

Energy Does Not Survive Every Description: Build the Missing-Energy Ledger

Registered research scaffold — no numerical result

Brecht Corbeel

Absolute Digital Publishers

absolutedigitalpublishers.com

September 20, 2026

Abstract

A calorimeter’s unlogged heat, a spin’s shifting energy in a rotating apparatus, and a redshifted photon are routinely filed under one headline: energy conservation has failed somewhere. This draft registers, without simulating it, one exact operator meant to settle the first kind of case and to make precise why the same operator has no jurisdiction over the third. For a completely positive, trace-preserving channel Λ carrying a fine quantum description (Hilbert space $\mathcal{H}_f$, Hamiltonian H_f) to a coarser one ($\mathcal{H}_c$, H_c), the construction is the pullback defect $D_\Lambda := H_f - \Lambda^\dagger(H_c)$, where $\Lambda^\dagger$ is the channel’s Heisenberg dual, and its state-dependent reading $\Delta E_\Lambda(\rho) := \text{Tr}[\rho D_\Lambda]$. Both are shown, not assumed, to carry energy units and to satisfy $\Delta E_\Lambda(\rho) = 0$ for every fine state if and only if $D_\Lambda = 0$ as an operator identity. Three claims are proved by direct construction rather than asserted: a time-independent passive relabeling forces $D_\Lambda \equiv 0$ identically, for any unitary and any Hamiltonian, before any state is chosen; tracing out a bath returns $D_\Lambda = I_S \otimes H_B + H_{SB}$ exactly, reducing to a bath’s bare energy for an uncorrelated eigenstate and to an unresolvable joint quantity once system and bath are entangled; and a time-dependent frame carries a distinct, separately Hermitian correction term $i\hbar\dot{U}U^\dagger$ that a naive relabeling silently drops, worked exactly for a resonant two-level exchange (reproducing a 25.0ns swap time and a $10.3\mu\text{eV}$ energy swing for stated illustrative circuit-QED frequencies) and for the textbook rotating-frame reduction of a precessing spin. The construction is then shown not to extend to curved spacetime: a global conserved energy there requires an exact timelike Killing vector, which a coarse-graining channel cannot supply; gravitational energy has no local tensorial density, only coordinate-dependent pseudotensors; the ADM and Bondi masses of the same spacetime differ by exactly the flux radiated between them, never by an unexplained gap; and cosmological redshift is the observational

signature of an absent time-translation symmetry in Friedmann-Lemaître-Robertson-Walker spacetimes, not a failure of the local conservation law $\nabla_\mu T^{\mu\nu} = 0$, which continues to hold throughout. No simulation was run. Every quoted number is an exact analytic evaluation or a finite-precision structural check performed by an accompanying zero-dependency program, which also verifies the tensor-algebraic core of the Killing/current-conservation argument, the Friedmann fluid equation’s static and radiation limits, and a Noether-charge commutator on the solvable exchange model; the coupled-oscillator and spin-bath numerical comparison this construction is registered against, and the kill criterion that comparison would have to survive, remain future work, stated as such throughout. **Keywords:** quantum channel; Heisenberg

dual map; Noether’s theorem; Stinespring dilation; Lindblad equation; circuit quantum electrodynamics; Killing vector; ADM mass; Bondi mass; gravitational pseudotensor; cosmological redshift.

Series: Revolutionary Life 75, Series D, D08.

Status: Draft scaffold. The coupled-oscillator and spin-bath numerical comparison, the genuinely open (non-unitary) coarse-graining implementation, and external review are incomplete.

1 A Pullback Built From What a Coarser Account Cannot See

The object of this draft is a pair of running physical settings, treated separately throughout: a small quantum system coupled to a bosonic bath of oscillators, of the kind the Caldeira-Leggett model uses to study dissipation [11], and a small system coupled to a bath of spins, of the kind the central-spin/spin-bath literature studies for decoherence [12], each subjected in turn to a time-independent relabeling, a time-dependent frame,

and a genuinely information-discarding channel. The baseline against which every claim below is checked is Noether charge accounting, not a newly proposed conserved quantity: nothing here asserts that energy conservation has been extended, weakened, or generalized. What is built is a diagnostic for when a coarser description's own Hamiltonian disagrees with a finer one's, stated as an operator identity that can be evaluated exactly rather than estimated.

Fix a fine Hilbert space $\mathcal{H}_f$ and a Hermitian fine Hamiltonian H_f on it, in energy units throughout. For the running two-body setting, $\mathcal{H}_f = \mathcal{H}_S \otimes \mathcal{H}_B$ for a system S and a bath B (oscillators in one case, spins in the other), with

$$H_f = H_S \otimes I_B + I_S \otimes H_B + H_{SB},$$

three Hermitian terms: the system's own generator, the bath's own generator, and a coupling. A fine state is a density operator ρ on $\mathcal{H}_f$.

Definition 1.1 (Coarse-graining channel and its Heisenberg dual). *A coarse-graining channel is a completely positive, trace-preserving map $\Lambda : \mathcal{B}(\mathcal{H}_f)_* \rightarrow \mathcal{B}(\mathcal{H}_c)_*$ between trace-class operators on a fine space $\mathcal{H}_f$ and a coarse space $\mathcal{H}_c$, together with a Hermitian coarse Hamiltonian H_c on $\mathcal{H}_c$. Its Heisenberg dual $\Lambda^\dagger : \mathcal{B}(\mathcal{H}_c) \rightarrow \mathcal{B}(\mathcal{H}_f)$ is the unique bounded map satisfying*

$$\text{Tr} [\Lambda(\rho) O_c] = \text{Tr} [\rho \Lambda^\dagger(O_c)] \quad (1)$$

for every fine state ρ and every coarse operator O_c .

A time-independent unitary relabeling ($\mathcal{H}_c = \mathcal{H}_f$, $\Lambda(\rho) = W\rho W^\dagger$), the identity channel, and the partial trace over the bath ($\Lambda(\rho) = \text{Tr}_B[\rho]$, $\mathcal{H}_c = \mathcal{H}_S$) are all instances of Definition 1.1; Section 2 treats the first two and Section 3 the third. Appendix A shows, from an explicit Kraus decomposition, that $\Lambda^\dagger$ is always unital, Hermiticity-preserving, and positivity-preserving, and that Eq. (1) is exactly what an accompanying zero-dependency program verifies on an explicit random channel rather than assumes.

Definition 1.2 (The pullback defect and its reading).

$$D_\Lambda := H_f - \Lambda^\dagger(H_c), \quad \Delta E_\Lambda(\rho) := \text{Tr} [\rho D_\Lambda].$$

Proposition 1.1 (Units and self-adjointness). *D_Λ is Hermitian on $\mathcal{H}_f$ and carries units of energy; $\Delta E_\Lambda(\rho)$ is a real number in units of energy for every fine state ρ .*

Proof. H_f is Hermitian by hypothesis and in energy units. $\Lambda^\dagger$ is Hermiticity-preserving (Appendix A, Lemma A.2(ii)), so $\Lambda^\dagger(H_c)$ is Hermitian; H_c is in energy units and $\Lambda^\dagger$ introduces no dimensionful factor, so $\Lambda^\dagger(H_c)$ is in energy units. A difference of two Hermitian, same-unit operators is Hermitian and in that unit. $\Delta E_\Lambda(\rho) = \text{Tr}[\rho D_\Lambda]$ is then real because ρ and D_Λ are both Hermitian and $\rho \geq 0$, $\text{Tr} \rho = 1$. $\square$

Proposition 1.2 (Exact vanishing criterion). *$\Delta E_\Lambda(\rho) = 0$ for every fine state ρ if and only if $D_\Lambda = 0$ as an operator identity, i.e. if and only if $\Lambda^\dagger(H_c) = H_f$ exactly.*

Proof. The "if" direction is immediate. For "only if," Appendix A, Lemma A.1, proves by polarization that a bounded Hermitian operator is determined by its expectation value on every pure state, hence on every density operator; applying this to D_Λ gives the claim. $\square$

Proposition 1.2 is the whole of this construction's ambition: it licenses reading a single computed number, $\Delta E_\Lambda(\rho)$, as a certificate that two fully specified Hamiltonian descriptions of the identical physical situation agree exactly, never approximately and never as a statistical tendency over an ensemble. It does not, by itself, say anything about where a nonzero reading comes from; that is what Sections 2–5 establish case by case.

Lemma 1.1 (Expectation-value restatement). *For every fine state ρ ,*

$$\Delta E_\Lambda(\rho) = \langle H_f \rangle_\rho - \langle H_c \rangle_{\Lambda(\rho)}. \quad (2)$$

Proof. $\Delta E_\Lambda(\rho) = \text{Tr}[\rho H_f] - \text{Tr}[\rho \Lambda^\dagger(H_c)] = \langle H_f \rangle_\rho - \text{Tr}[\Lambda(\rho) H_c]$, using Eq. (1) on the second term, and $\text{Tr}[\Lambda(\rho) H_c] = \langle H_c \rangle_{\Lambda(\rho)}$ by definition of an expectation value. $\square$

Lemma 1.1 is the operational content behind the word "ledger": $\Delta E_\Lambda(\rho)$ is exactly the gap between what the fine description's own Hamiltonian assigns to ρ and what the coarse description's own Hamiltonian assigns to the state ρ has actually become under Λ . Nothing in Definition 1.2 or Proposition 1.2 adjudicates whether that gap, once nonzero, is a bookkeeping omission, a dropped correction term, or a question the construction was never entitled to ask; Sections 2–6 work through exactly that sorting, case by case, with the last of them stating precisely where the sorting stops.

2 The Identity a Renaming Must Satisfy

Before any coarse-graining that discards a degree of freedom is trusted, the construction of Section 1 has to survive the case where nothing at all is discarded: a fixed, time-independent change of labels on the same fine space. Model this as $\mathcal{H}_c = \mathcal{H}_f$, $\Lambda(\rho) := W\rho W^\dagger$ for a fixed unitary W , with the coarse Hamiltonian defined self-consistently as the same physical generator carried through the same relabeling,

$$H_c := W H_f W^\dagger.$$

This is not a special or contrived choice of H_c ; it is the only choice consistent with calling Λ a relabeling rather than a genuine change of physics, since H_c is required to generate, in the new labels, the identical dynamics H_f generates in the old ones.

Proposition 2.1 (Passive relabeling forces $D_\Lambda \equiv 0$). *For $\Lambda(\rho) = W\rho W^\dagger$ and $H_c = WH_fW^\dagger$, with W any unitary on $\mathcal{H}_f$ and H_f any Hermitian operator on $\mathcal{H}_f$,*

$$D_\Lambda = 0$$

identically, as an operator equation, before any state is chosen.

Proof. From the defining pairing Eq. (1), $\text{Tr}[W\rho W^\dagger O_c] = \text{Tr}[\rho W^\dagger O_c W]$ holds for every ρ and every O_c , so $\Lambda^\dagger(O_c) = W^\dagger O_c W$. Substituting $O_c = H_c = WH_fW^\dagger$,

$$\Lambda^\dagger(H_c) = W^\dagger(WH_fW^\dagger)W = (W^\dagger W)H_f(W^\dagger W) = H_f,$$

using only $W^\dagger W = I$. Hence $D_\Lambda = H_f - \Lambda^\dagger(H_c) = H_f - H_f = 0$. $\square$

Proposition 2.1 is checked, not merely asserted, by an accompanying program: for random Hermitian H_f and random unitary W at several dimensions, the matrix D_Λ built from Definition 1.2 comes out at machine precision, never as an approximation valid in some limit. This is the covariance requirement every candidate construction in this literature has to survive before its behavior under a genuine coarse-graining is worth trusting: a change of basis, a renaming of which combination of amplitudes is called "excited," or any other relabeling that discards no degree of freedom must never, by itself, manufacture a nonzero reading. A construction that returned $\Delta E_\Lambda(\rho) \neq 0$ purely from a renaming, with nothing thrown away, would be reporting an artifact of notation as though it were physics, and every subsequent section of this draft would be built on an unsound floor.

Remark 2.1 (The trivial floor, and what it is not). Taking $W = I$ in Proposition 2.1 recovers the identity channel, $\Lambda(\rho) = \rho$, $H_c = H_f$, for which $D_\Lambda = 0$ trivially. This special case is the operator-language shadow of the oldest secure fact the subject has: a system evolving under its own Hamiltonian, with nothing relabeled and nothing coarse-grained away, conserves that Hamiltonian's own expectation value exactly, because H_f commutes with itself and generates its own time evolution. This is an instance, at time-translation symmetry, of Noether's theorem connecting a continuous symmetry of the dynamics to a conserved quantity along it [1]. Nothing in this draft derives, extends, or needs a new proof of that theorem; it is assumed at this one degenerate limit, as the floor every other case below is measured against. Proposition 2.1 is a strictly different

statement: it is not that H_f is conserved under its own dynamics, but that the pullback defect vanishes under a passive relabeling for *any* H_f , including one whose own dynamics does something complicated with it.

3 What a Trace Over a Bath Actually Returns

Return to the genuine coarse-graining set aside in Section 2: a coarse observer who keeps the system S and discards the bath B entirely, $\Lambda(\rho) = \text{Tr}_B[\rho]$, $\mathcal{H}_c = \mathcal{H}_S$, $H_c = H_S$. Unlike the relabeling of Section 2, this Λ genuinely reduces the space of states, and it is the operation both running settings of Section 1 — a system coupled to an oscillator bath, a system coupled to a spin bath — are actually subjected to whenever "the bath" is declared out of scope.

Proposition 3.1 (Exact bath identity). *For $\Lambda = \text{Tr}_B$, $\mathcal{H}_c = \mathcal{H}_S$, $H_c = H_S$,*

$$D_\Lambda = I_S \otimes H_B + H_{SB}, \quad (3)$$

an exact operator identity, not an approximation valid in some regime.

Proof. The partial trace has Kraus operators $K_b = I_S \otimes \langle b|$ for an orthonormal bath basis $\{|b\rangle\}$, satisfying $\sum_b K_b^\dagger K_b = I_S \otimes I_B$ by completeness of the basis. The dual, from $\Lambda^\dagger(O_S) = \sum_b K_b^\dagger O_S K_b$ (Appendix A), is $\Lambda^\dagger(H_S) = \sum_b (I_S \otimes |b\rangle) H_S (I_S \otimes \langle b|) = H_S \otimes \sum_b |b\rangle \langle b| = H_S \otimes I_B$, the ordinary embedding of a system operator into the joint space. Subtracting from $H_f = H_S \otimes I_B + I_S \otimes H_B + H_{SB}$ gives Eq. (3) directly. $\square$

Both sides of Eq. (3) carry the same units by inspection, since H_B and H_{SB} are each already in energy units in the fine description; the identity is exact algebra, not a limit, and holds regardless of whether B is a single mode, a continuum of oscillators, or an extensive spin bath.

Corollary 3.1 (Uncorrelated bath). *If $H_{SB} = 0$ and $\rho = \rho_S \otimes |b\rangle \langle b|$ for any system state ρ_S and any bath eigenstate $|b\rangle$ with $H_B |b\rangle = E_b |b\rangle$, then $\Delta E_\Lambda(\rho) = E_b$ for every choice of ρ_S .*

Proof. $\Delta E_\Lambda(\rho) = \text{Tr}[\rho(I_S \otimes H_B)] = \text{Tr}_S[\rho_S] \langle b|H_B|b\rangle = 1 \cdot E_b$. $\square$

Corollary 3.1 is the case in which "missing energy" has an address: a fixed, named pile of energy sitting in a part of the world a coarse description agreed in advance not to track, entirely independent of what the system is doing. Nothing is unaccounted for here in any troubling sense; the coarse observer simply keeps no receipt for energy that was never claimed as theirs.

Remark 3.1 (Entangled bath, where the ledger stops resolving). Once $H_{SB} \neq 0$ and ρ is entangled across the S - B cut, $\text{Tr}[\rho H_{SB}]$ need not vanish, and, more importantly, it need not be assignable to "the system's energy" or "the bath's energy" individually: it is a property of the correlations in the joint state that a bipartite split has no further vocabulary to divide. Quantum thermodynamics has had to confront exactly this ambiguity in defining work, heat, and internal energy for open systems with non-negligible system-bath correlations [6], and the discipline adopted here is the one adopted there: $\Delta E_\Lambda(\rho)$ reports the total defect exactly and does not overreach into an assignment the formalism does not license. This is a stated boundary of the construction's interpretive reach, not a defect in Proposition 3.1, which remains exact regardless of how entangled ρ is.

An accompanying program builds Eq. (3) explicitly for a two-level system coupled to a three-level bath with random Hermitian H_S, H_B, H_{SB} , confirms the dual-map construction of $\Lambda^\dagger(H_S)$ against the direct embedding $H_S \otimes I_B$ at machine precision, and separately confirms Corollary 3.1 for a diagonalized H_B and several random ρ_S and bath eigenstates.

4 An Exchange Solved Without Approximation

Proposition 3.1 is exact but abstract; one solvable case makes it concrete and, at the same time, exhibits the Noether charge actually responsible for the coupling's structure. What follows is an exact analytic evaluation of a fully stated model, not a simulation and not a measurement.

Let S and B each be a two-level system on resonance, $H_S = (\hbar\omega/2)\sigma_z^S$, $H_B = (\hbar\omega/2)\sigma_z^B$, with $\sigma_z|e\rangle = +|e\rangle$, $\sigma_z|g\rangle = -|g\rangle$, coupled by an excitation-conserving exchange term of the kind underlying the Jaynes-Cummings model of a two-level system linearly coupled to a bosonic or spin-like partner [10],

$$H_{SB} = \hbar g(\sigma_+^S \sigma_-^B + \sigma_-^S \sigma_+^B). \quad (4)$$

Proposition 4.1 (Noether charge of the exchange). *The total excitation number $N = n_S \otimes I_B + I_S \otimes n_B$, with $n = \text{diag}(0, 1)$ in the $\{|g\rangle, |e\rangle\}$ basis, satisfies $[H_f, N] = 0$ exactly.*

Proof. H_S and H_B are each diagonal in the n -eigenbasis and so commute with N . H_{SB} raises one excitation while lowering the other, so $\sigma_+^S \sigma_-^B$ maps an N -eigenspace to itself (one excitation moves from B to S , the total is unchanged), and likewise for its conjugate term; an operator mapping every eigenspace of N to itself commutes with N . $\square$

Proposition 4.1 is the $U(1)$ symmetry — conservation of total excitation number — that makes the model solvable at all; it is verified as an exact zero matrix commutator by an accompanying program, not assumed from the form of Eq. (4) alone. Because H_{SB} conserves N , dynamics starting from one excitation stays confined to $\text{span}\{|e, g\rangle, |g, e\rangle\}$, on which $H_S + H_B$ vanishes (each basis state carries $+\hbar\omega/2 - \hbar\omega/2 = 0$) and H_{SB} acts as $\hbar g$ times the swap between the two. Starting from $|\psi(0)\rangle = |e, g\rangle$,

$$|\psi(t)\rangle = \cos(gt)|e, g\rangle - i \sin(gt)|g, e\rangle \quad (5)$$

solves $i\hbar d|\psi\rangle/dt = H_f|\psi\rangle$ exactly for all t , confirmed by direct substitution rather than by numerical integration: an accompanying program evaluates both sides of the Schrödinger equation from Eq. (5) and from the matrix H_f independently and finds them equal at machine precision at several sampled times.

From Eq. (5), $\langle H_f \rangle$ is exactly time-independent, as it must be for any state evolving under its own generator, and equal to $\langle e, g|H_f|e, g\rangle = 0$ since H_{SB} is purely off-diagonal in this basis. The system-only expectation is $\langle H_S \rangle(t) = (\hbar\omega/2)[\cos^2(gt) - \sin^2(gt)] = (\hbar\omega/2)\cos(2gt)$, so, by Proposition 3.1 restricted to this state and Lemma 1.1,

$$\Delta E_\Lambda(t) = \langle H_f \rangle - \langle H_S \rangle(t) = -\frac{\hbar\omega}{2}\cos(2gt). \quad (6)$$

At $t = 0$ this equals $-\hbar\omega/2$, the bath's own ground-state energy, matching Corollary 3.1 at the product state $|e, g\rangle$; at the swap time $t^* = \pi/(2g)$, the state is $|g, e\rangle$ up to phase and $\Delta E_\Lambda(t^*) = +\hbar\omega/2$, the bath's excited-state energy. A system-only observer watching only $\langle H_S \rangle(t)$ across this half-period sees the tracked energy swing by a full $\hbar\omega$ with no term inside H_S to explain it; Eq. (6) accounts for the entire swing at every instant, not only at the two endpoints.

An illustrative numerical evaluation. Taking $\omega/2\pi = 5$ GHz and $g/2\pi = 10$ MHz, both inside the range of frequencies and coupling strengths reported across circuit quantum electrodynamics architectures [7], the swap time is $t^* = \pi/(2g) = 25.0$ ns and the ledger's swing amplitude is $\hbar\omega/2 \approx 1.66 \times 10^{-24}$ J ≈ 10.3 μ eV, both exact arithmetic evaluations of Eqs. (6) reproduced by the accompanying program to more than ten significant figures; neither number is a simulated or measured quantity, and both are offered as illustrative of the construction's scale rather than as a claim about any specific measured device.

5 The Correction a Time-Dependent Relabeling Cannot Skip

A third regime sits between the two already treated and needs its own construction rather than an application of Definition 1.2. Nothing is discarded — $\mathcal{H}_c = \mathcal{H}_f$ still — but the relabeling itself now carries explicit time dependence, because the coarse observer describes the identical system from a frame in motion relative to the one defining H_f : a rotating platform, a driven interaction picture, a ramping field. This is genuinely a different operation from Sections 2–3, and treating it as a third instance of D_Λ evaluated on some state would misdescribe what is actually going on, for a reason made precise in Remark 5.1 below.

Let $U(t)$ be a time-dependent unitary and $|\psi'(t)\rangle := U(t)|\psi(t)\rangle$, with $|\psi(t)\rangle$ solving $i\hbar d|\psi\rangle/dt = H_f|\psi\rangle$. Differentiating the product and using $|\psi(t)\rangle = U(t)^\dagger|\psi'(t)\rangle$,

$$i\hbar \frac{d|\psi'\rangle}{dt} = (UH_fU^\dagger + i\hbar \dot{U}U^\dagger)|\psi'\rangle,$$

so the operator actually generating $|\psi'(t)\rangle$'s evolution is

$$H_c(t) := UH_fU^\dagger + i\hbar \dot{U}U^\dagger, \quad (7)$$

not the bare conjugate $UH_fU^\dagger$ alone. Write $\Delta_{\text{frame}}(t) := i\hbar \langle \dot{U}(t)U(t)^\dagger \rangle$ for the correction term's expectation value.

Proposition 5.1 (The frame term is Hermitian). *$i\hbar \dot{U}(t)U(t)^\dagger$ is Hermitian and carries units of energy, for any differentiable unitary path $U(t)$.*

Proof. Differentiating $U(t)U(t)^\dagger = I$ gives $\dot{U}U^\dagger + U\dot{U}^\dagger = 0$, i.e. $\dot{U}U^\dagger = -U\dot{U}^\dagger = -(\dot{U}U^\dagger)^\dagger$: $\dot{U}U^\dagger$ is anti-Hermitian. Then $(i\hbar \dot{U}U^\dagger)^\dagger = -i\hbar (\dot{U}U^\dagger)^\dagger = -i\hbar (-\dot{U}U^\dagger) = i\hbar \dot{U}U^\dagger$. Units: $\dot{U}U^\dagger$ is a rate (time⁻¹, since U is dimensionless), and $\hbar \times \text{rate}$ is energy. $\square$

An accompanying program confirms Proposition 5.1 by central finite difference on a genuinely time-dependent unitary path built from two non-commuting generators, at several sample times, to the expected finite-difference precision.

Remark 5.1 (Why this is not an instance of D_Λ). Fix any single instant t and set $\Lambda_t(\rho) := U(t)\rho U(t)^\dagger$, with the naive pulled-back Hamiltonian $H_c^{\text{naive}}(t) := U(t)H_fU(t)^\dagger$ standing in for "the coarse Hamiltonian." By Proposition 2.1, applied at that fixed t , $D_{\Lambda_t} = H_f - \Lambda_t^\dagger(H_c^{\text{naive}}(t)) = 0$ identically, for every t , regardless of how fast or how nontrivially U depends on time. The pullback defect of Section 1 is blind to time dependence by construction, because it compares two Hamiltonians at one instant, not two generators of dynamics

across an interval. The genuine discrepancy a moving frame introduces is not a nonzero D_{Λ_t} ; it is the gap between $H_c^{\text{naive}}(t)$ and the operator in Eq. (7) that actually generates $|\psi'(t)\rangle$'s motion, namely $\Delta_{\text{frame}}(t)$ itself. This is exactly the distinction the object of Section 1 requires: a time-independent relabeling, a time-dependent frame, and a genuinely open coarse-graining are three different operations, and only the first and third are instances of D_Λ .

A case tying the correction to a century-old reduction. Let $H_f = (\hbar\omega_0/2)\sigma_z$, a spin precessing at its Larmor frequency, and $U(t) = \exp(i\omega_r t \sigma_z/2)$, a frame rotating about the same axis at rate ω_r , the transformation underlying Rabi's original treatment of a spin in a gyrating field [8] and Bloch's phenomenological equations for the same reduction [9]. Since $U(t)$ is built from σ_z , it commutes with H_f , so $UH_fU^\dagger = H_f$ exactly; but $\dot{U}U^\dagger = i\omega_r \sigma_z/2$ exactly, giving

$$H_c(t) = H_f + i\hbar \left(i\frac{\omega_r}{2} \sigma_z \right) = \frac{\hbar(\omega_0 - \omega_r)}{2} \sigma_z, \quad (8)$$

the standard detuned Hamiltonian governing spin precession in a frame rotating with an applied field. Two limits confirm the recovery rather than an accident of the specific numbers chosen: at $\omega_r = 0$, $\Delta_{\text{frame}} \equiv 0$ and $H_c = H_f$, the already-established baseline of Remark 2.1; at $\omega_r = \omega_0$, exact corotation, $H_c = 0$, the ordinary resonance condition. An accompanying program builds Eq. (8) directly from Eq. (7) at several ω_r including these two limits and finds agreement with the closed form at machine precision.

The failure mode this section exists to name sits exactly at the line drawn in Remark 5.1: an observer who correctly computes $UH_fU^\dagger$ — a legitimate, Proposition 2.1-licensed operator identity — and then mistakes that operator for the rotating frame's own generator, rather than using Eq. (7), will find a discrepancy with no term in their own bookkeeping to explain it. That discrepancy is not missing energy in the sense of Section 3; it is a dropped term whose physical origin is whatever classical control enacts $U(t)$ in the first place, already doing ordinary, accountable work on the description rather than on the system's conserved energy.

6 Where a Killing Vector Does the Work No Channel Can

Nothing in Sections 1–5 touches curved spacetime, and this section exists to keep it that way explicitly, because a construction built around the phrase "missing energy" is exactly the kind of object this subject's cranks reach for. The settled physics is stated first, and nothing below qualifies it: the local conservation law of general

relativity,

$$\nabla_\mu T^{\mu\nu} = 0, \quad (9)$$

is a differential identity forced by the Einstein field equations through the contracted Bianchi identity, and it holds in every spacetime — flat, curved, static, or radiating — without exception. Both sides of Eq. (9) are tensors of the same type at a point, so dimensional and index consistency is automatic; nothing that follows questions Eq. (9) or needs it weakened.

A local law is not, by itself, a global conserved number. That requires a timelike Killing vector field ξ^μ , satisfying $\nabla_{(\mu}\xi_{\nu)} = 0$, an exact continuous isometry of the full geometry rather than a convenient coordinate choice.

Proposition 6.1 (Killing current conservation). *If $\nabla_\mu T^{\mu\nu} = 0$ and ξ^μ is a Killing vector, then $J^\mu := T^{\mu\nu}\xi_\nu$ satisfies $\nabla_\mu J^\mu = 0$.*

Proof. $\nabla_\mu J^\mu = (\nabla_\mu T^{\mu\nu})\xi_\nu + T^{\mu\nu}\nabla_\mu\xi_\nu = T^{\mu\nu}\nabla_\mu\xi_\nu$, using Eq. (9). The Killing equation makes $\nabla_\mu\xi_\nu$ purely antisymmetric, since its symmetric part vanishes by hypothesis. Relabeling the dummy indices $\mu \leftrightarrow \nu$ in $T^{\mu\nu}\nabla_\mu\xi_\nu$ and using $T^{\mu\nu} = T^{\nu\mu}$, $\nabla_\nu\xi_\mu = -\nabla_\mu\xi_\nu$, gives $T^{\mu\nu}\nabla_\mu\xi_\nu = -T^{\mu\nu}\nabla_\nu\xi_\mu$, so the quantity equals its own negative and vanishes. $\square$

The step doing the work in Proposition 6.1 — a symmetric tensor contracted against an antisymmetric one vanishes identically — is exactly what an accompanying program verifies, in exact integer arithmetic rather than floating-point tolerance, for random symmetric and antisymmetric matrices at several sizes, together with the degenerate case $\nabla_\mu\xi_\nu \equiv 0$: a constant Killing field in Cartesian Minkowski coordinates, for which the same identity holds trivially and J^μ reduces to the ordinary flat-spacetime conserved energy current with no Killing-equation subtlety needed at all. This is the exact recovery of ordinary conserved energy in the flat limit, and the $H = 0$ limit of Proposition 6.2 below is the analogous static-case recovery.

Given $\nabla_\mu J^\mu = 0$, the charge $Q = \int_\Sigma J^\mu d\Sigma_\mu$ on a Cauchy slice Σ is independent of which slice is chosen, by the divergence theorem, provided no flux escapes through the boundary. Komar’s construction turns this into a covariant surface-integral charge built directly from ξ^μ [14], and the general statement tying conserved quantities to Killing vectors is standard material in the field’s own references [13].

Remark 6.1 (No local density, only coordinate-dependent pseudotensors). Away from an exact Killing vector, general relativity does not generally supply a local, coordinate-independent tensorial density for the energy carried by the gravitational field itself; only pseudotensors serve that role, and a pseudotensor’s value at a point depends on the coordinate system

used to compute it, a point treated carefully in modern reviews of quasi-local energy-momentum [18]. This is a fact about gravitational field energy specifically, distinct from Eq. (9), which is exact and coordinate-independent for matter stress-energy throughout.

Two well-separated quantities follow from weaker, asymptotic versions of the Killing requirement, and neither should be conflated with the other or with the Killing charge above. At spatial infinity of an asymptotically flat spacetime, the Arnowitt-Deser-Misner construction defines a total mass-energy from a surface integral needing only asymptotic flatness, not an exact Killing symmetry everywhere [15], and Witten’s spinor argument proves this ADM mass is never negative for matter satisfying the dominant energy condition [16]. At null infinity, a related but distinct quantity, the Bondi mass, is defined from data on outgoing null hypersurfaces; Bondi, van der Burg, and Metzner’s original construction shows it is non-increasing along retarded time, decreasing by exactly the flux of energy radiated away [17]. An accompanying program checks the structural content of that mass-loss law — that a manifestly non-negative flux forces a monotonic, non-increasing mass function, with the total drop equal to the total integrated flux — on an explicit illustrative flux profile; this checks the shape of the argument, not Bondi’s specific coefficients, which are not reproduced here. The ADM and Bondi masses of the same spacetime agree before any radiation has left and differ afterward by exactly this radiated flux: a computed, physically real outgoing quantity, never an unexplained gap between two competing bookkeepers.

Proposition 6.2 (FLRW conservation, the static limit, and the redshift exponent). *For a perfect fluid with equation of state $p = w\rho$ in a Friedmann-Lemaître-Robertson-Walker spacetime with scale factor $a(t)$ and Hubble rate $H = \dot{a}/a$, Eq. (9) gives*

$$\dot{\rho} + 3H(1+w)\rho = 0, \quad (10)$$

solved by $\rho(a) = \rho_0(a/a_0)^{-3(1+w)}$. For radiation, $w = 1/3$, giving $\rho \propto a^{-4}$; with comoving photon number conserved, $n_\gamma \propto a^{-3}$, so energy per photon scales as $\rho/n_\gamma \propto a^{-1}$, exactly the redshift relation $E_{\text{obs}}/E_{\text{emit}} = a_{\text{emit}}/a_{\text{obs}} = 1/(1+z)$. Setting $H = 0$ in Eq. (10) gives $\dot{\rho} = 0$ identically, for any w : the static-spacetime recovery of ordinary conservation.

Proof. Eq. (10) is the $\nu = 0$ component of Eq. (9) evaluated on the FLRW metric with $T^\mu{}_\nu = \text{diag}(-\rho, p, p, p)$, standard in general relativity and cosmology texts [13]. Writing $d\rho/da = \dot{\rho}/\dot{a}$, Eq. (10) becomes $a d\rho/da = -3(1+w)\rho$, solved by the stated power law since $a d(a^n)/da = na^n$ for $n = -3(1+w)$. The remaining claims follow by substituting $w = 1/3$ and $H = 0$ respectively. $\square$

An accompanying program confirms the power-law solution of Eq. (10) by finite difference for radiation, dust, and a curvature-like fluid, confirms the $H = 0$ substitution gives zero exactly for each, and confirms the redshift exponent bookkeeping $(-4 - (-3) = -1)$ and a sample ratio $E_{\text{obs}}/E_{\text{emit}} = 1/(1+z)$ in exact rational arithmetic.

FLRW spacetimes in comoving coordinates have no timelike Killing vector: the metric depends explicitly on time through $a(t)$, and no coordinate change removes that dependence globally. There is consequently no FLRW analog of the Killing charge, the ADM mass, or the Bondi mass for a quantity like the universe's total photon population, and cosmological redshift is the observational signature of that absent symmetry, not a violation of Eq. (9), which Proposition 6.2 shows continues to hold throughout. Objections that treat this redshift as though "expanding space" itself were doing physical work smuggle in properties general relativity does not assign to the expansion, a confusion addressed directly in the literature on this specific point [19, 20].

What is shared with Sections 1–5, and what is not. The resemblance is real and stops exactly here: both settings share the logical shape that a globally conserved charge requires the generator of the relevant symmetry to be exactly present in the description doing the accounting, and that a nonzero "defect" is only meaningful once that absence is named precisely — a discarded bath Hamiltonian in Section 3, a missing timelike Killing vector here. The load-bearing difference is that Λ in Definition 1.1 is a choice: a laboratory observer decides how much of a bath to track, and a more generous choice returns a smaller or zero defect for the identical physical situation, as Corollary 3.1 already showed. The absence of a timelike Killing vector in an FLRW spacetime is not a choice available to any observer; it is a geometric fact about the manifold that no coarse-graining channel, however generously defined, can restore. Nothing in D_Λ plays the role of a finer channel that would recover a redshifted photon's energy the way tracing further into a bath recovers a swapped excitation in Section 4, and this draft does not attempt to unify the two pictures into one formalism.

7 The Comparison This Result Would Have to Survive

The construction is rejected as a physically useful diagnostic, rather than merely a correct but idle piece of algebra, if either of two things happens once it is tested beyond the cases already checked in Sections 2–4: D_Λ fails to vanish under a passive relabeling in some tested configuration, or $\Delta E_\Lambda(\rho)$ disagrees with the energy actually absent from a retained description, once that absent energy can be computed independently of

the operator itself.

The first half is already closed, not open. Proposition 2.1 is exact, finite-dimensional linear algebra, true for every unitary and every Hermitian operator before any state or dimension is fixed; an accompanying program confirms it at several dimensions to machine precision. A numerical test that found $D_\Lambda \neq 0$ under a passive relabeling would not have discovered new physics — it would have found a bug in that test's own implementation of $\Lambda^\dagger$, since the identity does not depend on any physical assumption a richer simulation could violate. Reporting this half of the criterion as though it awaited empirical confirmation would misstate what kind of claim Proposition 2.1 is.

The second half is the one with actual empirical content, and it has been checked only in the single exactly solvable exchange of Section 4 and the abstract bath identity of Section 3, both closed, unitary, exactly diagonalizable systems. The registered numerical program — not run in this draft — compares $\Delta E_\Lambda(\rho(t))$, computed from D_Λ , against a directly computed energy account built independently from a simulated trajectory (tracking each oscillator mode's or each bath spin's energy explicitly, without reference to the operator), for the two running settings of Section 1: a small system coupled to several bosonic modes, and a small system coupled to several spins, each under more than one coarse-graining of the same underlying dynamics — different choices of which modes or spins count as "system," and, unlike every construction proved exactly above, at least one genuinely open coarse-graining generated by a Lindblad dissipator rather than a closed unitary dilation. Agreement in the closed, unitary case is not automatic evidence for the open case: Proposition 3.1 and its proof use a unitary H_f and an exact partial trace throughout, and a Lindblad-generated channel's dual map, while still unital and positivity-preserving by the general argument of Appendix A, has not been shown here to admit a Hamiltonian H_c whose pullback reproduces an independently computed energy account rather than merely an internally consistent one.

Whether the construction survives that comparison is exactly the sharpest thing this draft does not yet know, and Section 9 names precisely what stands between the checks already closed above and a completed run of it.

8 Four Places This Construction Could Be Mistaken for Another

The adjoint identity behind D_Λ is elementary, which makes it easy to mistake for something already established elsewhere or, worse, for something this draft has

no license to claim. Four specific collisions are checked here rather than left implicit.

A bound over a family, against one number for a named pair. Cirstoiu, Korzekwa, and Jennings addressed how far a quantum channel invariant under a group G associated with a Noether charge can still deviate from conserving that charge, bounding the deviation, over the entire convex set of G -symmetric channels, by a measure of how far the channel sits from being a closed, unitary evolution [2]. That is a bound over a family of channels sharing a symmetry; D_Λ is a single exact number for one fully named channel and one fully named pair of Hamiltonians, with no symmetry of Λ itself assumed anywhere in Sections 1–5. The two are not competing answers to the same question: theirs bounds a whole class, this one evaluates a single checkable instance, and nothing here extends, tightens, or reproduces their bound.

A flow on couplings, against a channel on states. A companion piece in this program treats Wilson’s renormalization group, a flow on a space of coupling constants generated by integrating out short-distance degrees of freedom, as a discrete replicator process near a fixed point, graded claim by claim against natural selection’s own logic. That construction’s object is a trajectory through coupling space; nothing in it is a completely positive map between density operators, and nothing here computes a beta function, a fixed point, or a replicator equation. Renormalization-group coarse-graining and the channel Λ of Definition 1.1 both discard information, but discarding information is the only thing the two operations share: one acts on a space of theories, the other on a space of quantum states of one fixed system, and an energy-defect operator for the second says nothing about flow near a fixed point of the first.

An accounting identity, against an argument about time’s direction. A second companion piece follows the thermodynamic arrow of time back through Boltzmann and Loschmidt’s dispute, Landauer’s erasure bound, and fluctuation theorems, asking what licenses treating dissipation as a one-way process at all. D_Λ needs none of that: Proposition 2.1 and Proposition 3.1 are time-symmetric algebraic identities, holding identically whether Λ is thought of as running forward or backward, and nothing in this draft invokes entropy production, an irreversibility bound, or a preferred temporal direction. A reader looking for this draft to explain why energy bookkeeping and time’s arrow are connected will not find that question addressed here at all.

A selection functional, against a bookkeeping operator. A third companion piece reads decoherence, cosmic mode selection, and gravitational collapse as one cascade of selection functionals, audited against the biological conditions for selection to operate. Section 6’s

treatment of cosmological redshift is not a fourth stage of that cascade: it makes no claim about variation, heredity, or differential persistence among cosmological modes, and its only content is the narrower, purely geometric fact that an FLRW spacetime lacks the time-like Killing vector a global conserved energy would require. The two pieces can discuss the same equation, $\dot{\rho} + 3H(1+w)\rho = 0$, without making the same argument about it.

None of Stinespring’s dilation theorem, Lindblad’s generator classification, or the Gorini-Kossakowski-Sudarshan classification for finite-dimensional systems is claimed as this draft’s own contribution [3–5]; they are the inherited machinery Appendix A leans on to justify that $\Lambda^\dagger$ behaves as Definition 1.1 requires. What remains, once these four collisions and this inherited machinery are set aside, is narrow and stated precisely: the pullback defect itself, proved to vanish under passive relabeling and proved exact under partial trace; the explicit separation of a time-dependent frame’s correction term from either of those two cases; and the boundary drawn in Section 6 between a coarse-graining choice and an absent spacetime symmetry.

9 What This Operator Has Not Yet Been Asked to Survive

Everything computed in this draft is an exact operator identity, an exact analytic evaluation of a solvable model, an exact integer-arithmetic tensor identity, or a finite-difference check of a closed-form expression against itself. None of it is the coupled-oscillator and spin-bath comparison Section 7 describes, and this section states, without qualification, what stands between the two.

The first gate is the numerical program itself. No many-mode coupled-oscillator or many-spin-bath dynamics has been implemented beyond the single exactly diagonalizable two-level exchange of Section 4. Until it exists, the sentence in Section 7 describing what the registered program would compare — $\Delta E_\Lambda(\rho(t))$ against an independently tracked energy account — names a comparison that has not been run for any system larger than two qubits.

The second gate is whether the construction of Section 1 extends usefully past a closed, unitary dilation at all. Every proved case in this draft — passive relabeling, partial trace, the exact exchange — has Λ built from a unitary or an exact partial trace on a closed joint system. Appendix A shows $\Lambda^\dagger$ is unital, Hermiticity-preserving, and positivity-preserving for any completely positive, trace-preserving map, including one generated by a Lindblad dissipator, but this draft has not exhibited a genuinely open channel together with a coarse Hamiltonian H_c whose pullback re-

produces anything an independent observer would recognize as "the missing energy," rather than an internally consistent number with no external check. The claim in Section 7 that agreement in the closed case is not automatic evidence for the open case remains, for this reason, the operative uncertainty rather than a rhetorical caution.

The third gate is the direct energy account itself. A number computed independently of D_Λ — each oscillator mode's or bath spin's energy, tracked explicitly from a simulated trajectory — does not yet exist for any system in this draft's registered scope. Without it, there is nothing for $\Delta E_\Lambda(\rho(t))$ to be compared against beyond the two cases, Corollary 3.1 and Eq. (6), where the comparison could be done by hand because the model was small enough to diagonalize exactly.

The fourth gate is the kill criterion's empirically meaningful half, named in Section 7: whether $\Delta E_\Lambda(\rho)$ continues to equal the independently computed missing energy once numerical integration error, an imperfect dilation, or a genuinely open channel are in play, rather than only in the exactly solvable cases checked here. The first half of that same criterion — passive relabeling leaving D_Λ nonzero — is not an open gate; Proposition 2.1 closes it as exact algebra, and any implementation reporting otherwise would be reporting a bug in itself, a distinction Section 7 draws explicitly so that closing gates one through three is not mistaken for reopening a settled one.

The fifth gate is scope inside Section 5 itself: the rotating-frame recovery of Eq. (8) is worked only for a generator $U(t)$ built from σ_z , which commutes with H_f and therefore leaves the bare pullback $UH_fU^\dagger$ unchanged. A genuinely non-commuting time-dependent frame, where both terms of Eq. (7) vary nontrivially and neither vanishes on its own, has not been worked as a second concrete case, and the claim that Remark 5.1's distinction generalizes beyond the commuting example is, for now, an extrapolation from one solved instance rather than a second checked one.

The sixth gate is the priority search of Section 8: four specific, named collisions were checked directly, not an exhaustive systematic review of the open-quantum-systems, quantum-thermodynamics, and mathematical-relativity literatures each touches. The claim that this draft's narrow contribution — the pullback defect itself and the three-way separation of Section 5 — has not been found elsewhere is only as strong as the four comparisons actually made.

The seventh gate is external review by a named subject-matter reviewer with expertise spanning open quantum systems, quantum thermodynamics, and the conserved-quantities literature of mathematical relativity, who has not yet been approached. Before any statement in this draft is reported as a finding rather than a registered scaffold, all seven gates above must close:

the coupled-oscillator and spin-bath simulation; a working construction for a genuinely open coarse-graining channel; an independent direct energy account; the empirical half of the kill criterion tested against both; a second, non-commuting worked frame example; a completed priority search; and external review. As of this draft, all seven are open, and any one of them being open is sufficient reason to withhold a numerical claim.

10 What One Exact Operator Is Entitled to Claim

Sorted by what actually supports each: that a completely positive, trace-preserving channel admits a Heisenberg dual satisfying Eq. (1), and that this dual is unital, Hermiticity-preserving, and positivity-preserving, is inherited machinery, owed to Stinespring, Lindblad, and Gorini, Kossakowski, and Sudarshan [3–5], proved again in Appendix A rather than merely cited. That energy is conserved under a system's own unitary evolution with nothing coarse-grained away is likewise inherited, from Noether's theorem at the degenerate floor of Remark 2.1 [1]. What belongs to this draft is narrower and stated three times, once for each regime: the pullback defect D_Λ itself, proved rather than assumed to vanish identically under passive relabeling (Proposition 2.1) and proved exact under partial trace (Proposition 3.1), with both endpoints of the exactly solvable exchange of Section 4 matching the closed form to more than ten significant figures; the explicit separation, proved rather than asserted, of a time-dependent frame's own correction term $\Delta_{\text{frame}}(t)$ from anything D_Λ itself can register (Remark 5.1); and the boundary of Section 6, where the same logical shape — a conserved charge requires the generator of its symmetry to be present in the accounting — runs into a spacetime that simply lacks that generator, a geometric fact no coarse-graining channel can supply.

The kill criterion of Section 7 is only half addressed, and the half that remains open is the one carrying empirical content. Passive relabeling does not manufacture a defect: Proposition 2.1 closes that question as algebra, not as a numerical tendency that could still be overturned by a larger simulation. Whether $\Delta E_\Lambda(\rho)$ continues to equal an independently computed missing energy once a genuinely open coarse-graining channel and a many-mode system replace the closed, exactly solvable cases checked here is not settled, and Section 9 names this as the sharpest of the seven open items rather than a closed question dressed up as one.

What survives, narrowly, is a sorting procedure rather than a new conserved quantity: a reported energy defect is either exactly zero because nothing was discarded, as Section 2 always gives; a computable, boundedly interpretable number attributable

to a named bath and interaction, as Section 3 gives for an uncorrelated bath and refuses to over-resolve for an entangled one; a dropped inertial term a moving frame's own control system is already paying for, as Section 5 gives; or a question this construction has no jurisdiction to answer at all, because the missing symmetry belongs to a spacetime this operator carries no coordinates for, as Section 6 gives. The calorimeter's unlogged heat named at the outset of this program is the second case. A spinning apparatus's shifting reading is the third. A redshifted photon is the fourth. What this draft adds to a subject already possessing Noether's theorem, Stinespring's dilation, and Killing's equation is not a new way any of them conserves energy; it is one checkable procedure for deciding, before the word "missing" is written down, which of the four a given claim actually is.

A Two Lemmas the Construction Leans On

Section 1 states, without proof at the point of use, that a Hermitian operator is determined by its expectation values and that a channel's Heisenberg dual is well-behaved. Both are proved here.

Lemma A.1 (An operator is determined by its expectation values). *Let A, B be bounded Hermitian operators on a Hilbert space $\mathcal{H}$. If $\langle \psi | A | \psi \rangle = \langle \psi | B | \psi \rangle$ for every $|\psi\rangle \in \mathcal{H}$, then $A = B$.*

Proof. Let $C = A - B$, Hermitian, with $\langle \psi | C | \psi \rangle = 0$ for every $|\psi\rangle$. For any $|\phi\rangle, |\chi\rangle \in \mathcal{H}$, the polarization identity gives

$$\begin{aligned} \langle \phi | C | \chi \rangle &= \frac{1}{4} \left[\langle \phi + \chi | C | \phi + \chi \rangle - \langle \phi - \chi | C | \phi - \chi \rangle \right. \\ &\quad \left. - i \langle \phi + i\chi | C | \phi + i\chi \rangle + i \langle \phi - i\chi | C | \phi - i\chi \rangle \right], \end{aligned}$$

and every term on the right vanishes by hypothesis, so $\langle \phi | C | \chi \rangle = 0$ for all ϕ, χ , hence $C = 0$. This proves the claim for pure states; for a general density operator $\rho = \sum_i p_i |\psi_i\rangle \langle \psi_i|$, $\text{Tr}[\rho C] = \sum_i p_i \langle \psi_i | C | \psi_i \rangle = 0$ follows by linearity, so testing over density operators is equivalent to testing over pure states. $\square$

Proposition 1.2 applies Lemma A.1 to $A = D_\Lambda$, $B = 0$: since every pure state is a density operator, $\Delta E_\Lambda(\rho) = 0$ for every ρ is exactly the hypothesis of the lemma, giving $D_\Lambda = 0$.

Lemma A.2 (The Heisenberg dual is well-behaved). *Let $\Lambda(\rho) = \sum_i K_i \rho K_i^\dagger$ be completely positive and trace-preserving, so $\sum_i K_i^\dagger K_i = I_f$, and define $\Lambda^\dagger(O) := \sum_i K_i^\dagger O K_i$. Then: (i) $\Lambda^\dagger(I_c) = I_f$ (unital); (ii) $\Lambda^\dagger(O)$ is Hermitian whenever O is; (iii) $\Lambda^\dagger(O) \succeq 0$ whenever $O \succeq 0$; (iv) $\text{Tr}[\Lambda(\rho)O] = \text{Tr}[\rho \Lambda^\dagger(O)]$ for every ρ, O .*

Proof. (i) $\Lambda^\dagger(I_c) = \sum_i K_i^\dagger K_i = I_f$ by trace preservation of Λ . (ii) $\Lambda^\dagger(O)^\dagger = \sum_i (K_i^\dagger O K_i)^\dagger = \sum_i K_i^\dagger O^\dagger K_i = \Lambda^\dagger(O)$ when $O^\dagger = O$. (iii) For any $|\psi\rangle$, $\langle \psi | \Lambda^\dagger(O) | \psi \rangle = \sum_i \langle \psi | K_i^\dagger O K_i | \psi \rangle = \sum_i \langle K_i \psi | O | K_i \psi \rangle \geq 0$ since $O \succeq 0$. (iv) $\text{Tr}[\Lambda(\rho)O] = \text{Tr}[\sum_i K_i \rho K_i^\dagger O] = \sum_i \text{Tr}[K_i^\dagger O K_i \rho] = \text{Tr}[\rho \sum_i K_i^\dagger O K_i] = \text{Tr}[\rho \Lambda^\dagger(O)]$, using cyclicity of the trace. $\square$

Every dual map used in this draft — conjugation by $W^\dagger$ in Section 2, the partial-trace embedding in Section 3 — is the Kraus-operator construction of Lemma A.2 for a specific choice of K_i , so properties (i)–(iv) apply directly rather than needing to be re-derived case by case. An accompanying program builds an explicit random channel from a Stinespring dilation (an isometry $V : \mathcal{H}_f \rightarrow \mathcal{H}_c \otimes \mathcal{H}_E$, with $K_e := (I_c \otimes |e\rangle)V$) and confirms all four properties numerically: property (iii) is checked not by computing eigenvalues but by comparing $\langle \psi | \Lambda^\dagger(|v\rangle \langle v|) | \psi \rangle$ against the manifestly non-negative sum $\sum_e |\langle v | K_e | \psi \rangle|^2$ directly, which the proof of (iii) shows these must equal.

B Mapping Each Claim to Its Check

`check_missing_energy_operator.py` performs no coupled-oscillator or spin-bath simulation and implements no genuinely open coarse-graining channel. In its default `-self-check` mode it performs the following checks, each described in the sections above:

1. Lemma A.2's four dual-map properties (unitality, Hermiticity preservation, positivity preservation, and the defining pairing Eq. (1)), built from an explicit random Stinespring dilation at dimensions $d_f = 4$, $d_c = 3$, $d_e = 3$ over five trials, with worst deviations at or below 10^{-14} for every property;
2. Proposition 2.1, $D_\Lambda \equiv 0$ under passive relabeling, checked for random Hermitian generators and random unitaries at dimensions 2 through 5, with worst deviation 7.8×10^{-15} ;
3. Proposition 3.1 and Corollary 3.1, checked for a two-level system coupled to a three-level bath: trace preservation, the embedding $\Lambda^\dagger(H_S) = H_S \otimes I_B$, the full identity $D_\Lambda = I_S \otimes H_B + H_{SB}$, and the uncorrelated-bath limit $\Delta E_\Lambda(\rho) = E_b$, each at or below 10^{-8} deviation;
4. Proposition 4.1 ($[H_f, N] = 0$ exactly, zero to machine precision), direct substitution of Eq. (5) into the Schrödinger equation (deviation 7×10^{-43}), and Eq. (6) against the same substitution, together with the illustrative numbers $t^* = 25.0$ ns and a ledger swing amplitude of 1.6565×10^{-24} J = $10.339 \mu\text{eV}$ for $\omega/2\pi = 5$ GHz, $g/2\pi = 10$ MHz;

5. Proposition 5.1, checked by finite difference on a genuinely time-dependent unitary path (worst Hermiticity deviation 2.5×10^{-10}), and Eq. (8) at $\omega_r \in \{0, 0.4, \omega_0, 2.0\}$ against the direct construction of Eq. (7), including both required limits, at or below 10^{-9} deviation;
6. Proposition 6.1’s algebraic core, a symmetric tensor contracted against an antisymmetric one, verified in exact integer arithmetic (not floating-point tolerance) at sizes 3, 4, and 5 over six trials each, always exactly zero, together with the exact flat-spacetime degenerate case;
7. Proposition 6.2, checked by finite difference for radiation, dust, and a curvature-like fluid (relative deviations near 10^{-10}), the exact $H = 0$ static limit, and the exact rational-arithmetic redshift exponent and sample ratio $E_{\text{obs}}/E_{\text{emit}} = 1/(1+z) = 1/3$ for $a_{\text{emit}} = 1$, $a_{\text{obs}} = 3$;
8. the structural content of the Bondi mass-loss discussion in Section 6: an illustrative non-negative flux profile forces a monotonically non-increasing mass function (worst increase -7×10^{-14} , i.e. none), with the total drop matching the total integrated flux to 10^{-9} .

The saved output, `scaffold_validation_output.txt`, is the program’s real captured standard output from a `-self-check` run; `scientific_result` and `energy_ledger_simulation_run` are both recorded as `false`, and no field reports a coarse-grained energy defect for any many-body system as a physical measurement. The program’s `-simulate` option exits with status 2 and prints the seven-item list of unimplemented prerequisites matching Section 9; it never proceeds to produce a number for the registered coupled-oscillator or spin-bath comparison. A later implementation must not reuse the name `check_missing_energy_operator.py` while silently relaxing these eight checks; it must add the open-channel, direct-energy-accounting, and kill-criterion tests named in Section 9, adopt a versioned output schema, and replace `requirements.lock` with an exact, platform-tested dependency lock for whichever numerical-integration or master-equation stack the simulation requires.

References

- [1] E. Noether (1971). Invariant variation problems. *Transport Theory and Statistical Physics* 1(3), 186–207. <https://doi.org/10.1080/00411457108231446>.
- [2] C. Cirstoiu, K. Korzekwa, and D. Jennings (2020). Robustness of Noether’s principle: maximal disconnects between conservation laws and symmetries in quantum theory. *Physical Review X* 10, 041035. <https://doi.org/10.1103/PhysRevX.10.041035>.
- [3] W. F. Stinespring (1955). Positive functions on C^* -algebras. *Proceedings of the American Mathematical Society* 6(2), 211–216. <https://doi.org/10.1090/S0002-9939-1955-0069403-4>.
- [4] G. Lindblad (1976). On the generators of quantum dynamical semigroups. *Communications in Mathematical Physics* 48, 119–130. <https://doi.org/10.1007/BF01608499>.
- [5] V. Gorini, A. Kossakowski, and E. C. G. Sudarshan (1976). Completely positive dynamical semigroups of N -level systems. *Journal of Mathematical Physics* 17(5), 821–825. <https://doi.org/10.1063/1.522979>.
- [6] S. Vinjanampathy and J. Anders (2016). Quantum thermodynamics. *Contemporary Physics* 57(4), 545–579. <https://doi.org/10.1080/00107514.2016.1201896>.
- [7] A. Blais, A. L. Grimsmo, S. M. Girvin, and A. Wallraff (2021). Circuit quantum electrodynamics. *Reviews of Modern Physics* 93, 025005. <https://doi.org/10.1103/RevModPhys.93.025005>.
- [8] I. I. Rabi (1937). Space quantization in a gyrating magnetic field. *Physical Review* 51, 652–654. <https://doi.org/10.1103/PhysRev.51.652>.
- [9] F. Bloch (1946). Nuclear induction. *Physical Review* 70, 460–474. <https://doi.org/10.1103/PhysRev.70.460>.
- [10] E. T. Jaynes and F. W. Cummings (1963). Comparison of quantum and semiclassical radiation theories with application to the beam maser. *Proceedings of the IEEE* 51(1), 89–109. <https://doi.org/10.1109/PROC.1963.1664>.
- [11] A. O. Caldeira and A. J. Leggett (1983). Path integral approach to quantum Brownian motion. *Physica A: Statistical Mechanics and its Applications* 121(3), 587–616. [https://doi.org/10.1016/0378-4371\(83\)90013-4](https://doi.org/10.1016/0378-4371(83)90013-4).
- [12] N. V. Prokof’ev and P. C. E. Stamp (2000). Theory of the spin bath. *Reports on Progress in Physics* 63(4), 669–726. <https://doi.org/10.1088/0034-4885/63/4/204>.
- [13] R. M. Wald (1984). *General Relativity*. University of Chicago Press, Chicago.

- [14] A. Komar (1959). Covariant conservation laws in general relativity. *Physical Review* 113, 934–936. <https://doi.org/10.1103/PhysRev.113.934>.
- [15] R. Arnowitt, S. Deser, and C. W. Misner (2008). Republication of: the dynamics of general relativity. *General Relativity and Gravitation* 40, 1997–2027. <https://doi.org/10.1007/s10714-008-0661-1>.
- [16] E. Witten (1981). A new proof of the positive energy theorem. *Communications in Mathematical Physics* 80, 381–402. <https://doi.org/10.1007/BF01208277>.
- [17] H. Bondi, M. G. J. van der Burg, and A. W. K. Metzner (1962). Gravitational waves in general relativity, VII. Waves from axi-symmetric isolated systems. *Proceedings of the Royal Society A* 269, 21–52. <https://doi.org/10.1098/rspa.1962.0161>.
- [18] L. B. Szabados (2009). Quasi-local energy-momentum and angular momentum in general relativity. *Living Reviews in Relativity* 12, 4. <https://doi.org/10.12942/lrr-2009-4>.
- [19] M. J. Francis, L. A. Barnes, J. B. James, and G. F. Lewis (2007). Expanding space: the root of all evil? *Publications of the Astronomical Society of Australia* 24(2), 95–102. <https://doi.org/10.1071/AS07019>.
- [20] T. M. Davis and C. H. Lineweaver (2004). Expanding confusion: common misconceptions of cosmological horizons and the superluminal expansion of the universe. *Publications of the Astronomical Society of Australia* 21(1), 97–109. <https://doi.org/10.1071/AS03040>.