46	22	93
11.80	58	23	98
12.25	6	23	105
12.71	-6	23	121
13.16	6	22	135
13.61	-17	21	147
14.07	6	20	160
14.52	26	19	178
14.97	-12	19	199
15.43	-19	19	221
15.88	8	21	227
16.34	7	23	250
16.79	14	26	275
17.24	-33	28	295
17.70	6	30	312
18.15	26	32	336
18.61	-26	33	363
19.06	-6	35	405
19.51	8	41	421
19.97	26	55	435
20.42	57	88	477
20.87	-116	155	519
21.33	-432	282	573

For channel indices $i, j = 0, \dots, 42$, the covariance approximation is

$$C_{ij} = \sigma_i \sigma_j Q_{|i-j|}. \tag{A.1}$$

The complete published rounded lag sequence used here is

$$\begin{aligned}(Q_0,\dots,Q_{42}) = (&1.000,\ 0.176,\ -0.203,\ 0.145,\ 0.077,\ -0.005,\ -0.022,\\ &0.032,\ 0.053,\ 0.025,\ -0.003,\ 0.007,\ 0.029,\ 0.029,\\ &0.003,\ -0.002,\ 0.016,\ 0.020,\ 0.011,\ 0.002,\ 0.007,\\ &0.011,\ 0.009,\ 0.003,\ -0.004,\ -0.001,\ 0.003,\ 0.003,\\ &-0.001,\ -0.003,\ 0.000,\ 0.003,\ 0.009,\ 0.015,\ 0.008,\\ &0.003,\ -0.002,\ 0.000,\ -0.006,\ -0.006,\ 0.000,\ 0.002,\ 0.008).\end{aligned}$$

These coefficients are from the covariance approximation described in Section 3.3 of Fixsen et al. [13]. The code explicitly checks positive definiteness after rounding. An invertible table-level covariance is not a reconstruction of all calibration, foreground and preprocessing uncertainties.

The conditional fit parameters are given below. The central intervals are estimates plus or minus 1.959964 conditional standard errors. They are not the collaboration's final confidence intervals. The y and μ columns come from separate fits; the table is not evidence that all listed distortions were simultaneously identified.

Table 9: Conditional GLS template fits. Temperature changes are in microkelvin; distortion and grey-amplitude coefficients are in parts per million. Galaxy coefficients are dimensionless.

Fit	Parameter	Estimate	1σ	95% lower	95% upper
baseline	ΔT (μK)	0.3179	9.5763	-18.4514	19.0872
baseline	g	-0.0007	0.0302	-0.0599	0.0586
y	ΔT (μK)	0.4188	9.7317	-18.6551	19.4926
y	g	0.0006	0.0370	-0.0719	0.0731
y	$10^6 y$	-0.2268	3.8958	-7.8625	7.4088
mu	ΔT (μK)	-8.9924	27.9380	-63.7500	45.7651
mu	g	0.0023	0.0314	-0.0592	0.0639
mu	$10^6 \mu$	-13.1121	36.9625	-85.5572	59.3330
grey	ΔT (μK)	-7.8345	35.3106	-77.0421	61.3731
grey	g	0.0018	0.0319	-0.0608	0.0644
grey	$10^6 a_{\text{grey}}$	11.8633	49.4579	-85.0723	108.7989

A.3 CMS normalized-ratio comparisons

Every pre-FSR full-phase-space experimental ratio used in the calculation is included here, together with the publication's FEWZ reference and the conditional marginal pull. Ratios are dimensionless. The differential spectrum in the figures divides each ratio by its mass-bin width.

Table 10: All 26 CMS pre-FSR normalized-ratio comparisons. Ratios and errors are multiplied by 10^3. Pulls include directional marginal FEWZ uncertainty; no full-covariance global test is claimed.

Channel	Mass (GeV)	$10^3 R$	$10^3 \sigma$	FEWZ	Pull
muon	15–20	780	69	812	-0.415
muon	20–30	533	34	494	+0.987
muon	30–40	147	8	141	+0.607
muon	40–50	62	4	55	+1.603
muon	50–60	30	2	28	+0.920
muon	60–76	32	2	33	-0.461
muon	76–86	58	3	58	+0.000
muon	86–96	857	26	844	+0.432
muon	96–106	41	3	52	-3.463
muon	106–120	12	1	13	-0.960
muon	120–150	8.8	0.9	6.9	+2.069
muon	150–200	2.7	0.6	2.7	+0.000
muon	200–600	1.5	0.4	1.3	+0.497
electron	15–20	508	238	812	-1.264
electron	20–30	559	97	494	+0.656
electron	30–40	131	21	141	-0.466
electron	40–50	47	7	55	-1.110
electron	50–60	27	4	28	-0.246
electron	60–76	36	5	33	+0.594
electron	76–86	64	26	58	+0.231
electron	86–96	834	60	844	-0.159
electron	96–106	53	25	52	+0.040
electron	106–120	14	3	13	+0.331
electron	120–150	5.4	1.2	6.9	-1.242
electron	150–200	2.3	0.7	2.7	-0.569
electron	200–600	1.7	0.5	1.3	+0.797

A.4 Published summaries and conditional XENON translation

The summary table lists standardized residuals and two-sided Gaussian diagnostics. Measurements, uncertainties and conventional references are specified in Sections 12 and 13 and in the companion CSV. Correlated systematic effects not supplied in the source summaries are not reconstructed by this table.

Table 11: Seven published-summary interfaces. The last two use supplied neutrino benchmark parameters; none is a new microscopic MNT prediction.

Comparison	Pull	Two-sided diagnostic p
Optical height shift	+0.3136	0.7539
Moving muon lifetime	-0.3178	0.7506
Galileo redshift parameter	+0.0766	0.9389
MICROSCOPE differential fall	-0.5463	0.5849
Muon anomaly / 2025 reference	+0.6047	0.5454
Daya Bay mixing benchmark	-0.1588	0.8738
Daya Bay splitting benchmark	+0.6833	0.4944

Only the heavy-mediator XENON rows enter the contact-scale conversion. All scales below refer to the explicitly defined normalization in (11.2); they are not predicted MNT particle masses.

Table 12: Conditional contact-scale translation of the heavy-mediator spin-independent limit. The operator convention is fixed in the text; it is not a microscopic coupling determination.

m_χ (GeV)	σ_N^{90} (cm^2)	μ_N (GeV)	$\Lambda_{\min}$ (TeV)
3	8.674×10^{-42}	0.71516	1.6442
4	1.294×10^{-43}	0.76048	4.8510
5	1.328×10^{-44}	0.79054	8.7386
6	2.532×10^{-45}	0.81193	13.4024
7	7.523×10^{-46}	0.82794	18.3317
8	3.289×10^{-46}	0.84036	22.7134
9	2.015×10^{-46}	0.85029	25.8253
10	1.455×10^{-46}	0.85840	28.1453
11	1.213×10^{-46}	0.86515	29.5726
12	9.788×10^{-47}	0.87086	31.3047

A.5 Finite-model preparations

The next table contains all 36 new nonlinear preparations. The microscopic edge coefficients are identical across the table. Initial conditions and graph connectivity vary according to the declared plan. The energy error is $\max_t |H(t) - H(0)|/\max(1, |H(0)|)$ in model units. The charge error is absolute in the stated model convention. These numerical measures are not experimental uncertainty estimates.

Table 13: All 36 conservative stress trajectories. Energy error is normalized by $\max(1, |H_0|)$; it is not always a fractional energy error. Reversal is a maximum absolute state-coordinate discrepancy.

Graph	Amp.	Rep.	H_0	Energy error	Reversal error
ring6	0.05	0	0.12895	6.44×10^{-11}	1.89×10^{-11}
ring6	0.05	1	0.09619	3.49×10^{-11}	1.52×10^{-11}
ring6	0.05	2	0.17226	1.30×10^{-10}	3.51×10^{-11}
ring6	0.35	0	1.92857	7.71×10^{-10}	2.09×10^{-10}
ring6	0.35	1	0.60581	1.41×10^{-10}	2.34×10^{-11}
ring6	0.35	2	1.67454	4.54×10^{-10}	3.02×10^{-10}
ring6	1.2	0	4.93111	3.84×10^{-10}	1.39×10^{-9}
ring6	1.2	1	5.35127	2.70×10^{-10}	1.08×10^{-9}
ring6	1.2	2	3.10160	1.60×10^{-10}	3.15×10^{-10}
path6	0.05	0	0.09098	2.85×10^{-11}	2.73×10^{-11}
path6	0.05	1	0.13177	6.73×10^{-11}	2.75×10^{-11}
path6	0.05	2	0.12260	4.56×10^{-11}	1.96×10^{-11}
path6	0.35	0	0.23459	3.95×10^{-11}	1.80×10^{-11}
path6	0.35	1	0.76539	1.88×10^{-10}	1.85×10^{-11}
path6	0.35	2	0.59328	1.61×10^{-10}	2.91×10^{-11}
path6	1.2	0	3.12315	2.93×10^{-10}	5.99×10^{-10}
path6	1.2	1	0.12432	4.08×10^{-11}	1.41×10^{-11}
path6	1.2	2	4.78140	1.51×10^{-10}	1.10×10^{-9}
star6	0.05	0	0.09446	4.84×10^{-11}	1.72×10^{-11}
star6	0.05	1	0.10410	5.32×10^{-11}	2.50×10^{-11}
star6	0.05	2	0.12554	9.25×10^{-11}	2.55×10^{-11}
star6	0.35	0	0.66233	4.21×10^{-10}	1.20×10^{-10}
star6	0.35	1	0.76249	1.98×10^{-10}	4.17×10^{-11}
star6	0.35	2	0.19652	1.22×10^{-10}	4.51×10^{-11}
star6	1.2	0	2.06383	1.50×10^{-10}	1.18×10^{-10}
star6	1.2	1	3.70677	2.43×10^{-10}	1.07×10^{-9}
star6	1.2	2	3.36091	3.41×10^{-10}	1.85×10^{-10}
chord6	0.05	0	0.18929	1.11×10^{-10}	2.46×10^{-11}
chord6	0.05	1	0.20957	1.22×10^{-10}	2.54×10^{-11}
chord6	0.05	2	0.21984	7.39×10^{-11}	1.70×10^{-11}
chord6	0.35	0	2.64580	7.89×10^{-10}	7.90×10^{-10}
chord6	0.35	1	2.21132	5.21×10^{-10}	4.06×10^{-10}
chord6	0.35	2	0.44614	1.41×10^{-10}	2.38×10^{-11}
chord6	1.2	0	3.37968	2.89×10^{-10}	9.05×10^{-10}
chord6	1.2	1	8.78737	3.84×10^{-10}	6.91×10^{-9}
chord6	1.2	2	9.01591	4.05×10^{-10}	3.25×10^{-9}

A.6 Selected neutrino forecast rows

The following subset shows the density 2.8 g cm^{-3} neutrino forecasts at three CP phases; all 120 rows, including antineutrinos and vacuum, appear in the CSV. The conventional mixing parameters are supplied inputs, and the values are probabilities rather than detector event counts.

Table 14: Selected conventional matter-propagation forecasts at 1,300 km. These are not observed DUNE counts or microscopic mass predictions.

δ_{CP}	E (GeV)	$P_{\mu e}$	$P_{\mu\mu}$	$P_{\mu\tau}$
-90	0.5	0.133610	0.043078	0.823313
-90	1	0.037634	0.391283	0.571084
-90	2	0.078318	0.184368	0.737314
-90	2.5	0.078856	0.002635	0.918509
-90	3	0.067178	0.051942	0.880880
-90	5	0.030101	0.479998	0.489902
0	0.5	0.055076	0.058432	0.886492
0	1	0.034623	0.379217	0.586160
0	2	0.053109	0.192702	0.754189
0	2.5	0.062742	0.003419	0.933839
0	3	0.058117	0.049324	0.892559
0	5	0.029943	0.476559	0.493497
90	0.5	0.004914	0.043078	0.952008
90	1	0.006857	0.391283	0.601860
90	2	0.044369	0.184368	0.771262
90	2.5	0.048066	0.002635	0.949299
90	3	0.043730	0.051942	0.904328
90	5	0.022879	0.479998	0.497123

B. Statistical derivations and uncertainty propagation

B.1 Eliminating a linear amplitude

For a known shape vector $f(\omega)$ at a fixed nonlinear parameter ω, let the model be $m = Af$. The quadratic objective is

$$Q(A,\omega) = (y - Af)^T C^{-1}(y - Af) \tag{B.1}$$
$$= y^T C^{-1} y - 2Af^T C^{-1} y + A^2 f^T C^{-1} f. \tag{B.2}$$

Differentiation gives

$$\widehat{A}(\omega) = \frac{f^T C^{-1} y}{f^T C^{-1} f}, \qquad Q_{\text{prof}}(\omega) = y^T C^{-1} y - \frac{(f^T C^{-1} y)^2}{f^T C^{-1} f}. \tag{B.3}$$

The DESI optimizer searches only over Ω_m, with this exact amplitude elimination. Its reported joint parameter covariance uses the local Jacobian J,

$$V_\vartheta \simeq (J^T C^{-1} J)^{-1}. \tag{B.4}$$

This is a local linearized covariance, not a guarantee that every nonlinear posterior is exactly Gaussian. The archived bootstrap independently checks the distribution of the minimized discrepancy under the fitted reference and fixed covariance.

B.2 Whitening a correlated linear model

Let $C = LL^T$ be a Cholesky factorization and $y = Xb + \epsilon$. Define $\widetilde{y} = L^{-1}y$ and $\widetilde{X} = L^{-1}X$. Then

$$(y - Xb)^T C^{-1}(y - Xb) = \|\widetilde{y} - \widetilde{X}b\|^2. \tag{B.5}$$

This gives the familiar generalized least-squares estimator

$$\widehat{b} = (X^T C^{-1} X)^{-1} X^T C^{-1} y, \qquad V_b = (X^T C^{-1} X)^{-1}. \tag{B.6}$$

The FIRAS implementation additionally rescales columns by their whitened norms before solving. A temperature derivative and a dimensionless distortion template can differ greatly in numerical magnitude; column scaling improves conditioning without changing the fitted physical model. The scaled design condition number is recorded for every fit.

B.3 Projected identifiability of an added template

Let X_0 contain temperature and Galactic nuisance templates, and let s be a candidate additional spectral shape. Eliminating the nuisances produces the precision operator

$$P = C^{-1} - C^{-1} X_0 (X_0^T C^{-1} X_0)^{-1} X_0^T C^{-1}. \tag{B.7}$$

The information in the added amplitude is $s^T P s$. If s lies exactly in the nuisance span, this quantity vanishes and its amplitude is not identifiable from the spectrum. Two remaining templates s_1, s_2 can be compared through

$$r_{12} = \frac{s_1^T P s_2}{\sqrt{(s_1^T P s_1)(s_2^T P s_2)}}. \tag{B.8}$$

The recorded value for the y and μ shapes is +0.824046. This is a *projected template correlation*. It must not be described as the correlation of fitted y and μ parameters in the separate fits; a simultaneous two-parameter covariance generally has a different sign convention and would require its own calculation.

B.4 A held-out group with correlated noise

Partition a data vector into training indices c and held-out indices t. At fixed parameters, Gaussian conditioning gives

$$\mathbb{E}[y_t \mid y_c, \vartheta] = m_t + C_{tc} C_{cc}^{-1} (y_c - m_c), \tag{B.9}$$

$$S = C_{tt} - C_{tc} C_{cc}^{-1} C_{ct}. \tag{B.10}$$

Let J_c, J_t be model derivatives at the training fit. The derivative of the conditional mean is

$$G = J_t - C_{tc}C_{cc}^{-1}J_c. \tag{B.11}$$

To first order, parameter uncertainty adds

$$S_{\text{pred}} \simeq S + GV_{\vartheta,c}G^T. \tag{B.12}$$

This is the predictive covariance used in the seven DESI group holdouts. For this particular input, correlations are principally within redshift groups; keeping the general conditional formula avoids relying on that special pattern. Training fits overlap between holdouts, so their predictive scores must not be treated as seven independent experiments or summed into a new global probability.

B.5 A change of measurement units cannot improve a correct likelihood

Let a nonsingular linear transformation T change units or coordinates of a data vector. Set $y' = Ty$, $m' = Tm$ and $C' = TCT^T$. Since $C'^{-1} = T^{-T}C^{-1}T^{-1}$,

$$(y' - m')^T C'^{-1}(y' - m') = (y - m)^T C^{-1}(y - m). \tag{B.13}$$

A numerical "improvement" produced only by changing units and not transforming uncertainties consistently is therefore an implementation error. The same reasoning explains why coordinate-time rescaling alone cannot repair a physical cosmological mismatch when both predictions and measurements are transformed consistently.

B.6 Detector-efficiency cancellation and a rank deficiency

For a normalized collider ratio, a universal η cancels exactly in (10.3). Hence $\partial R_i/\partial\eta = 0$, and the Fisher information for that parameter from these ratios alone is zero. No number of additional normalized bins determines an overall efficiency to which the observable is insensitive.

For an absolute rate $\lambda = \epsilon_d g \eta r_P$ with both η and r_P unknown, the transformation $\eta \to a\eta$, $r_P \to r_P/a$ leaves λ unchanged. More measurements at the same uncalibrated normalization do not remove that degeneracy. An independent source model, a calibration, or a measurable state-dependent shape is required. This is an inference limitation of the specified observable, not a prohibition on testing the physical proposal.

C. Scope and result ledger

The following table is also provided as `results/domain_scorecard.csv`. It is the appropriate compact summary for a website because it identifies both the calculation and its evidentiary class.

Table 15: Twenty coverage areas, not twenty independent confirmations. Status is deliberately not compressed into a percentage.

ID / area	What was done	Scope
D01: Conservative finite dynamics	Executed: mathematical verification 36 new trajectories across four graphs; 376 new structural/stress assertions	Finite model only; does not establish physical particle identity
D02: Proper-time and clock covariance	Executed: mathematical and summary comparisons 12 new nonlinear clock comparisons and retained clock suite	Time transformation is not a derivation of the spacetime metric
D03: Coupled lattice dispersion	Executed: conditional model relation 50 random directions; axial-minus-diagonal coefficient -1/36	Physical field identification and dimensional lattice scale not fixed
D04: Quantum probabilities and no-signalling	Retained conditional tests rerun Kraus/positivity/quantum evolution checks conditional on Hilbert-space assumptions	Born rule not uniquely derived from finite classical Hamiltonian
D05: Gauge structure and chirality	Structural audit; not physically closed Pure-gradient scalar link plaquette equals identity	No unique SU(3)xSU(2)xU(1) representations derived
D06: Fundamental constants and masses	No new independent prediction Historical arithmetic checks retained with correct dependency labels	No new alpha or particle mass predicted independently
D07: Atomic clocks	Executed: summary comparison Optical residual +0.314 sigma	Conventional weak-field relation and supplied geometry
D08: Muon lifetime	Executed: summary comparison Moving lifetime residual -0.318 sigma	Rounded gamma; no full beam-data likelihood
D09: Muon magnetic anomaly	Executed: reference comparison 2025 experimental/SM reference residual +0.605 sigma	SM reference dependent; not an MNT loop calculation
D10: Universal free fall and redshift	Executed: summary comparisons MICROSCOPE -0.546 sigma; Galileo +0.077 sigma	Assumed conventional metric null, not unique evidence
D11: GW / electromagnetic propagation	Executed: published interval check Zero deviation lies in GW170817 interval	Emission assumptions; not Gaussian significance
D12: GWOSC raw strain	Not run on real detector files Local pipeline passes synthetic HDF5 and phase-recovery controls	Acquisition failed in workspace; no measured raw-strain result
D13: CMS Drell-Yan measurements	Executed: published-table comparison 24 of 26 conditional marginal pulls within 2 sigma	Published FEWZ reference; full cross-bin covariance absent

Continued

ID / area	What was done	Scope
D14: CMS 100k event reconstruction	Not run on real event file 1,200 synthetic rows used only for parser/rounding controls	No full sample downloaded; no particle mass prediction
D15: XENONnT dark-matter search	Executed: limit-table translations 20 rows checked; 10 heavy-mediator scale translations	No event-level recast or MNT recoil spectrum
D16: CMB monopole spectrum	Executed: 43-channel residual GLS Separate y, mu and grey fits compatible with zero in conditional approximation	Not full calibration likelihood; no MNT emission spectrum
D17: DESI BAO expansion	Executed: full supplied covariance Fixed comparator tension; fitted LCDM chi2=10.284 for nominal 11 df	Reference cosmology fitted; no microscopic parameter adjustment
D18: Cosmic age	Executed: supplied-background calculation 13.7906925 Gyr from declared parameters	Not a new independent age measurement or prediction
D19: Neutrinos and DUNE	Summary comparisons plus forecasts Two Daya Bay benchmark comparisons; 120 propagation forecasts	No observed DUNE far-detector spectrum fitted
D20: Compact objects, BBN and CMB anisotropy	Not newly evaluated in this campaign Literature requirements and required physical maps documented	No numerical validation score assigned

The complete machine-readable assertion ledger is `results/verification_ledger.csv`. Each row records its suite, origin, criterion, pass status, numerical value or discrepancy when applicable, tolerance and scope. Repeating a run does not add an independent assertion or physical observation. A thousand random initial conditions would test robustness of the specified dynamics, not create a thousand new measured particles.

D. Data provenance, checksums and acquisition boundaries

D.1 What entered the calculation workspace

The numerical inputs fall into two classes. The DESI and XENON tables, conventional summary inputs and baseline parameter files were available in the inherited reproducibility package. The DESI mean and covariance were checked against the public text. FIRAS and CMS numerical tables were transcribed from browser-readable primary products, with transcription provenance recorded. The FIRAS residual and noise information was also checked against its published table. CMS page-image retrieval was attempted but did not succeed; the parsed table was cross-checked against the publication's separate normalized-differential table.

These are published measurement products. They are not locally downloaded raw detector files. A reader can independently compare the supplied text and CSV files with the linked original sources. Local SHA-256 hashes identify the exact transcription and archived numerical snapshot; they are not represented as publisher-signed hashes of the original PDFs.

D.2 What did not enter the calculation workspace

No raw GWOSC HDF5 strain file, full CMS educational event CSV, XENON event-level likelihood dataset or DUNE far-detector oscillation sample was acquired here. The failed raw-file attempt is preserved in the execution logs. The synthetic raw-pipeline fixtures are labeled at creation and are never substituted for those missing observations.

For the optional GW150914 acquisition, the publisher checksum expectations are

Hanford H1, 32 seconds at 4,096 Hz
`H-H1_LOSC_4_V2-1126259446-32.hdf5`
MD5: `50441a42c13fc1f14e5c4ea5527f1515`.

Livingston L1, 32 seconds at 4,096 Hz
`L-L1_LOSC_4_V2-1126259446-32.hdf5`
MD5: `361ae6a040a9fef7897b1e0124d5b0a1`.

A new local acquisition records an additional SHA-256 hash and UTC retrieval time. The expected MD5 values serve as an integrity comparison with the publisher release, not a general cryptographic claim of security against deliberate collision attacks.

D.3 Model and code identity

The analysis plan fixes the microscopic coefficients, graph preparations, random seeds, statistical reference models, exclusions, fit ranges and failure policy. It was created for this execution using already-public data; it is not an independently timestamped preregistration before the measurements existed.

The final code/input manifest was recorded before the final staged numerical execution. Changes made while developing the pipeline are documented in `protocol/CHANGELOG.md`. In particular, enforcing a physical lag interval rather than a sample-count-dependent one is a software-definition correction, not a fitted change to a microscopic parameter. Publication tables are generated from output JSON and CSV files. Formatting changes to the report do not constitute new scientific tests.

The delivered snapshot has a separate manifest checked by `verify_snapshot.py`. Rerunning the scientific code changes timestamps and generated files, so preserve an untouched copy when verifying the original delivery. A successful manifest check establishes file identity; it does not certify a physical theory.

E. Executable methods: selected source listings

The package contains all scientific scripts, not only the selections printed here. These listings expose the central numerical steps without requiring an online service. The methods include explicit result labels so generated data cannot silently become a measured observation. The scripts below are the tested source; long lines are wrapped typographically for the page.

E.1 DESI covariance fit and refitted bootstrap

```
"""DESI DR2 compressed BAO: fixed comparison, transparent LCDM fit and held-out groups."""
from __future__ import annotations
import json
import numpy as np
import pandas as pd
from scipy.linalg import cho_factor, cho_solve
from scipy.optimize import minimize_scalar
from scipy.stats import chi2
from numpy.polynomial.legendre import leggauss
from common import *

def run():
 raw=np.loadtxt(DATA/'desi_dr2_mean.txt',dtype=str);z=raw[:,0].astype(float);y=raw[:,1].astype(float);kind=raw[:,2]
 C=np.loadtxt(DATA/'desi_dr2_cov.txt');ch=Checks('desi');cf=cho_factor(C);Ci=cho_solve(cf,np.eye(len(y)))
 ch.add('13 measurements with matching covariance',C.shape==(13,13) and len(y)==13)
 ch.small('Covariance symmetry',np.max(abs(C-C.T)),1e-14)
 ch.add('Covariance is positive definite',np.linalg.eigvalsh(C)[0]>0)
 Or=.000092;c=299792.458;A0=c/(67.4*147.1)
 def shape(om,nquad=80):
  g,w=leggauss(nquad) if nquad!=80 else (g80,w80)
  zz=z[:,None]*(g+1)/2
  E=np.sqrt(om*(1+z)**3+Or*(1+z)**4+1-om-Or)
  Ez=np.sqrt(om*(1+zz)**3+Or*(1+zz)**4+1-om-Or)
  I=z/2*np.sum(w/Ez,axis=1)
  return np.where(kind=='DM_over_rs',I,np.where(kind=='DH_over_rs',1/E,(I*I*z/E)**(1/3)))
 g80,w80=leggauss(80)
 def fit(obs,idx=None):
  idx=np.arange(13) if idx is None else np.asarray(idx);yy=obs[idx];cc=C[np.ix_(idx,idx)];inv=np.linalg.inv(cc)
  def score(om,ret=False):
   f=shape(om)[idx];a=float(f@inv@yy/(f@inv@f));r=yy-a*f;s=float(r@inv@r)
   return (s,a) if ret else s
  op=minimize_scalar(score,bounds=(.05,.6),method='bounded',options={'xatol':2e-10})
  ss,aa=score(op.x,True);return float(op.x),aa,ss
 fixed=A0*shape(.315);r0=y-fixed;fixedchi=float(r0@Ci@r0)
 om,A,bestchi=fit(y);best=A*shape(om)
 def jac(om,A):
  eps=1e-5
  return np.c_[shape(om), A*(shape(om+eps)-shape(om-eps))/(2*eps)]
 J=jac(om,A);pcov=np.linalg.inv(J.T@Ci@J)
 ch.small('80 vs 160 point distance quadrature',np.max(abs(shape(om)-shape(om,160))),1e-12)
```

```
ch.small('Fixed benchmark reproduces retained chi-square',abs(fixedchi-28.5062),1e-4)
ch.add('Fitted comparator does not worsen fixed comparator',bestchi<=fixedchi)
ch.small('Amplitude normal equation',abs(shape(om)@Ci@(y-best)),1e-8)
# A model-dependent age, explicitly using supplied H0 and not claiming a new measurement.
from scipy.integrate import quad
def age(om,H0):
 I=quad(lambda a:1/np.sqrt(Or/a**2+om/a+(1-om-Or)*a*a),0,1,epsabs=1e-12)[0]
 return I/(H0*1000/3.0856775814913673e22)/31557600/1e9
rng=np.random.default_rng(20260907);L=np.linalg.cholesky(C);bt=[]
for n in range(1000):
 obs=best+L@rng.normal(size=13);oi,ai,ss=fit(obs);bt.append({'replication':n,'Omega_m':oi,'amplitude':ai,'chi2_min':ss})
b=np.array([x['chi2_min'] for x in bt]);bootp=(1+np.sum(b>=bestchi))/(len(b)+1)
cv=[];cvrows=[]
for zg in np.unique(z):
 test=np.flatnonzero(z==zg);train=np.flatnonzero(z!=zg);ot,at,ss=fit(y,train);pred=at*shape(ot)
 Ccc=C[np.ix_(train,train)];Ctc=C[np.ix_(test,train)];Ctt=C[np.ix_(test,test)];icc=np.linalg.inv(Ccc)
 conditional=pred[test]+Ctc@icc@(y[train]-pred[train]);S=Ctt-Ctc@icc@Ctc.T
 JJ=jac(ot,at);V=np.linalg.inv(JJ[train].T@icc@JJ[train]);G=JJ[test]-Ctc@icc@JJ[train];Sp=S+G@V@G.T
 rr=y[test]-conditional;sspred=float(rr@np.linalg.solve(Sp,rr))
 cv.append({'held_out_z':zg,'training_rows':len(train),'test_rows':len(test),'Omega_m_train':ot,'amplitude_train':at,'training_chi2':ss,'predictive_mahalanobis':sspred,'nominal_p':chi2.sf(sspred,len(test))})
 for j,ii in enumerate(test):cvrows.append({'z':z[ii],'quantity':kind[ii],'observed':y[ii],'held_out_prediction':conditional[j],'predictive_sigma':np.sqrt(Sp[j,j]),'predictive_pull':rr[j]/np.sqrt(Sp[j,j])})
 ch.add(f'Positive held-out covariance z={zg}',np.linalg.eigvalsh(Sp)[0]>0)
# Wrong expansion shape as a power check, amplitude profiled but matter-only E(z).
e=(1+z)**1.5;integ=2*(1-1/np.sqrt(1+z));wrong=np.where(kind=='DM_over_rs',integ,np.where(kind=='DH_over_rs',1/e,(integ**2*z/e)**(1/3)))
aw=wrong@Ci@y/(wrong@Ci@wrong);wrongchi=float((y-aw*wrong)@Ci@(y-aw*wrong))
ch.add('Matter-only negative control is separated from best comparator',wrongchi>bestchi+25)
rows=[{'z':float(z[i]),'quantity':kind[i],'observed':float(y[i]),'marginal_sigma':float(np.sqrt(C[i,i])),'fixed_prediction':float(fixed[i]),'fixed_pull':float(r0[i]/np.sqrt(C[i,i])),'fitted_LCDM_prediction':float(best[i]),'fitted_pull':float((y[i]-best[i])/np.sqrt(C[i,i]))} for i in range(13)]
save_rows('desi_points.csv',rows);save_rows('desi_bootstrap.csv',bt);save_rows('desi_leave_group_out.csv',cv);save_rows('desi_heldout_points.csv',cvrows)
res={'role':'Published compressed full-covariance comparison. LCDM optimization is not fitting or validating microscopic MNT.','N':13,'fixed':{'H0':67.4,'rd':147.1,'Omega_m':.315,'Omega_r':Or,'A':A0,'chi2':fixedchi,'nominal_df':13,'tail_probability':chi2.sf(fixedchi,13),'age_Gyr_from_supplied_background':age(.315,67.4)},'fitted_comparator':{'Omega_m':om,'Omega_m_sigma_linearized':np.sqrt(pcov[1,1]),'A_c_over_H0rd':A,'A_sigma_linearized':np.sqrt(pcov[0,0]),'H0_times_rd_km_s':c/A,'chi2':bestchi,'nominal_df':11,'tail_probability':chi2.sf(bestchi,11),'bootstrap_tail_probability':bootp,'bootstrap_MC_sigma':np.sqrt(bootp*(1-bootp)/1001),'parameter_order':['A','Omega_m'],'parameter_covariance':pcov,'bootstrap_replications':1000},'matter_only_negative_control':{'chi2':
```

```
    wrongchi,'amplitude':aw},'heldout_caveat':'Seven overlapping training sets;
    group predictive scores are not independent or a historical blind test.
    Parameter uncertainty uses a linearized covariance.','
    points_within_one_sigma_fixed':int(sum(abs(r0/np.sqrt(np.diag(C)))<=1)),'
    points_within_two_sigma_fixed':int(sum(abs(r0/np.sqrt(np.diag(C)))<=2))}
 save_json('desi_results.json',res)
 plt.figure(figsize=(7.2,4.2));x=np.arange(13);plt.axhline(0,linewidth=.8);plt.
    plot(x,r0/np.sqrt(np.diag(C)),'o',label='Frozen supplied background');plt.plot
    (x,(y-best)/np.sqrt(np.diag(C)),'s',label='Fitted LCDM comparator');plt.
    axhline(2,linestyle=':',linewidth=.8);plt.axhline(-2,linestyle=':',linewidth
    =.8);plt.xticks(x,[f'{zz:g}\n{kk[:2]}' for zz,kk in zip(z,kind)],fontsize=7);
    plt.ylabel('Marginal residual / sigma');plt.xlabel('Redshift and distance type
     (correlations retained in likelihood)');plt.legend(fontsize=8);figure('
    desi_pulls')
 plt.figure(figsize=(7,4));plt.hist(b,bins=35,density=True,histtype='step',label='
    1,000 refitted simulations');xx=np.linspace(0,35,300);plt.plot(xx,chi2.pdf(xx
    ,11),label='Chi-square, 11 degrees');plt.axvline(bestchi,linestyle='--',label=
    'Measured minimum');plt.xlabel('Minimum chi-square');plt.ylabel('Probability
    density');plt.legend(fontsize=8);figure('desi_bootstrap')
 plt.figure(figsize=(7,4));plt.errorbar(np.arange(13),[v['observed']-v['
    held_out_prediction'] for v in cvrows],yerr=[v['predictive_sigma'] for v in
    cvrows],fmt='o');plt.axhline(0,linewidth=.8);plt.xlabel('Held-out measurement
    index');plt.ylabel('Observed minus prediction (distance / sound horizon)');
    figure('desi_holdout')
 ch.finish();print(json.dumps(res,default=lambda x:x.tolist() if hasattr(x,'tolist
    ') else float(x),indent=2))
if __name__=='__main__':run()
```

E.2 FIRAS residual generalized least squares

```
"""Conditional GLS template reanalysis of FIRAS published residuals and covariance
    ."""
from __future__ import annotations
import json
import numpy as np
from scipy.linalg import toeplitz
from scipy.stats import chi2
from common import *
H=6.62607015e-34;C_LIGHT=299792458.;KB=1.380649e-23

def planck(nu,T):
 x=H*nu/(KB*T);return (2*H*nu**3/C_LIGHT**2)/np.expm1(x)/1e-23

def run():
 d=np.loadtxt(DATA/'firas_monopole_transcribed.txt');nu=d[:,0]*100*C_LIGHT;res=d
    [:,2];sig=d[:,3];gal=d[:,4];T=2.725
 q=np.array(json.loads((DATA/'firas_correlation_lags.json').read_text()));R=
    toeplitz(q);cov=np.outer(sig,sig)*R;L=np.linalg.cholesky(cov);inv=np.linalg.
    inv(cov);ch=Checks('firas')
 ch.add('All 43 residual channels present',len(d)==43 and len(q)==43)
 ch.add('Published rounded Toeplitz covariance positive definite',np.linalg.
    eigvalsh(R)[0]>0)
 x=H*nu/(KB*T);B=planck(nu,T);D=B*x/T/(-np.expm1(-x));Y=T*D*(x/np.tanh(x/2)-4);MU
    =-T*D/x
```

```
ch.small('Analytic temperature derivative agrees with central difference',np.max(abs((planck(nu,T+1e-5)-planck(nu,T-1e-5))/(2e-5)-D))/np.max(abs(D)),1e-8)
def fit(X,y,covariance=cov):
 LL=np.linalg.cholesky(covariance);W=np.linalg.solve(LL,X);yy=np.linalg.solve(LL,y);sc=np.linalg.norm(W,axis=0);Ws=W/sc
 bb=np.linalg.lstsq(Ws,yy,rcond=None)[0];b=bb/sc;V=np.linalg.inv(Ws.T@Ws)/np.outer(sc,sc);rr=y-X@b;ss=float(rr@np.linalg.solve(covariance,rr))
 return b,V,ss,rr,float(np.linalg.cond(Ws))
models={'baseline':(np.c_[D,gal],['delta_T_K','galaxy_amplitude']), 'y':(np.c_[D,gal,Y],['delta_T_K','galaxy_amplitude','y']), 'mu':(np.c_[D,gal,MU],['delta_T_K','galaxy_amplitude','mu']), 'grey':(np.c_[D,gal,B],['delta_T_K','galaxy_amplitude','fractional_amplitude'])}
rows=[];fits={};nuis={}
for name,(X,names) in models.items():
 b,V,ss,rr,cond=fit(X,res);bd,Vd,ssd,_,_=fit(X,res,np.diag(sig**2));base=fit(models['baseline'][0],res)[2]
 fits[name]={'parameters':dict(zip(names,b)),'standard_errors':dict(zip(names,np.sqrt(np.diag(V)))),'covariance':V,'chi2':ss,'nominal_df':43-len(b),'nominal_tail_diagnostic':chi2.sf(ss,43-len(b)),'scaled_design_condition_number':cond,'delta_chi2_from_baseline':base-ss,'diagonal_covariance_sensitivity':{'parameters':dict(zip(names,bd)),'standard_errors':dict(zip(names,np.sqrt(np.diag(Vd)))),'chi2':ssd}}
 for j,n in enumerate(names):rows.append({'fit':name,'parameter':n,'estimate':b[j],'conditional_sigma':np.sqrt(V[j,j]),'lower_95_Gaussian':b[j]-1.95996398454*np.sqrt(V[j,j]),'upper_95_Gaussian':b[j]+1.95996398454*np.sqrt(V[j,j]),'diagonal_cov_estimate':bd[j]})
 ch.small(f'GLS normal equations {name}',np.linalg.norm(X.T@inv@rr)/(1+np.linalg.norm(X.T@inv@res)),1e-10)
 ch.add(f'Residual chi-square finite {name}',ss>=0 and np.isfinite(ss))
 if name!='baseline':ch.add(f'Added template cannot increase minimum {name}',ss<=base+1e-10)
 nuis[name]=(b,V,rr)
# No blackbody-temperature discovery: column 2 was constructed with a temperature already supplied.
constructed_residual=d[:,1]*1000-planck(nu,T)
# Rounded wavenumbers and brightness mean exact reconstruction need not equal stored residual.
save_rows('firas_channels.csv',[{'wavenumber_cm-1':d[i,0],'frequency_GHz':nu[i]/1e9,'published_residual_kJy_sr':res[i],'sigma_kJy_sr':sig[i],'galaxy_template_kJy_sr':gal[i],'temperature_template_kJy_sr_K':D[i],'y_template_kJy_sr':Y[i],'mu_template_kJy_sr':MU[i],'baseline_fitted_residual_kJy_sr':nuis['baseline'][2][i],'y_fitted_residual_kJy_sr':nuis['y'][2][i]} for i in range(43)])
# Repeated simulations are calibration of this Gaussian estimator, not repeated experiments.
rng=np.random.default_rng(202609071);X,names=models['y'];truth=np.array([2e-5,.015,1e-5]);N=4000
E=L@rng.normal(size=(43,N));obs=X@truth[:,None]+E
V=np.linalg.inv(X.T@inv@X);est=V@X.T@inv@obs;se=np.sqrt(np.diag(V));zs=(est-truth[:,None])/se[:,None];coverage=np.mean(abs(zs)<=1.95996398454,axis=1)
sim=[{'parameter':names[j],'injected':truth[j],'mean_recovered':np.mean(est[j]),'empirical_sd':np.std(est[j],ddof=1),'analytic_sd':se[j],'mean_z':np.mean(zs[j]),'sd_z':np.std(zs[j],ddof=1),'coverage_95':coverage[j]} for j in range(3)]
for j in range(3):
 ch.small(f'Injection ensemble standardized mean {names[j]}',abs(np.mean(zs[j]))
```

```
    ,.08)
  ch.small(f'Injection ensemble 95 percent coverage {names[j]}',abs(coverage[j
    ]-.95),.025)
 # Projected y/mu degeneracy demonstrates why fits are separate.
 X0=models['baseline'][0];P=inv-inv@X0@np.linalg.solve(X0.T@inv@X0,X0.T@inv)
 corr=float(Y@P@MU/np.sqrt((Y@P@Y)*(MU@P@MU)))
 result={'role':'Published-residual Gaussian-template analysis; not raw FIRAS
    calibration and not an MNT spectrum prediction.','channels':43,'
    correlation_min_eigenvalue':float(np.linalg.eigvalsh(R)[0]),'
    reference_temperature_K':T,'fits':fits,'projected_y_mu_template_correlation':
    corr,'injection_simulations':N,'injection_coverage':sim,'systematic_boundary':
    'Conditional errors omit full historical calibration and sky-component
    systematics. Published residuals were already processed/fitted. Nominal tail
    probabilities are diagnostic, not replacement collaboration confidence limits.
    ','absolute_EQEF_efficiency':'Not identifiable from a grey amplitude without
    an independent emission model and absolute calibration.','
    grey_energy_loss_sign':'The fitted grey coefficient is delta I/I; a loss-only
    interpretation would use epsilon=-coefficient and require a separately
    justified physical model.'}
 save_json('firas_results.json',result);save_rows('firas_fit_parameters.csv',rows)
    ;save_rows('firas_injection_coverage.csv',sim)
 np.savetxt(OUT/'firas_covariance_kJy2.txt',cov,fmt='%.12g')
 plt.figure(figsize=(7.2,4.3));plt.errorbar(nu/1e9,res,yerr=sig,fmt='.',label='
    Published FIRAS residual');plt.plot(nu/1e9,models['y'][0]@nuis['y'][0],label='
    Temperature + Galaxy + y fit');plt.axhline(0,linewidth=.8);plt.xlabel('
    Frequency (GHz)');plt.ylabel('Residual intensity (kJy / sr)');plt.legend(
    fontsize=8);figure('firas_residuals')
 plt.figure(figsize=(7.2,4.1));plt.plot(nu/1e9,Y*1e-5,label='Compton y = 10^-5');
    plt.plot(nu/1e9,MU*1e-5,label='Chemical potential mu = 10^-5');plt.plot(nu/1e9
    ,D*1e-5,label='Temperature change = 10 microK');plt.xlabel('Frequency (GHz)');
    plt.ylabel('Template change (kJy / sr)');plt.legend(fontsize=8);figure('
    firas_templates')
 plt.figure(figsize=(6.8,4));plt.hist(zs[2],bins=40,density=True,histtype='step',
    label='4,000 synthetic y fits');t=np.linspace(-4,4,300);plt.plot(t,np.exp(-t*t
    /2)/np.sqrt(2*np.pi),label='Unit normal');plt.xlabel('(Recovered y - injected
    y) / conditional sigma');plt.ylabel('Density');plt.legend(fontsize=8);figure('
    firas_coverage')
 ch.finish();print(json.dumps(result,default=lambda v:v.tolist() if hasattr(v,'
    tolist') else float(v),indent=2))
if __name__=='__main__':run()
```

E.3 The complete execution runner

```
#!/usr/bin/env python3
"""Reproduce the published-table analyses and finite-model verification offline.

By default both the six new suites and the six retained suites are executed.
Raw detector acquisition is a separate explicit command in code/raw_detectors.py.
No failed acquisition is replaced by synthetic observations.
"""
from __future__ import annotations
import argparse, csv, hashlib, importlib.metadata, json, os, platform, subprocess,
     sys, time
from datetime import datetime, timezone
```

```
from pathlib import Path
ROOT = Path(__file__).resolve().parent
RESULTS = ROOT / 'results'
NEW = [('desi','analyze_desi.py'),('firas','analyze_firas.py'),('cms','analyze_cms.py'),
       ('stress','stress_model.py'),('constraints','analyze_constraints.py'),('raw_fixture','test_raw_pipelines.py')]
PRIOR = ROOT / 'prior/MNT_EQEF_Final_Source_Package'

def main() -> int:
    parser = argparse.ArgumentParser(description=__doc__)
    parser.add_argument('--skip-prior', action='store_true', help='Run 738 new assertions only; do not claim prior checks rerun.')
    parser.add_argument('--collect-only', action='store_true', help='Aggregate existing outputs without claiming a new execution.')
    parser.add_argument('--prior-only', action='store_true', help='Execute retained suites; aggregate previously executed new results.')
    args = parser.parse_args()
    if args.prior_only and args.skip_prior: parser.error('Choose only one stage option')
    RESULTS.mkdir(exist_ok=True)
    start = datetime.now(timezone.utc).isoformat(); timer = time.monotonic()
    env = dict(os.environ, OPENBLAS_NUM_THREADS='1', OMP_NUM_THREADS='1', MKL_NUM_THREADS='1')
    if not args.collect_only:
        for name, filename in ([] if args.prior_only else NEW):
            print('Running '+name, flush=True)
            with (RESULTS/(name+'_execution.log')).open('w') as log:
                subprocess.run([sys.executable, str(ROOT/'code'/filename)], cwd=ROOT,
                               stdout=log, stderr=subprocess.STDOUT, check=True, env=env)
            summary = json.loads((RESULTS/(name+'_check_summary.json')).read_text())
            if summary['failed']:
                raise RuntimeError(name+' failed assertions; inspect results, no pass is declared')
        if not args.prior_only:
            (RESULTS/'execution_new.json').write_text(json.dumps({'started_utc':start,'completed_utc':datetime.now(timezone.utc).isoformat(),'suites':[x[0] for x in NEW],'all_executed':True},indent=2)+'\n')
        if not args.skip_prior:
            prior_start=datetime.now(timezone.utc).isoformat()
            print('Running six retained suites', flush=True)
            with (RESULTS/'prior_reexecution.log').open('w') as log:
                subprocess.run([sys.executable,str(PRIOR/'run_all.py')],cwd=PRIOR,
                               stdout=log,stderr=subprocess.STDOUT,check=True,env=env)
            (RESULTS/'execution_prior.json').write_text(json.dumps({'started_utc':prior_start,'completed_utc':datetime.now(timezone.utc).isoformat(),'all_executed':True},indent=2)+'\n')
    rows=[]; suites=[]
    for name,_ in NEW:
        s=json.loads((RESULTS/(name+'_check_summary.json')).read_text())
        suites.append({'suite':name,'assertions':s['total'],'passed':s['passed'],'
```

```
origin':'new implementation'})
        with (RESULTS/(name+'_checks.csv')).open() as f:
            for r in csv.DictReader(f):
                rows.append({'origin':'new','suite':r['suite'],'check':r['name'],'passed':r['passed'],
                             'error_or_value':r['error'],'tolerance':r['tolerance'],'detail':r['kind'],
                             'scope':'Mathematical / software assertion; not an independent experiment'})
    if not args.skip_prior:
        with (PRIOR/'results/All_946_Verification_Checks.csv').open() as f:
            for r in csv.DictReader(f): rows.append({'origin':'retained',**r})
        for name in dict.fromkeys(r['suite'] for r in rows if r['origin']=='retained'):
            rr=[r for r in rows if r['origin']=='retained' and r['suite']==name]
            suites.append({'suite':name,'assertions':len(rr),'passed':sum(str(r['passed']).lower()=='true' for r in rr),'origin':'retained suite'})
    for filename,data in [('verification_ledger.csv',rows),('suite_counts.csv',suites)]:
        with (RESULTS/filename).open('w',newline='') as f:
            w=csv.DictWriter(f,fieldnames=data[0]);w.writeheader();w.writerows(data)
    failed=[r for r in rows if str(r['passed']).lower()!='true']
    versions={}
    for name in ['numpy','scipy','pandas','matplotlib','h5py','sympy','mpmath']:
        versions[name]=importlib.metadata.version(name)
    result={'edition':'OV-1.0','action':'collect existing results only' if args.collect_only else 'actual offline execution',
            'stage_records':{k:json.loads((RESULTS/f'execution_{k}.json').read_text()) for k in ['new','prior'] if (RESULTS/f'execution_{k}.json').exists()},
            'started_utc':start,'finished_utc':datetime.now(timezone.utc).isoformat(),
            'elapsed_seconds':time.monotonic()-timer,'python':sys.version,'platform':platform.platform(),
            'libraries':versions,'BLAS_threads_per_subprocess':1,'assertions':len(rows),'passed':len(rows)-len(failed),'failed':failed,
            'new_assertions':sum(r['origin']=='new' for r in rows),
            'retained_assertions':sum(r['origin']=='retained' for r in rows),
            'synthetic_replications':{'DESI_refitted_bootstraps':1000,'FIRAS_coverage_trials':4000},
            'published_measurement_values':{'DESI_BAO':13,'FIRAS_residual_channels':43,'CMS_muon_and_electron_bins':26},
            'other_products':{'XENON_limit_rows':20,'summary_comparisons':7,'GW170817_interval':1,'neutrino_forecast_points':120},
            'new_independent_parameter_free_MNT_physical_alignments':0,
            'real_raw_GWOSC_pipeline_executed':False,'real_raw_CMS100k_pipeline_executed':False,
            'raw_pipeline_statement':'Synthetic-fixture tested only in this edition; optional local runs save separate results.',
            'statistical_statement':'Do not aggregate these heterogeneous items into an experimental accuracy percentage.'}
    (RESULTS/'campaign_summary.json').write_text(json.dumps(result,indent=2)+'\n')
    print(json.dumps(result,indent=2))
    return 1 if failed else 0
```

```
if __name__=='__main__':raise SystemExit(main())
```

E.4 Where the other implementations reside

File	Purpose
code/stress_model.py	Independent nonlinear Hamiltonian implementation, finite-graph stresses, reversibility, clock changes, spectral and scalar-loop checks.
code/analyze_cms.py	Published Drell–Yan table comparison, directional marginal errors, common-efficiency cancellation and cross-channel diagnostics.
code/analyze_constraints.py	Summary measurements, heavy-mediator scale conversion and conventional neutrino forecasts.
code/raw_detectors.py	Optional real-file acquisition, input validation, GW conditioning and fixed-template profile, CMS interval mass reconstruction.
code/test_raw_pipelines.py	Labeled synthetic HDF5 and CSV fixtures, injected timing/phase recovery, corrupted-input controls.
code/common.py	Shared result-writing, figure export and assertion-record functions.
prior/.../run_all.py	Reexecution of the six retained suites and normalization of their 946-assertion ledger.

The reference scripts intentionally avoid an automatic "confirmed theory" decision. An analysis which lacks a specified MNT observable records that boundary; an optional raw-file acquisition which fails exits with an error. The numerical and physical interpretation are therefore both inspectable.

References

[1] R. Zwanzig, "Nonlinear generalized Langevin equations," *Journal of Statistical Physics* **9**, 215–220 (1973). doi:10.1007/BF01008729.

[2] H. Mori, "Transport, Collective Motion, and Brownian Motion," *Progress of Theoretical Physics* **33**, 423–455 (1965). doi:10.1143/PTP.33.423.

[3] A. O. Caldeira and A. J. Leggett, "Path integral approach to quantum Brownian motion," *Physica A* **121**, 587–616 (1983). doi:10.1016/0378-4371(83)90013-4.

[4] E. Noether, "Invariante Variationsprobleme," *Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen, Mathematisch-Physikalische Klasse*, 235–257 (1918). English translation by M. A. Tavel, "Invariant Variation Problems," arXiv:physics/0503066.

[5] R. Kubo, "The fluctuation-dissipation theorem," *Reports on Progress in Physics* **29**, 255–284 (1966). doi:10.1088/0034-4885/29/1/306.

[6] A. Einstein, "Zur Elektrodynamik bewegter Körper," *Annalen der Physik* **17**, 891–921 (1905). doi:10.1002/andp.19053221004. Relativistic clock comparisons are used as an established reference, not as a claimed derivation from EQEF.

[7] A. Einstein, "Die Grundlage der allgemeinen Relativitätstheorie," *Annalen der Physik* **354**, 769–822 (1916). doi:10.1002/andp.19163540702.

[8] G. Cowan, K. Cranmer, E. Gross and O. Vitells, "Asymptotic formulae for likelihood-based tests of new physics," *European Physical Journal C* **71**, 1554 (2011). arXiv:1007.1727; doi:10.1140/epjc/s10052-011-1554-0.

[9] M. Abdul-Karim et al. (DESI Collaboration), "DESI DR2 Results II: Measurements of Baryon Acoustic Oscillations and Cosmological Constraints," *Physical Review D* **112**, 083515 (2025). arXiv:2503.14738; doi:10.1103/tr6y-kpc6.

[10] DESI DR2 BAO likelihood data distributed in *CobayaSampler/bao_data*, directory `desi_bao_dr2`. Repository directory; GC-combined mean vector; GC-combined covariance. Accessed 6 September 2026; local input bytes recorded separately.

[11] N. Aghanim et al. (Planck Collaboration), "Planck 2018 results. VI. Cosmological parameters," *Astronomy & Astrophysics* **641**, A6 (2020). arXiv:1807.06209; doi:10.1051/0004-6361/201833910.

[12] NASA Goddard Space Flight Center, LAMBDA, *COBE/FIRAS CMB monopole spectrum*, version 1 (May 2005). Numerical product based on Fixsen et al. (1996) and Fixsen and Mather (2002). Official numerical table; Product documentation. Browser-accessed 6–7 September 2026; locally transcribed columns and exact bytes are documented separately.

[13] D. J. Fixsen, E. S. Cheng, J. M. Gales, J. C. Mather, R. A. Shafer and E. L. Wright, "The Cosmic Microwave Background Spectrum from the Full COBE/FIRAS Data Set," *Astrophysical Journal* **473**, 576 (1996). arXiv:astro-ph/9605054; doi:10.1086/178173.

[14] M. Planck, "Ueber das Gesetz der Energieverteilung im Normalspectrum," *Annalen der Physik* **309**, 553–563 (1901). doi:10.1002/andp.19013090310.

[15] S. Chatrchyan et al. (CMS Collaboration), "Measurement of the Drell–Yan Cross Section in pp Collisions at $\sqrt{s}$ = 7 TeV," *Journal of High Energy Physics* **10**, 007 (2011). arXiv:1108.0566; doi:10.1007/JHEP10(2011)007. Tables 8–10 supply the measurement and reference ratios used here.

[16] E. Aprile et al. (XENON Collaboration), "First Search for Light Dark Matter in the Neutrino Fog with XENONnT," *Physical Review Letters* **134**, 111802 (2025). arXiv:2409.17868; public release doi:10.5281/zenodo.14522707.

[17] XENONnT Collaboration, *light_wimp_data_release*. Official public repository; SI limit CSV; Light-mediator limit CSV. Accessed 6 September 2026.

[18] C. W. Chou, D. B. Hume, T. Rosenband and D. J. Wineland, "Optical Clocks and Relativity," *Science* **329**, 1630–1633 (2010). doi:10.1126/science.1192720; NIST-hosted primary paper. Figure 3 supplies the height-step comparison used here.

[19] P. Delva et al., "Gravitational Redshift Test Using Eccentric Galileo Satellites," *Physical Review Letters* **121**, 231101 (2018). arXiv:1812.03711; doi:10.1103/PhysRevLett.121.231101.

[20] D. M. Webber et al. (MuLan Collaboration), "Measurement of the Positive Muon Lifetime and Determination of the Fermi Constant to Part-per-Million Precision," *Physical Review Letters* **106**, 041803 (2011). arXiv:1010.0991; doi:10.1103/PhysRevLett.106.041803. Independent rest-lifetime input.

[21] J. Bailey et al., "Measurements of relativistic time dilatation for positive and negative muons in a circular orbit," *Nature* **268**, 301–305 (1977). doi:10.1038/268301a0. Published positive-muon lifetime and rounded abstract Lorentz factor; no full beam-distribution reanalysis.

[22] P. Touboul et al., "MICROSCOPE Mission: Final Results of the Test of the Equivalence Principle," *Physical Review Letters* **129**, 121102 (2022). arXiv:2209.15487.

[23] D. P. Aguillard et al. (Muon $g-2$ Collaboration), "Measurement of the Positive Muon Anomalous Magnetic Moment to 127 ppb" (2025). arXiv:2506.03069. Experimental world-average comparison as reported in the cited release.

[24] R. Aliberti et al., "The anomalous magnetic moment of the muon in the Standard Model: an update" (2025). arXiv:2505.21476. Reference prediction and hadronic-input qualifications.

[25] B. P. Abbott et al., "Gravitational Waves and Gamma-rays from a Binary Neutron Star Merger: GW170817 and GRB 170817A," *Astrophysical Journal Letters* **848**, L13 (2017). arXiv:1710.05834.

[26] F. P. An et al. (Daya Bay Collaboration), "Precision Measurement of Reactor Antineutrino Oscillation at Kilometer-Scale Baselines by Daya Bay," *Physical Review Letters* **130**, 161802 (2023). arXiv:2211.14988; doi:10.1103/PhysRevLett.130.161802.

[27] L. Wolfenstein, “Neutrino oscillations in matter,” *Physical Review D* **17**, 2369–2374 (1978). doi:10.1103/PhysRevD.17.2369.

[28] S. Navas et al. (Particle Data Group), “Review of Particle Physics,” *Physical Review D* **110**, 030001 (2024), with the 2025 update. doi:10.1103/PhysRevD.110.030001; official review index. Particle conventions, kinematics, cosmology and statistical methods are used as external effective-physics references.

[29] B. Abi et al. (DUNE Collaboration), “Deep Underground Neutrino Experiment (DUNE), Far Detector Technical Design Report, Volume II: DUNE Physics” (2020). arXiv:2002.03005.

[30] Fermi National Accelerator Laboratory, “DUNE will use liquid-argon time projection chamber technology both near and far” (May 2026). Official experiment-status report. Accessed 6 September 2026.

[31] Gravitational Wave Open Science Center, *GW150914*, public event and strain-data release. Official event page; Versioned catalog record. Browser-accessed 6–7 September 2026. Large raw files were not acquired in this workspace.

[32] Gravitational Wave Open Science Center, *GW150914 strain and illustrative numerical-relativity waveform*. Hanford 32-second HDF5; Livingston 32-second HDF5; Illustrative waveform. URLs and publisher checksums are pinned in the optional local acquisition code.

[33] B. P. Abbott et al. (LIGO Scientific Collaboration and Virgo Collaboration), “A guide to LIGO–Virgo detector noise and extraction of transient gravitational-wave signals,” *Classical and Quantum Gravity* **37**, 055002 (2020). arXiv:1908.11170; doi:10.1088/1361-6382/ab685e.

[34] T. McCauley, *Events with two muons from 2010*, CERN Open Data Portal (2014). Record 303; doi:10.7483/OPENDATA.CMS.4M97.3SQ9. Educational reduction of 100,000 dimuon events; its documentation states that the sample is not suitable for a full physics analysis. No large event CSV was acquired here.

[35] P. J. Mohr, D. B. Newell, B. N. Taylor and E. Tiesinga, “CODATA recommended values of the fundamental physical constants: 2022,” *Reviews of Modern Physics* **97**, 025002 (2025). doi:10.1103/RevModPhys.97.025002; arXiv:2409.03787.

[36] Bureau International des Poids et Mesures, *The International System of Units (SI)*, ninth edition (2019), official online brochure and updates. SI brochure. Used for defining-constant and unit conventions, not as a microscopic prediction.

[37] C. N. Yang and R. L. Mills, “Conservation of Isotopic Spin and Isotopic Gauge Invariance,” *Physical Review* **96**, 191–195 (1954). doi:10.1103/PhysRev.96.191.

[38] K. G. Wilson, “Confinement of quarks,” *Physical Review D* **10**, 2445 (1974). doi:10.1103/PhysRevD.10.2445.

[39] J. Kogut and L. Susskind, "Hamiltonian formulation of Wilson's lattice gauge theories," *Physical Review D* **11**, 395 (1975). doi:10.1103/PhysRevD.11.395.

[40] S. Weinberg, "A Model of Leptons," *Physical Review Letters* **19**, 1264–1266 (1967). doi:10.1103/PhysRevLett.19.1264.

[41] H. B. Nielsen and M. Ninomiya, "A no-go theorem for regularizing chiral fermions," *Physics Letters B* **105**, 219–223 (1981). doi:10.1016/0370-2693(81)91026-1.

[42] G. Lindblad, "On the generators of quantum dynamical semigroups," *Communications in Mathematical Physics* **48**, 119–130 (1976). doi:10.1007/BF01608499.

[43] V. Gorini, A. Kossakowski and E. C. G. Sudarshan, "Completely positive dynamical semigroups of N-level systems," *Journal of Mathematical Physics* **17**, 821–825 (1976). doi:10.1063/1.522979.

[44] J. S. Bell, "On the Einstein Podolsky Rosen paradox," *Physics Physique Fizika* **1**, 195–200 (1964). doi:10.1103/PhysicsPhysiqueFizika.1.195.

[45] J. F. Clauser, M. A. Horne, A. Shimony and R. A. Holt, "Proposed Experiment to Test Local Hidden-Variable Theories," *Physical Review Letters* **23**, 880 (1969); erratum **24**, 549 (1970). doi:10.1103/PhysRevLett.23.880.

[46] C. Barceló, S. Liberati and M. Visser, "Analogue Gravity," *Living Reviews in Relativity* **8**, 12 (2005); updated review, arXiv:gr-qc/0505065, version 4 (2024). Analogue propagation geometry is distinguished from autonomous Einstein dynamics.

[47] T. Jacobson, "Thermodynamics of Spacetime: The Einstein Equation of State," *Physical Review Letters* **75**, 1260–1263 (1995). doi:10.1103/PhysRevLett.75.1260; arXiv:gr-qc/9504004.

[48] P. Virtanen et al., "SciPy 1.0: fundamental algorithms for scientific computing in Python," *Nature Methods* **17**, 261–272 (2020). arXiv:1907.10121; doi:10.1038/s41592-019-0686-2.