Faithfulness, Not Algebra Type, Controls the Rapid-Maintenance Singularity

Concrete compact-Dirac all-time control, faithful-family saturation, and exact no-go boundaries

Lluis Eriksson
Independent Researcher

Research draft — 5 August 2026

Abstract

Recent finite-dimensional maintenance bounds show that first-order leakage from the support of a rank-deficient target produces a universal free-energy loss $k_\mathrm{B}Tat\log(1/t)+O(t)$, and therefore logarithmically divergent exact-refresh power. Local algebras in algebraic quantum field theory (AQFT), however, are typically type III, while the vacuum and thermal reference states used in local physics are normally faithful. This creates an apparent mismatch between a rank-boundary theorem and its proposed field-theoretic interpretation.

We resolve the mismatch in ten steps. First, on every sigma-finite properly infinite factor we construct an explicit bounded quantum Markov semigroup, a nonfaithful normal target, and a faithful invariant reference state for which the Araki relative entropy is exactly binary. The resulting regulator-free local-algebra theorem gives

$$\frac{k_\mathrm{B}T}{t}\big[D(\varphi_0\|\eta)-D(\varphi_t\|\eta)\big]=k_\mathrm{B}T\kappa_+\log(1/t)+O(1).$$

Second, we derive this QMS from a relativistically localized microscopic model. A normalized compactly supported smear of the free charged Dirac field is coupled by an even exchange Hamiltonian to a stream of thermal fermionic probes. In the weak-collision scaling, the repeated reduced channels converge in norm to the QMS, with

$$\kappa_+=g^2\bar{n},\qquad \kappa_-=g^2(1-\bar{n}),\qquad \frac{\kappa_+}{\kappa_-}=e^{-\beta\Omega}.$$

The interaction and the limiting channel preserve gauge parity and fix every spacelike separated observable. Third, we remove the fresh-probe and divergent-collision idealization. The same local CAR mode is coupled, with a fixed smooth spatial form factor, to one free relativistic Dirac field in a thermal KMS state. For each coupling λ, the combined dynamics is a single autonomous bounded perturbation. In van Hove time $\tau=\lambda^2t$, its reduced interaction-picture maps converge in completely bounded norm to the QMS with on-shell rates

$$\Gamma_-=2\pi J_g(\Omega)[1-n_\beta(\Omega)],\qquad \Gamma_+=2\pi J_g(\Omega)n_\beta(\Omega),\qquad \frac{\Gamma_+}{\Gamma_-}=e^{-\beta\Omega}.$$

Thus the white-noise probe stream is not assumed: it emerges from a continuous finite-bandwidth bath in the Davies limit. Fourth, we quantify that limit at nonzero coupling. If the retarded and thermal kernels have finite first moments M_0,M_1, the exact quadratic CAR dynamics obeys, for every finite T,

$$\sup_{0\le\tau\le T}\|\widetilde{\mathcal{R}}_\lambda(\tau)-e^{\tau\mathcal{L}_D}\|_\mathrm{cb}\le 16\lambda^2M_1(1+M_0T)^2.$$

The induced finite-coupling correction to the maintenance ledger is at most $k_{\mathrm{B}}T\tau^{-1}[h_2(\epsilon/2)+\epsilon|\log q_*|]$, with $\epsilon = 16\lambda^2 M_1(1+M_0T)^2$. Hence the logarithmic law survives the joint limit whenever $\lambda^2\log(1/|\lambda|) = o(\tau\log(1/\tau))$. Fifth, we discharge the all-time assumptions for one displayed compact Dirac datum. For the massless form factor $g_R = (-\Delta)^2(b_R\xi)$, with an explicit bump b_R supported in a ball, Paley–Wiener estimates, threshold order $J_{g_R}(\omega) = O(\omega^{10})$, and $J_{g_R}(\Omega) > 0$ verify the fermionic resonance hypotheses and give

$$\sup_{t\geq 0}\|\mathcal{R}_\lambda(t) - e^{t(\mathcal{L}_S+\lambda^2\mathcal{L}_D)}\|_{\mathrm{cb}} \leq C|\lambda|.$$

Under additional reduced-residue, remainder, and stationary-matching hypotheses, this improves to

$$\sup_{\tau\geq 0}\|\widetilde{\mathcal{R}}_\lambda(\tau) - e^{\tau\mathcal{L}_D}\|_{\mathrm{cb}} \leq C\lambda^2.$$

An exact-pole resonance approximation is asymptotically exact; complete positivity requires, and is stated only under, a separate GKLS-embedding hypothesis. We also prove the complementary Zeno obstruction: at fixed coupling every bounded autonomous dilation leaks across a system support only at order t^2, so the Davies rapid singularity is necessarily kinetic rather than a literal fixed-coupling $t\downarrow 0$ law. Sixth, we construct a background-covariant two-Dirac-field completion on every globally hyperbolic spin spacetime. A compactly supported, gauge-neutral off-diagonal bundle coupling gives a Green-hyperbolic operator, a natural CAR scattering morphism defined by relative Cauchy evolution, causal factorization, and exact spacelike triviality. We also prove a locality obstruction: every nonzero local multiplier has infinite one-particle rank and therefore cannot equal the rank-one exchange used by the solvable reservoir model. A stationary Feshbach reduction recovers the same Davies generator, while quantifying the unavoidable finite-coupling correction from the complementary field modes. Seventh, for a finite-dimensional target softened by weight ε on its kernel, we prove an exact entropy-slope identity and the cutoff law

$$-\frac{d}{dt}F(e^{t\mathcal{L}}\rho_\varepsilon)\bigg|_{t=0} = k_{\mathrm{B}}T\,a\log(1/\varepsilon) + O(1).$$

Thus soft support does not remove the obstruction: it replaces the temporal logarithm by a state-space logarithm. Eighth, we prove the sharp faithful-state no-go and its optimal replacement. No fixed faithful vacuum/KMS restriction can have outward support leakage, but the exact type-III model admits faithful states φ_{t^α} for every $t>0$ with

$$\mathcal{P}_{t^\alpha}(t) = k_{\mathrm{B}}T\kappa_+\min\{\alpha,1\}\log(1/t) + O(1).$$

Thus faithful floors closing at least linearly recover the full coefficient without ever using a nonfaithful state at nonzero regulator. Ninth, we prove a uniform joint-limit theorem for regulator families. The diagonal rank-boundary term is governed neither by leakage a_n nor refresh time t_n separately, but by $a_n\log(1/t_n)$. Tenth, we formulate the honest split-regulated interface for targets not given intrinsically on the local algebra. Type III permits but does not force the singularity: fixed faithful states obey the no-go, while conditioned targets and faithful vanishing-floor families give the two precise operational corridors.

Keywords. algebraic quantum field theory; split property; type III factors; quantum Markov semigroups; support leakage; entropy perturbation; nonequilibrium free energy; noncommuting limits

1 The problem and the result map

Let ρ be a finite-dimensional target state, $P = \operatorname{supp}\rho$, $Q = \mathbf{1}-P$, and let $T_t = e^{t\mathcal{L}}$ be an uncontrolled positive trace-preserving semigroup. The outward support-leakage coefficient is

$$a_{\mathcal{L}}(\rho) = \mathrm{Tr}[Q\mathcal{L}(\rho)]\,. \tag{1}$$

For $a_{\mathcal{L}}(\rho) > 0$, the support-extending entropy perturbation gives [1, 29]

$$F(\rho) - F(T_t\rho) = k_{\mathrm{B}}T\, a_{\mathcal{L}}(\rho) t \log(1/t) + O(t), \qquad t \downarrow 0, \tag{2}$$

where $F(\sigma) = \mathrm{Tr}(H\sigma) - k_{\mathrm{B}}TS(\sigma)$. Under the fresh-target-cell ledger of Ref. [29], division by t converts (2) into divergent rapid exact-refresh power.

The AQFT problem is not a routine replacement of matrices by operators. In a local type-III algebra there is no canonical trace or density matrix, and hence no unregularized local von Neumann entropy. A split inclusion supplies an intermediate type-I factor, but the factor, a finite corner inside it, a target support, and the induced dynamics are additional data. Moreover, the physically standard local vacuum and KMS states are normally faithful under the usual hypotheses. Their support is the identity, so the particular coefficient in (1) has no complement on which to leak.

The paper separates four statements that are often conflated:

1. *Algebraic leakage:* population can leave the support of a nonfaithful normal state on any von Neumann algebra, including a type-III factor.

2. *Entropic singularity:* converting leakage into $t\log(1/t)$ requires an entropy or relative-free-energy functional and its domain.

3. *Split regularization:* a type-I implementation permits density-matrix bookkeeping, but its support boundary may depend on the split and cutoff.

4. *Continuum survival:* a singularity is intrinsic only if its coefficient is stable under a declared regulator limit.

The algebraic AQFT result is the exact regulator-free construction in Theorem 3.1. The first microscopic result is Theorem 4.1: a compactly localized free-Dirac CAR mode and a thermal fermionic probe stream generate the same QMS as a norm limit of unitary collisions. The stronger autonomous result is Theorem 5.1: one continuous relativistic Dirac KMS reservoir with a fixed smooth form factor generates the QMS in the weak-coupling limit, without probe resets or a collision Hamiltonian diverging as $\delta^{-1/2}$. The quantitative upgrade is Theorem 6.2: first moments of the bath kernels give a closed $O(\lambda^2)$ cb error on every van Hove window. Definition 6.4, Lemma 6.5, and Theorem 6.6 give a fully specified compact Dirac datum and an unconditional $O(|\lambda|)$ approximation on the entire half-line. Hypothesis 6.8 and Theorem 6.9 isolate the extra reduced data needed for the sharper $O(\lambda^2)$ order, while Proposition 6.10 proves why neither absolute approximation turns the exact fixed-coupling Zeno corner into first-order leakage. The covariant completion is Theorem 7.1; Proposition 7.2 and Theorem 7.5 give the local-rank obstruction and stationary Feshbach bridge. Proposition 3.3 rules out the support mechanism at a fixed faithful state, while Theorem 8.5 recovers its exact coefficient along faithful vanishing-floor families. The finite-dimensional soft-boundary identity is Theorem 8.1; the regulator scaling result is Theorem 9.2. The separate split statement, Theorem 10.2, remains deliberately conditional because a cutoff code is not fixed by the existence theorem.

2 Support leakage is an algebraic notion

We first remove traces and density matrices from the definition of leakage.

Proposition 2.1 (Intrinsic support-leakage derivative)**.** *Let $\mathcal{M}$ be a von Neumann algebra, $(\Phi_t)_{t\geq 0}$ a normal unital positive semigroup on $\mathcal{M}$, and φ a normal state. Write $p = s(\varphi)$ for its support projection and $q = \mathbf{1} - p$. If the right derivative exists, define*

$$a_\Phi(\varphi) = \lim_{t\downarrow 0} \frac{(\Phi_{t,*}\varphi)(q)}{t}, \tag{3}$$

where $\Phi_{t,*}\varphi = \varphi \circ \Phi_t$. Then $a_\Phi(\varphi) \geq 0$. If φ is faithful, $p = \mathbf{1}$ and $a_\Phi(\varphi) = 0$ identically.

Proof. Normal positivity gives $(\Phi_{t,*}\varphi)(q) \geq 0$ for $t \geq 0$, while the definition of the support projection gives $\varphi(q) = 0$. A right derivative at a minimum is nonnegative. For a faithful state, $q = 0$. □

Remark 2.2 (Type III is not the deciding property)*.* A type-III factor has many proper projections and many nonfaithful normal states. Proposition 2.1 therefore permits positive intrinsic leakage on a type-III factor. Conversely, a faithful state on a matrix algebra also has $p = \mathbf{1}$ and zero support-leakage coefficient. The relevant dichotomy is faithfulness versus a proper support boundary, not type I versus type III.

Remark 2.3 (What the proposition does not prove)*.* Equation (3) is representation independent, but it is only a mass derivative. It does not by itself define a local von Neumann entropy on a type-III algebra and does not prove a $t \log t$ term for Araki relative entropy. Such a claim requires support and modular-domain hypotheses for the chosen pair of normal functionals [11, 26].

The next section supplies those ingredients in an exact, regulator-free model. It is an existence theorem rather than a claim about every local state or every physical bath.

In finite dimension, if the predual generator exists in norm, (3) reduces to (1). If $\mathcal{L}$ is GKLS,

$$\mathcal{L}(\rho) = -i[K, \rho] + \sum_j \gamma_j \left(L_j \rho L_j^\dagger - \tfrac{1}{2}\{L_j^\dagger L_j, \rho\} \right), \tag{4}$$

then

$$a_{\mathcal{L}}(\rho) = \sum_j \gamma_j \, \|Q L_j \sqrt{\rho}\|_2^2 \,. \tag{5}$$

The Hamiltonian and anticommutator terms have zero Q-trace. Equation (5) is the positive no-cancellation formula used in the preceding maintenance papers [29, 30].

3 An exact regulator-free theorem on a type-III local algebra

We now show that the missing modular statement can be proved exactly for a nontrivial local model. No trace, density matrix, split factor, finite corner, or limiting argument is used.

Theorem 3.1 (Exact type-III relative-entropy maintenance singularity)**.** *Let $\mathcal{M}$ be a sigma-finite properly infinite von Neumann factor. Fix rates $\kappa_+, \kappa_- > 0$, put $\lambda = \kappa_+ + \kappa_-$, and define*

$$r = \frac{\kappa_-}{\lambda}, \qquad 1 - r = \frac{\kappa_+}{\lambda}. \tag{6}$$

Then there exist orthogonal normal states φ_0, φ_1, a faithful normal state $\eta_r = r\varphi_0 + (1-r)\varphi_1$, and a bounded normal quantum Markov semigroup $(\Phi_t)_{t\geq 0}$ such that:

(i) $\eta_r \circ \Phi_t = \eta_r$ for every $t \geq 0$;

(ii) starting from φ_0, the predual orbit is exactly

$$\varphi_t = \Phi_{t,*}\varphi_0 = (1-p_t)\varphi_0 + p_t\varphi_1, \qquad p_t = (1-r)(1-e^{-\lambda t}); \tag{7}$$

(iii) the intrinsic leakage coefficient is $a_\Phi(\varphi_0) = \kappa_+$;

(iv) for Araki relative entropy,

$$D(\varphi_t \| \eta_r) = (1-p_t)\log\frac{1-p_t}{r} + p_t \log\frac{p_t}{1-r}; \tag{8}$$

(v) the relative-free-energy restoration ledger obeys

$$\begin{aligned}\Pi_{\mathrm{rel}}(t) &:= \frac{k_{\mathrm{B}}T}{t}\left[D(\varphi_0\|\eta_r) - D(\varphi_t\|\eta_r)\right] \\ &= \frac{k_{\mathrm{B}}T}{t}\left[h_2(p_t) + p_t\log\frac{1-r}{r}\right] = k_{\mathrm{B}}T\kappa_+\log(1/t) + O(1).\end{aligned} \tag{9}$$

(vi) the instantaneous entropy-production rate is also exact:

$$-\frac{d}{dt}D(\varphi_t\|\eta_r) = p_t'\log\frac{(1-p_t)(1-r)}{p_t r} = \kappa_+\log(1/t) + O(1). \tag{10}$$

If $\mathcal{M} = \mathcal{A}(O)$ is a properly infinite local algebra in a Haag–Kastler representation, the semigroup can be implemented with coefficients in $\mathcal{A}(O)$. Its extension by the same bounded generator to the ambient algebra fixes $\mathcal{A}(O)'$, and hence every spacelike separated local algebra, pointwise.

Proof. Proper infiniteness supplies isometries $s_0, s_1 \in \mathcal{M}$ satisfying the Cuntz relations

$$s_i^* s_j = \delta_{ij}\mathbf{1}, \qquad s_0 s_0^* + s_1 s_1^* = \mathbf{1}.$$

Put $P = s_0 s_0^*$, $Q = s_1 s_1^*$, and $V = s_1 s_0^*$. Then

$$P + Q = \mathbf{1}, \qquad V^*V = P, \qquad VV^* = Q, \qquad V^2 = 0.$$

Sigma-finiteness gives a faithful normal state ϕ on the corner $P\mathcal{M}P$. Define

$$\varphi_0(x) = \phi(PxP), \qquad \varphi_1(x) = \phi(V^*xV). \tag{11}$$

Their support projections are P and Q, respectively.

Consider the bounded Heisenberg generator

$$\mathcal{L}(x) = \kappa_+\left(V^*xV - \tfrac{1}{2}\{P, x\}\right) + \kappa_-\left(VxV^* - \tfrac{1}{2}\{Q, x\}\right). \tag{12}$$

This is Lindblad form with noise operators $\sqrt{\kappa_+}V$ and $\sqrt{\kappa_-}V^*$, so it generates a norm-continuous normal unital completely positive semigroup. Direct substitution, using $V^2 = 0$, gives

$$\mathcal{L}_*\varphi_0 = \kappa_+(\varphi_1 - \varphi_0), \qquad \mathcal{L}_*\varphi_1 = \kappa_-(\varphi_0 - \varphi_1). \tag{13}$$

Thus the two-state span is invariant. Solving the classical rate equation proves (7); Equation (6) and (13) prove invariance of η_r. Since $\varphi_t(Q) = p_t = \kappa_+ t + O(t^2)$, item (iii) follows.

The state η_r is faithful. Indeed, if $x \geq 0$ and $\eta_r(x) = 0$, faithfulness of ϕ gives $PxP = 0$ and $V^*xV = 0$, hence also $QxQ = 0$. Positivity then kills the off-diagonal blocks, so $x = 0$.

It remains to prove the modular identity (8). Let $\theta_p = (1-p, p)$ and $\theta_r = (r, 1-r)$ be states on $\mathbb{C}^2$. The maps

$$E : \mathcal{M} \longrightarrow \mathbb{C}^2, \qquad E(x) = (\varphi_0(x), \varphi_1(x)),$$
$$\iota : \mathbb{C}^2 \longrightarrow \mathcal{M}, \qquad \iota(a, b) = aP + bQ$$

are normal unital completely positive, with ι a $*$-homomorphism. If $\omega_p = (1-p)\varphi_0 + p\varphi_1$, then

$$\theta_p \circ E = \omega_p, \quad \omega_p \circ \iota = \theta_p, \quad \theta_r \circ E = \eta_r, \quad \eta_r \circ \iota = \theta_r.$$

Data processing for Araki relative entropy applied once to each map yields both inequalities

$$D(\omega_p \| \eta_r) \le D(\theta_p \| \theta_r) \quad \text{and} \quad D(\theta_p \| \theta_r) \le D(\omega_p \| \eta_r).$$

Equality follows, and the classical binary expression is (8) [11, 12].

Finally, $D(\varphi_0 \| \eta_r) = -\log r$, so elementary algebra gives the first equality in (9). Since $p_t = \kappa_+ t + O(t^2)$ and $h_2(p) = p\log(1/p) + p + O(p^2)$, the asymptotic formula follows. Differentiating (8) proves (10). If $V \in \mathcal{A}(O)$, every $x \in \mathcal{A}(O)'$ commutes with V, P, Q, and (12) gives $\mathcal{L}(x) = 0$. The resulting operation is local in the strong, commutant-fixing sense used for local CP operations in AQFT [13]. □

Corollary 3.2 (Sharp faithfulness dichotomy)**.** *On every sigma-finite properly infinite factor, every prescribed leakage value $a > 0$ is realized by Theorem 3.1 after setting $\kappa_+ = a$; the reference bias $r \in (0, 1)$ can be prescribed independently by choosing $\kappa_- = ar/(1-r)$. By contrast, Proposition 2.1 forces $a_\Phi(\varphi) = 0$ for every faithful normal target and every positive semigroup. Thus proper infiniteness guarantees availability of the singular mechanism, while target faithfulness decides whether the support-leakage coefficient can be nonzero.*

Proposition 3.3 (Sharp faithful-state no-go)**.** *Let φ be a fixed faithful normal state on a von Neumann algebra and let $(\Phi_t)_{t\ge0}$ be any normal unital positive semigroup. Then the outward support coefficient of φ is identically zero. In finite dimension, if the generator is bounded, the free-energy derivative at $t = 0$ is finite. Therefore no fixed faithful vacuum or KMS restriction can exhibit the rank-boundary $t\log(1/t)$ mechanism.*

Proof. Faithfulness is equivalent to $s(\varphi) = \mathbf{1}$, so the complementary projection in (3) is zero. In finite dimension, $\log\rho$ is bounded for every fixed faithful density matrix ρ, and

$$\left.\frac{d}{dt} S(e^{t\mathcal{L}}\rho)\right|_{t=0} = -\operatorname{Tr}[\mathcal{L}(\rho)\log\rho]$$

is finite. A type-III modular entropy derivative may fail for separate domain reasons, but such a failure is not outward support leakage. This proves both the no-go and the distinction of mechanisms. □

Remark 3.4 (Why the invariant reference matters). The reference η_r is faithful and invariant, so monotonicity of $D(\,\cdot\|\eta_r)$ under $\Phi_{t,*}$ is automatic. Equation (9) prices restoration by the increase of the relative free-energy resource $k_{\mathrm{B}}TD(\,\cdot\|\eta_r)$. As in any maintenance ledger, interpreting this resource difference as minimum laboratory work additionally requires a specified reversible implementation; the theorem itself is an exact statement about the declared resource functional.

Remark 3.5 (Existence, not universality). The construction is intrinsic to the local factor and removes the split/corner regulator, but it deliberately chooses a nonfaithful target and an engineered bounded local dissipator. It does not say that the vacuum, a KMS state, or a semigroup derived from a particular field–bath Hamiltonian has this singularity. Rather, it proves that type III presents no modular or causal obstruction to the mechanism once a proper support boundary and a compatible local operation are specified.

4 Microscopic derivation from a relativistic fermionic probe stream

Theorem 3.1 used an algebraically selected partial isometry. We now realize the same two jumps with a compactly localized free Dirac field and derive their rates from unitary collisions with thermal probes. The construction combines three standard ingredients: bounded local CAR fields [14, 19], continuous limits of repeated interactions [20, 21], and causally localized system–probe couplings in AQFT [22].

Let O be a diamond in a four-dimensional globally hyperbolic spin spacetime. Consider the free charged Dirac field in a pure gauge-invariant quasi-free Hadamard representation. Write $\mathcal{F}(O)$ for the local field algebra and $\mathcal{A}(O) = \mathcal{F}(O)^0$ for its even, gauge-invariant observable algebra. Both local von Neumann algebras are type III in the standard representations, while the field net is graded local and the observable net is local [19].

Choose a nonzero spinor test-function class $[f]$ with compact support in O, normalized in the positive CAR form. The smeared charged field

$$a = \psi(f) \in \mathcal{F}(O) \tag{14}$$

is bounded and odd, and obeys

$$a^2 = 0, \qquad aa^* + a^*a = \mathbf{1}. \tag{15}$$

Thus $P = aa^*$ and $Q = a^*a$ are complementary projections in the observable algebra $\mathcal{A}(O)$, and the charged field a^* intertwines their corners.

Each fresh probe is a single fermionic CAR mode b, with Hamiltonian $H_B = \Omega b^*b$ and even Gibbs state

$$\tau_\beta(b^*b) = \bar{n} = \frac{1}{1+e^{\beta\Omega}}, \qquad \tau_\beta(b) = 0. \tag{16}$$

The system and probe are combined with the graded tensor product. Their gauge-neutral, parity-even exchange Hamiltonian is

$$h = a^*b + b^*a = h^*. \tag{17}$$

For a collision of duration δ, use the weak-collision scaling

$$U_\delta = e^{-ig\sqrt{\delta}\,h}, \qquad H_{\mathrm{int},\delta} = \frac{g}{\sqrt{\delta}}h, \tag{18}$$

and discard the probe afterwards. In a standard Klein representation of the graded product, the reduced Heisenberg channel on $\mathcal{A}(O)$ is

$$\mathcal{T}_\delta(x) = (\mathrm{id}\,\widehat{\otimes}\tau_\beta)\left[U_\delta^*(x\widehat{\otimes}\mathbf{1})U_\delta\right]. \tag{19}$$

Theorem 4.1 (Dirac–CAR collision-limit realization)**.** *The channels $\mathcal{T}_\delta$ are normal, unital, completely positive, and gauge preserving on $\mathcal{A}(O)$. In completely bounded norm,*

$$\mathcal{T}_\delta = \mathrm{id} + \delta\mathcal{L}_{\beta,g} + O(\delta^2), \qquad \delta \downarrow 0, \tag{20}$$

where

$$\mathcal{L}_{\beta,g}(x) = \kappa_+\left(axa^* - \tfrac{1}{2}\{P,x\}\right) + \kappa_-\left(a^*xa - \tfrac{1}{2}\{Q,x\}\right), \tag{21}$$

$$\kappa_+ = g^2\bar{n}, \qquad \kappa_- = g^2(1-\bar{n}). \tag{22}$$

For every finite T,

$$\sup_{0\le t\le T} \left\| \mathcal{T}_{t/n}^{n} - e^{t\mathcal{L}_{\beta,g}} \right\|_{\mathrm{cb}} = O_T(n^{-1}). \tag{23}$$

Moreover,

$$\frac{\kappa_+}{\kappa_-} = e^{-\beta\Omega}. \tag{24}$$

Every collision channel and the limiting QMS fix pointwise each even observable localized spacelike to O.

Proof. The even probe state makes the graded partial expectation in (19) a normal UCP map; in a Klein-transformed Hilbert-space representation it is the usual partial expectation. Gauge neutrality of h and gauge invariance of τ_β imply preservation of $\mathcal{A}(O)$.

Put $\theta = g\sqrt{\delta}$ and $X = x\widehat{\otimes}\mathbf{1}$. The bounded commutator expansion gives

$$U_\delta^* X U_\delta = X + i\theta[h, X] - \frac{\theta^2}{2}[h,[h,X]] + O(\theta^3).$$

Every odd power contains odd probe parity and vanishes after τ_β. Hence the actual remainder after the quadratic term is $O(\theta^4) = O(\delta^2)$, uniformly in cb norm on the unit ball. Using $b^2 = 0$, $bb^* + b^*b = \mathbf{1}$, and that the even x graded-commutes with probe operators,

$$(\mathrm{id}\,\widehat{\otimes}\tau_\beta)(hXh) = (1-\bar{n})a^*xa + \bar{n}axa^*,$$
$$(\mathrm{id}\,\widehat{\otimes}\tau_\beta)(h^2) = (1-\bar{n})Q + \bar{n}P.$$

Since $-\frac{1}{2}[h,[h,X]] = hXh - \frac{1}{2}\{h^2, X\}$, Equations (21)–(22) follow, including the absence of an extra factor of two.

The cb-norm remainder in (20), contractivity of the channels, and the boundedness of $\mathcal{L}_{\beta,g}$ give the Chernoff product estimate (23); equivalently, a telescoping comparison with $(\mathrm{id} + (t/n)\mathcal{L}_{\beta,g})^n$ gives an $O_T(n^{-1})$ bound. Equation (24) is immediate from (16).

Finally, let B be an even observable localized spacelike to O. Graded locality gives $[B, a] = [B, a^*] = 0$, hence $[B, h] = 0$. Therefore every U_δ, every $\mathcal{T}_\delta$, and their norm limit fix B pointwise. This is the causal-factorization property at the level of the nonselective channel. □

Corollary 4.2 (Microscopic relative-entropy maintenance law)**.** *Let ϕ be a faithful normal state on $P\mathcal{A}(O)P$, and define*

$$\varphi_0(x) = \phi(PxP), \qquad \varphi_1(x) = \phi(axa^*), \qquad \eta_\beta = (1-\bar{n})\varphi_0 + \bar{n}\varphi_1.$$

Then η_β is faithful and invariant under $e^{t\mathcal{L}_{\beta,g}}$. Starting from φ_0,

$$\varphi_t = (1-p_t)\varphi_0 + p_t\varphi_1, \qquad p_t = \bar{n}(1 - e^{-g^2 t}). \tag{25}$$

The Araki relative entropy is the binary divergence in (8) with $r = 1 - \bar{n}$, and

$$\frac{k_{\mathrm{B}}T}{t}\left[D(\varphi_0\|\eta_\beta) - D(\varphi_t\|\eta_\beta)\right] = k_{\mathrm{B}}Tg^2\bar{n}\log(1/t) + O(1). \tag{26}$$

Thus the coefficient of the maintenance singularity is the thermally activated probe absorption rate.

Proof. The charged partial isometry a^* implements mutually inverse $*$-isomorphisms between the P- and Q-corners of the even algebra, so the state construction and two-channel data-processing proof of Theorem 3.1 apply verbatim. Equations (22) give $\lambda = \kappa_+ + \kappa_- = g^2$, $r = \kappa_-/\lambda = 1 - \bar{n}$, and leakage $a_\Phi = \kappa_+ = g^2\bar{n}$. □

Remark 4.3 (Relativistic collision protocol). The single-step coupling is localized in O. A repeated protocol is obtained by placing causally ordered copies O_j in a timelike laboratory worldtube, using one fresh probe mode in each copy, and identifying the system algebras in the interaction picture by free Dirac evolution. Each finite protocol is a composition of localized unitary couplings and satisfies causal factorization, in the same operational spirit as the Fewster–Verch system–probe framework [22].

Remark 4.4 (What remains idealized). The derivation is microscopic and Hamiltonian at every finite δ, but the continuous Markov limit is a singular weak-collision limit: $H_{\mathrm{int},\delta}$ scales as $\delta^{-1/2}$, probes are freshly reset, and the interaction-picture identification is imposed. The next section removes both the resets and the divergent collision Hamiltonian by replacing the stream with one continuous finite-bandwidth Dirac reservoir.

5 Autonomous derivation from one relativistic Dirac reservoir

We now keep a single bath for the entire experiment. The only limiting operation will be the standard weak-coupling/long-time limit; there is no renewal of bath degrees of freedom and no interaction whose norm diverges as a time step is shortened. The construction uses the Davies weak-coupling theorem [4, 5], quasi-free CAR KMS states [18], and the boundedness of smeared Dirac fields.

Work in a stationary Minkowski laboratory. Retain the local system field a from (14) and set

$$\mathcal{N}_a = \{a, a^*\}'', \qquad \mathcal{R}_a = \mathcal{N}_a' \cap \mathcal{F}(O).$$

The CAR relations (15) make $\mathcal{N}_a$ a unital type-I_2 factor. Hence

$$\mathcal{F}(O) \simeq \mathcal{N}_a \,\overline{\otimes}\, \mathcal{R}_a, \qquad \mathcal{N}_a \simeq M_2(\mathbb{C}), \tag{27}$$

and the second tensor factor is a spectator for every dynamics below. Give the selected mode the bounded local Hamiltonian

$$H_S = \Omega Q, \qquad \alpha_t^S(a) = e^{-i\Omega t} a. \tag{28}$$

This is an engineered stationary laboratory mode, not the unrestricted free-Dirac time evolution of an exactly compact wave packet.

Let $\mathcal{F}_R$ be a second free charged Dirac field of mass $m_R \geq 0$, with relativistic one-particle dispersion $E_{\boldsymbol{p}} = \sqrt{|\boldsymbol{p}|^2 + m_R^2}$, free dynamics β_t, and gauge-invariant quasi-free β-KMS state ω_β. Choose smooth compactly supported spinor Cauchy data g in a laboratory ball and put

$$B = \chi(g) \in \mathcal{F}_R, \qquad V = a^* \widehat{\otimes} B + B^* \widehat{\otimes} a = V^*. \tag{29}$$

The field operators are bounded, so V is a bounded, parity-even, gauge-neutral interaction with a fixed smooth spatial form factor. For each $\lambda \in \mathbb{R}$, the derivation

$$\delta_\lambda = \delta_S + \delta_R + i\lambda[V, \cdot\,] \tag{30}$$

generates one autonomous system–reservoir automorphism group α_t^λ. In a Hamiltonian representation this is $H_\lambda = \Omega Q + H_R + \lambda V$. The same reservoir remains coupled at all times.

Define the thermal correlation functions

$$C_>(s) = \omega_\beta(\beta_s(B)B^*)\,, \qquad C_<(s) = \omega_\beta(\beta_s(B^*)B)\,. \tag{31}$$

For compactly supported smooth g in three spatial dimensions these functions are integrable: the on-shell oscillatory integrals decay as $O(|s|^{-3/2})$ for $m_R > 0$, and as $O(|s|^{-3})$ for $m_R = 0$. Equivalently, one may take $C_>, C_< \in L^1(\mathbb{R})$ as the explicit mixing hypothesis. Put

$$\Gamma_- = \int_{\mathbb{R}} e^{i\Omega s} C_>(s)\, ds, \qquad \Gamma_+ = \int_{\mathbb{R}} e^{-i\Omega s} C_<(s)\, ds. \tag{32}$$

If $\Omega > m_R$, the free-Dirac spectral representation gives

$$\Gamma_- = 2\pi J_g(\Omega)\, [1 - n_\beta(\Omega)], \qquad \Gamma_+ = 2\pi J_g(\Omega)\, n_\beta(\Omega), \tag{33}$$

$$n_\beta(\omega) = (1 + e^{\beta\omega})^{-1}, \qquad J_g(\omega) = \int \frac{d^3\boldsymbol{p}}{(2\pi)^3} \langle \widehat{g}(\boldsymbol{p}), \Pi_+(\boldsymbol{p}) \widehat{g}(\boldsymbol{p}) \rangle \delta(\omega - E_{\boldsymbol{p}}) \geq 0, \tag{34}$$

up to the conventional normalization absorbed into g. We impose the Fermi-golden-rule condition $J_g(\Omega) > 0$.

For $\tau \geq 0$, define the reduced interaction-picture map

$$\widetilde{\mathcal{R}}_\lambda(\tau) = \alpha^S_{-\tau/\lambda^2} \circ (\mathrm{id}\, \widehat{\otimes} \omega_\beta) \circ \alpha^\lambda_{\tau/\lambda^2}(\,\cdot\, \widehat{\otimes} \mathbf{1}). \tag{35}$$

Theorem 5.1 (Autonomous relativistic-bath realization)**.** *Under the mixing and Fermi-golden-rule hypotheses above, there is a real Lamb shift Δ such that, for every finite T,*

$$\lim_{\lambda \to 0} \sup_{0 \leq \tau \leq T} \left\| \widetilde{\mathcal{R}}_\lambda(\tau) - e^{\tau \mathcal{L}_D} \right\|_{\mathrm{cb}} = 0 \tag{36}$$

on $\mathcal{F}(O)$, and therefore on the gauge-invariant observable algebra $\mathcal{A}(O)$. The limiting Heisenberg generator is

$$\mathcal{L}_D(x) = i[\Delta Q, x] + \Gamma_+ \left(a x a^* - \tfrac{1}{2}\{P, x\} \right) + \Gamma_- \left(a^* x a - \tfrac{1}{2}\{Q, x\} \right). \tag{37}$$

It preserves gauge invariance, fixes $\mathcal{R}_a$ pointwise, and hence fixes every even system observable spacelike separated from O. The KMS condition gives

$$\frac{\Gamma_+}{\Gamma_-} = e^{-\beta\Omega}. \tag{38}$$

Proof. The matrix units P, Q, a, a^* generated by the CAR mode prove (27). Both H_S and V act trivially on $\mathcal{R}_a$, so every reduced map factors as a channel on $M_2(\mathbb{C})$ tensored with the identity on $\mathcal{R}_a$. The finite-system Davies theorem applies because (31) is integrable and the only nonzero Bohr frequencies of the coupling are $\pm\Omega$. Its norm convergence on M_2 therefore extends, with the same norm, to completely bounded convergence on $M_2 \overline{\otimes} \mathcal{R}_a$, proving (36).

The secular terms at $\pm\Omega$ give (37); the principal-value parts combine into ΔQ. Positivity of the Fourier transforms gives $\Gamma_\pm \geq 0$. The quasi-free Dirac two-point function gives (33)–(34); alternatively, the KMS boundary relation applied to B, B^* directly gives $\Gamma_+ = e^{-\beta\Omega}\Gamma_-$. Because V is even and gauge neutral, the finite-λ and limiting maps preserve the observable algebra. Finally, every even system observable spacelike to O commutes with a, a^*, P, Q, so it is fixed already by the full autonomous dynamics and by its Davies limit. □

Corollary 5.2 (Autonomous-bath maintenance law)**.** *Let $\Gamma = \Gamma_+ + \Gamma_-$, and define φ_0, φ_1 as in Corollary 4.2. Then*

$$\eta_D = \frac{\Gamma_-}{\Gamma} \varphi_0 + \frac{\Gamma_+}{\Gamma} \varphi_1 \tag{39}$$

is faithful and invariant. The Lamb shift does not alter the binary sector, and the predual orbit from φ_0 *is*

$$\varphi_\tau = (1 - p_\tau)\varphi_0 + p_\tau\varphi_1, \qquad p_\tau = \frac{\Gamma_+}{\Gamma}(1 - e^{-\Gamma\tau}). \tag{40}$$

Consequently,

$$\frac{k_{\mathrm{B}}T}{\tau}\left[D(\varphi_0\|\eta_D) - D(\varphi_\tau\|\eta_D)\right] = k_{\mathrm{B}}T\Gamma_+\log(1/\tau) + O(1). \tag{41}$$

The singular coefficient is the on-shell thermal absorption rate $2\pi J_g(\Omega)n_\beta(\Omega)$ *of one continuous relativistic bath.*

Remark 5.3 (What has and has not been removed). For every fixed λ, (30) is one time-homogeneous autonomous dynamics with a bounded interaction and a smooth spatial profile; no bath subsystem is discarded or reset. The profile g is fixed as $\lambda \to 0$, and no coupling grows like $\delta^{-1/2}$. The limiting QMS is nevertheless a weak-coupling statement on the long physical time $t = \tau/\lambda^2$; at finite λ the reduced dynamics retains memory and need not be a semigroup. The model also selects a stationary laboratory frame and the engineered local Hamiltonian ΩQ. Thus it is an autonomous relativistic reservoir derivation, not a claim of exact finite-coupling Markovianity or full general covariance.

6 An explicit finite-coupling error bound

The preceding Davies theorem is qualitative. The exchange model is quadratic, however, so its memory can be eliminated without a Born truncation. This gives a nonasymptotic error estimate at each nonzero λ. Let h_R be the positive-energy one-particle Dirac Hamiltonian and use the one-particle vector g that defines B. For $s \geq 0$, set

$$K(s) = e^{i\Omega s}\langle g, e^{-ih_R s}g\rangle, \qquad K_+(s) = e^{-i\Omega s}C_<(s), \tag{42}$$

and assume the explicit first-moment condition

$$M_j := \int_0^\infty s^j\left(|K(s)| + |K_+(s)|\right) ds < \infty, \qquad j = 0, 1. \tag{43}$$

This strengthens the L^1 hypothesis used in Theorem 5.1. It holds for the massless four-dimensional model above. In the massive model it can be enforced while retaining compactly supported smooth Cauchy data and $J_g(\Omega) > 0$ by imposing a zero of sufficient order on the form factor at the stationary threshold $\boldsymbol{p} = 0$. Put

$$\kappa = \int_0^\infty K(s)\, ds, \qquad \kappa_+ = \int_0^\infty K_+(s)\, ds. \tag{44}$$

Then $2\,\mathrm{Re}\,\kappa = \Gamma_+ + \Gamma_- =: \Gamma$, $2\,\mathrm{Re}\,\kappa_+ = \Gamma_+$, and the imaginary part of κ is the Lamb-shift convention in (37).

Lemma 6.1 (Exact scalar memory equations)**.** *The reduced interaction-picture channel on* $\mathcal{N}_a \simeq M_2(\mathbb{C})$ *has the form*

$$\widetilde{\mathcal{R}}_\lambda(\tau)(a) = v_\lambda(\tau)a, \qquad \widetilde{\mathcal{R}}_\lambda(\tau)(Q) = |v_\lambda(\tau)|^2 Q + z_\lambda(\tau)\mathbf{1}, \tag{45}$$

where $|v_\lambda| \leq 1$, $z_\lambda(0) = 0$, *and*

$$\dot{v}_\lambda(\tau) = -\int_0^{\tau/\lambda^2} K(s)v_\lambda(\tau - \lambda^2 s)\, ds, \tag{46}$$

$$\dot{z}_\lambda(\tau) = 2\,\mathrm{Re}\left[\overline{v_\lambda(\tau)}\int_0^{\tau/\lambda^2} K_+(s)v_\lambda(\tau - \lambda^2 s)\, ds\right]. \tag{47}$$

The Davies channel has the same form with

$$v_D(\tau) = e^{-\kappa\tau}, \qquad z_D(\tau) = \frac{\Gamma_+}{\Gamma}(1 - e^{-\Gamma\tau}). \tag{48}$$

Proof. The number-preserving quadratic Hamiltonian acts on the direct-sum one-particle space $\mathbb{C} \oplus \mathcal{H}_R$. Consequently its Heisenberg evolution is a Bogoliubov transformation without anomalous part: $a(t) = u_\lambda(t)a + B(f_{\lambda,t})$. Eliminating the reservoir component from the one-particle Schrödinger equation gives the Volterra equation for u_λ; setting $v_\lambda(\tau) = e^{i\Omega\tau/\lambda^2}u_\lambda(\tau/\lambda^2)$ gives (46). Quasi-freeness and gauge invariance kill the mixed first moments. Expanding $a(t)^*a(t)$ gives (45); differentiating its reservoir covariance gives (47). No factorization approximation is used. Replacing the delayed amplitudes by their value at τ, and extending the integrals to infinity, gives (48). $\square$

Theorem 6.2 (Finite-coupling Davies bound)**.** *Assume* (43)*. For every $T < \infty$ and every nonzero λ,*

$$\boxed{\sup_{0\le\tau\le T} \left\|\widetilde{\mathcal{R}}_\lambda(\tau) - e^{\tau\mathcal{L}_D}\right\|_{\mathrm{cb}} \le \epsilon_{\lambda,T} := 16\lambda^2 M_1(1 + M_0T)^2.} \tag{49}$$

Equivalently, in physical time the bound holds uniformly on $0 \le t \le T/\lambda^2$. It applies on the full local algebra $\mathcal{F}(O) \simeq M_2\overline{\otimes}\mathcal{R}_a$, with exactly the same constant.

Proof. Write $v = v_D$. Equation (46) and $|v_\lambda| \le 1$ give

$$|\dot{v}_\lambda(\tau)| \le M_0. \tag{50}$$

Subtract $\dot{v} = -\kappa v$. The delay defect is bounded by $\lambda^2 M_0 M_1$, while the omitted tail is $R_\lambda(\tau) = \int_{\tau/\lambda^2}^\infty |K(s)|ds$. Because $\operatorname{Re}\kappa = \Gamma/2 \ge 0$, variation of constants is contractive. Moreover,

$$\int_0^T R_\lambda(\tau)\, d\tau \le \lambda^2 \int_0^\infty s|K(s)|\, ds \le \lambda^2 M_1. \tag{51}$$

Hence

$$E_T := \sup_{0\le\tau\le T} |v_\lambda(\tau) - v(\tau)| \le \lambda^2 M_1(1 + M_0T). \tag{52}$$

Apply the same delay–tail decomposition to (47). Using $\big||v_\lambda|^2 - |v|^2\big| \le 2E_T$ gives

$$\sup_{0\le\tau\le T} |z_\lambda(\tau) - z_D(\tau)| \le 2\lambda^2 M_1(1 + M_0T)(1 + 2M_0T). \tag{53}$$

For completeness, this is just the integral of three terms: the delayed-amplitude defect $\lambda^2 M_0 M_1$, the tail controlled by (51), and the coefficient error $2M_0E_T$.

A unital gauge-covariant map on M_2 is fixed by its values on a, a^*, Q. Writing an arbitrary matrix, with arbitrary matrix-valued entries for a cb amplification, in this basis yields

$$\|\widetilde{\mathcal{R}}_\lambda - e^{\tau\mathcal{L}_D}\|_{\mathrm{cb}} \le 6E_T + 4|z_\lambda - z_D|. \tag{54}$$

Equations (52) and (53) give $\lambda^2 M_1[14 + 30M_0T + 16(M_0T)^2]$, which is bounded by the right-hand side of (49). Finally, $\|\Phi \otimes \mathrm{id}_{\mathcal{R}_a}\|_{\mathrm{cb}} = \|\Phi\|_{\mathrm{cb}}$, proving the local-algebra statement. $\square$

Corollary 6.3 (Stability of the maintenance law at finite coupling)**.** *Let $q_* := \min\{\Gamma_+/\Gamma, \Gamma_-/\Gamma\} > 0$, fix T, and suppose $\epsilon_{\lambda,T} \le 1$. If $\varphi_{\lambda,\tau}$ is the exact finite-coupling reduced state evolved from φ_0, then for $0 < \tau \le T$,*

$$\left|D(\varphi_{\lambda,\tau}\|\eta_D) - D(\varphi_\tau\|\eta_D)\right| \le h_2(\epsilon_{\lambda,T}/2) + \epsilon_{\lambda,T}|\log q_*|, \tag{55}$$

$$|\mathcal{P}_\lambda(\tau) - \mathcal{P}_D(\tau)| \le \frac{k_\mathrm{B}T}{\tau}\left[h_2(\epsilon_{\lambda,T}/2) + \epsilon_{\lambda,T}|\log q_*|\right]. \tag{56}$$

Here $h_2(x) = -x\log x - (1-x)\log(1-x)$, and both powers use the same relative-free-energy ledger as (41). Therefore

$$\lambda^2\log(1/|\lambda|) = o(\tau\log(1/\tau)) \tag{57}$$

is sufficient for the exact finite-coupling power to retain the leading coefficient $k_\mathrm{B}T\Gamma_+$ in a joint limit $(\lambda,\tau)\to(0,0)$.

Proof. The cb bound dualizes to trace-norm distance at most $\epsilon_{\lambda,T}$ on the selected M_2 mode. The spectator state is unchanged. The sharp two-dimensional entropy-continuity bound [10] gives $|S(\rho) - S(\sigma)| \le h_2(\|\rho-\sigma\|_1/2)$, while $|\operatorname{Tr}[(\rho-\sigma)\log\eta_D]| \le \|\rho-\sigma\|_1|\log q_*|$. This proves (55); division by τ proves (56). Since $h_2(c\lambda^2) = O(\lambda^2\log(1/|\lambda|))$, condition (57) makes the error negligible relative to $\log(1/\tau)$. □

Definition 6.4 (A compact explicit Dirac form factor)**.** Fix $\Omega > 0$, put $R = (2\Omega)^{-1}$, and let

$$b_R(\boldsymbol{x}) = \begin{cases} \exp\left[-\left(1-|\boldsymbol{x}|^2/R^2\right)^{-1}\right], & |\boldsymbol{x}| < R, \\ 0, & |\boldsymbol{x}| \ge R, \end{cases} \qquad g_R = \mathcal{N}(-\Delta)^2(b_R\xi), \tag{58}$$

where $0 \ne \xi \in \mathbb{C}^4$ and $\mathcal{N}$ normalizes the CAR one-particle norm. Thus $\widehat{g}_R(\boldsymbol{p}) = \mathcal{N}|\boldsymbol{p}|^4\widehat{b}_R(\boldsymbol{p})\xi$, and g_R is smooth and supported in the laboratory ball of radius R.

Lemma 6.5 (Araki–Wyss regularity and on-shell positivity)**.** *For the massless charged Dirac reservoir and the datum (58), choose the standard charge-conjugate unitary identification of the negative-energy Araki–Wyss fiber with the positive-energy multiplicity fiber. The two thermally dressed energy–angular form factors then have the common scalar parts*

$$\frac{u^5\widehat{b}_R(u\sigma)}{\sqrt{1+e^{-\beta u}}}, \qquad \frac{e^{-\beta u/2}u^5\widehat{b}_R(u\sigma)}{\sqrt{1+e^{-\beta u}}}, \tag{59}$$

times bounded positive-energy spin projectors and their charge conjugates. For every $0 < \theta' < \pi/\beta$, their energy translates and first two energy derivatives are uniformly in $L^2(\mathbb{R}\times S^2)$ for $|\operatorname{Im}\theta| \le \theta'$. Moreover,

$$J_{g_R}(\omega) = c_R\omega^{10}|\widehat{b}_R(\omega)|^2, \qquad c_R > 0, \tag{60}$$

up to the fixed normalization convention in (34); in particular $J_{g_R}(\Omega) > 0$ and $J_{g_R}(\omega) = O(\omega^{10})$ at threshold.

Proof. The massless radial spectral transform contributes one factor u, while $(-\Delta)^2$ contributes u^4. The two CAR thermal weights are $\sqrt{1-n_\beta(u)} = (1+e^{-\beta u})^{-1/2}$ and $\sqrt{n_\beta(u)} = e^{-\beta u/2}(1+e^{-\beta u})^{-1/2}$. A fixed unitary charge-conjugate gluing, including the harmless negative-fiber phase, turns the scalar threshold factor into u^5 on both sides of zero. The thermal square roots are analytic before their nearest zeros at $\operatorname{Im} u = \pm\pi/\beta$.

Paley–Wiener and integration by parts give, for every N and bounded complex strip,

$$|\partial_u^j \widehat{b}_R((u+\theta)\sigma)| \le C_{N,j,\theta'} e^{R|\operatorname{Im}\theta|}(1+|u|)^{-N}, \qquad j=0,1,2.$$

This proves uniform square integrability; the constants are not uniform as $\theta' \uparrow \pi/\beta$. Formula (34) and the massless shell measure give the power $\omega^2\omega^8 = \omega^{10}$. Since $\Omega r \le 1/2$ on $\operatorname{supp} b_R$,

$$\widehat{b}_R(\Omega) = 4\pi \int_0^R r^2 b_R(r) \frac{\sin(\Omega r)}{\Omega r}\, dr > 0.$$

Finally, $\int_{S^2} \Pi_+(\sigma) d\sigma = 2\pi \mathbf{1}$, because the term odd in σ integrates to zero. Hence the angular spinor factor is strictly positive for $\xi \neq 0$, proving $c_R > 0$ and the Fermi-golden-rule assertion. $\square$

Theorem 6.6 (Concrete all-time compact-Dirac approximation)**.** *Use the autonomous exchange model* (29) *with $m_R = 0$, $g = g_R$ from Definition 6.4, and a product initial state whose reservoir component is the β-KMS state. Include the second-order Lamb shift in $\mathcal{L}_D$. There are $\lambda_0 = \lambda_0(\beta, \Omega, R, \xi) > 0$ and $C < \infty$ such that*

$$\boxed{\sup_{t\ge 0} \left\| \mathcal{R}_\lambda(t) - e^{t(\mathcal{L}_S + \lambda^2 \mathcal{L}_D)} \right\|_{\mathrm{cb}} \le C|\lambda|, \qquad 0 < |\lambda| \le \lambda_0.} \tag{61}$$

Equivalently, after removal of the free oscillation, the same estimate holds uniformly for all kinetic times $\tau = \lambda^2 t$. It extends with identical cb norm to $\mathcal{F}(O) \simeq M_2 \overline{\otimes} \mathcal{R}_a$. No reduced-resolvent or reduced-remainder estimate is assumed.

Proof. Represent the graded tensor product by a Klein transformation. The odd–odd exchange becomes a finite sum of ordinary tensor products $G_j \otimes \phi_j(g_R)$, with bounded self-adjoint qubit matrices G_j and CAR field quadratures ϕ_j; the Klein parity commutes with the free dynamics and the even KMS state. Thus the reduced dynamics of even system observables lies in the spin–Fermi class of Ref. [7]. That paper's reservoir is explicitly a thermal free Fermi gas and already has the angular fiber $L^2(S^2)$; its deformation proof uses translation only in the energy variable. The Dirac charge and spin indices give a finite enlargement of that fiber. Lemma 6.5 verifies the translation-analyticity assumption directly.

For the exchange coupling, the Davies generator on M_2 has interaction-picture eigenvalues

$$0, \qquad -\Gamma, \qquad -\frac{\Gamma}{2} \pm i\Delta, \qquad \Gamma = \Gamma_+ + \Gamma_- > 0, \tag{62}$$

and detailed balance makes the zero eigenvector the Gibbs state. Equation (60) proves Fermi-golden-rule splitting. Analytic perturbation of the deformed Liouvillian gives the all-time reduced expansion

$$\mathcal{R}_\lambda(t) = \sum_j e^{itz_j(\lambda)} \widehat{P}_j(\lambda) + \widehat{E}_\lambda(t), \quad z_j(\lambda) = e_j + \lambda^2 z_j^{(2)} + O(\lambda^4), \quad \widehat{P}_j(\lambda) = \widehat{P}_j^{(0)} + O(|\lambda|), \tag{63}$$

with $\sup_{t\ge 0} \|\widehat{E}_\lambda(t)\| \le C|\lambda|$. The constants are uniform on the unit balls of system states and observables; the strip, and hence λ_0, may shrink as $\beta \to \infty$.

The second-order pole data in (63) are exactly those of $\mathcal{L}_S + \lambda^2 \mathcal{L}_D$, including Δ. For every decaying mode,

$$\sup_{t \ge 0} \lambda^4 t e^{-c\lambda^2 t} = O(\lambda^2),$$

while residue and contour errors are $O(|\lambda|)$; the stationary residue is covered by the same expansion. Reconstruction on the four matrix units gives the induced operator norm on maps of M_2, and the fixed-dimensional inequality $\|\Phi\|_{\mathrm{cb}} \le 2\|\Phi\|$ supplies the cb estimate. Tensoring the spectator identity preserves cb norm. This proves (61). $\square$

Remark 6.7 (What has actually been discharged)*.* Theorem 6.6 verifies a genuine all-time theorem for one displayed compactly supported Dirac datum, at the conservative resonance order $O(|\lambda|)$. It does not infer a second-order reduced residue estimate from parity alone. The next result records exactly the extra coefficient-level information needed for the sharper $O(\lambda^2)$ order.

Hypothesis 6.8 (Quadratic reduced resonance refinement)**.** *For a translation-analytic Dirac datum with the simple FGR resonances of Theorem 6.6, assume after applying the reservoir KMS expectation the reduced bounds*

$$z_j(\lambda) = e_j + \lambda^2 z_j^{(2)} + O(\lambda^4), \qquad \widehat{P}_j(\lambda) = \widehat{P}_j^{(0)} + O(\lambda^2), \qquad \|\widehat{R}_\lambda(t)\|_{\mathrm{cb}} \le C\lambda^2 e^{-c\lambda^2 t}.$$

Here hats denote reduced resonance residues and contour remainder, not unreduced Liouvillian projections. Assume separately the stationary matching condition

$$\|\rho_\infty(\lambda) - \rho_{\beta,S}\|_1 \le C\lambda^2,$$

where $\rho_\infty(\lambda)$ is the exact reduced interacting equilibrium state and $\rho_{\beta,S}$ is the invariant state of the second-order Davies generator. All constants are uniform for small real λ. These are additional coefficient-level hypotheses, not consequences claimed from Lemma 6.5.

Theorem 6.9 (Quadratic all-time Dirac–Davies refinement)**.** *Assume Hypothesis 6.8, use a product initial state with the reservoir in its β-KMS state, and include the second-order Lamb shift in $\mathcal{L}_D$. Let $\mathcal{R}_\lambda(t)$ denote the exact Schrödinger-picture reduced map and $\widetilde{\mathcal{R}}_\lambda(\tau)$ its interaction-picture form at $\tau = \lambda^2 t$. Then there are $\lambda_0, C < \infty$ such that, for $0 < |\lambda| \le \lambda_0$,*

$$\boxed{\sup_{\tau \ge 0} \|\widetilde{\mathcal{R}}_\lambda(\tau) - e^{\tau\mathcal{L}_D}\|_{\mathrm{cb}} \le C\lambda^2.} \tag{64}$$

Equivalently, in physical time the exact reduced map is $O(\lambda^2)$-close, uniformly for $t \ge 0$, to the Schrödinger-picture semigroup generated by $\mathcal{L}_S + \lambda^2\mathcal{L}_D$. The same estimate holds on $\mathcal{F}(O) \simeq M_2 \overline{\otimes} \mathcal{R}_a$.

Keeping the exact pole positions, residues, and stationary state gives a resonance approximation $\mathcal{R}_\lambda^{\mathrm{res}}(t)$, not automatically a quantum channel, with an exponentially decaying contour error. If, in addition, these reduced resonance data admit a trace-preserving GKLS embedding $e^{t\mathcal{M}(\lambda)}$ whose invariant state is the reduced interacting KMS state—as in the renormalized construction of Ref. [6]—then

$$\|\mathcal{R}_\lambda(t) - e^{t\mathcal{M}(\lambda)}\|_{\mathrm{cb}} \le C\big(|\lambda| + \lambda^2 t\big)e^{-c\lambda^2 t}, \qquad t \ge 0, \tag{65}$$

for some $c > 0$. Thus a GKLS-embeddable renormalized approximation is asymptotically exact, whereas the bare Davies error in (64) is uniformly small but need not vanish as $t \to \infty$. Without the embedding hypothesis, asymptotic exactness belongs to $\mathcal{R}_\lambda^{\mathrm{res}}(t)$, and no complete-positivity claim is made.

Proof. Analytic translation of the Araki–Wyss form factor moves the continuous Liouvillian spectrum below the real axis. The inverse resolvent contour then splits into the four reduced resonance residues and the remainder postulated in Hypothesis 6.8. The sign change $\lambda \mapsto -\lambda$ is implemented by reservoir parity, which leaves the KMS state invariant; hence the exact reduced map is even in λ. The additional reduced even-remainder bounds make this symmetry quantitative at the level required here. They give

$$z_j(\lambda) = e_j + \lambda^2 z_j^{(2)} + O(\lambda^4), \qquad \widehat{P}_j(\lambda) = \widehat{P}_j^{(0)} + O(\lambda^2).$$

The Fermi-golden-rule gap gives $\operatorname{Im} z_j^{(2)} \geq \gamma_{\mathrm{FGR}}$ away from the stationary resonance. Hence

$$\sup_{t\geq 0}\left| e^{itz_j(\lambda)} - e^{it(e_j+\lambda^2 z_j^{(2)})}\right| \leq C\lambda^4 \sup_{t\geq 0} te^{-c\lambda^2 t} \leq C'\lambda^2.$$

The stationary matching clause controls the zero resonance. Summing the finitely many residues proves the operator-norm version of (64); on M_2, operator norm and cb norm are equivalent with a fixed dimension constant. Tensoring with the spectator identity preserves cb norm. Keeping the exact resonance data and the interacting stationary state gives the exponentially accurate, but not automatically positive, resonance approximation. Under the additional GKLS-embedding clause, the renormalized construction of Ref. [6] gives (65). This is a conditional second-order refinement of Theorem 6.6: the reduced-residue, remainder, and stationary-matching clauses are exactly what upgrade its proved $O(|\lambda|)$ control to the displayed $O(\lambda^2)$ bound. □

Proposition 6.10 (Fixed-coupling Zeno obstruction)**.** *Let H_λ be any bounded autonomous system–reservoir Hamiltonian and let the initial joint state ρ_{SR} have system support in P. With $Q = \mathbf{1} - P$,*

$$\operatorname{Tr}\Big[(Q\otimes\mathbf{1})e^{-itH_\lambda}\rho_{SR}e^{itH_\lambda}\Big] = t^2 \operatorname{Tr}[(Q\otimes\mathbf{1})H_\lambda\rho_{SR}H_\lambda] + O(t^3). \tag{66}$$

Consequently the exact microscopic support-leakage derivative at fixed λ is zero. No bounded autonomous dilation can reproduce the Davies $t\log(1/t)$ law all the way down to physical time zero at fixed coupling.

Proof. Expand the two unitary factors to second order. The constant term vanishes because $(Q\otimes\mathbf{1})\rho_{SR} = 0$; cyclicity and $QP = 0$ kill the commutator term linear in t. The surviving quadratic double-commutator reduces to the displayed positive coefficient. Thus (64) is a uniform absolute approximation, but division by the refresh time is not uniform at the Zeno corner. The singular maintenance law remains a kinetic joint limit, as quantified by (57). □

7 Relative-Cauchy completion and the locality obstruction

The autonomous model is causal but tied to a stationary laboratory and to a selected rank-one CAR mode. We now place the coupling itself in the locally covariant framework of Brunetti–Fredenhagen–Verch [15], using the Dirac functor of Sanders [14], Green-hyperbolic propagation [16], and the system–probe scattering construction of Fewster–Verch [22]. The result both supplies the missing curved-spacetime dynamics and identifies a sharp obstruction: strict spacetime locality and exact one-particle rank one cannot hold simultaneously.

Let $\mathbf{M} = (M, g, \mathfrak{s})$ be an oriented, time-oriented, globally hyperbolic four-dimensional spin spacetime. Let $S_S, S_R \to M$ be two charged spinor bundles, with formally self-adjoint Dirac operators D_S, D_R, and write

$$D_0 = D_S \oplus D_R. \tag{67}$$

Choose a smooth compactly supported bundle morphism $\kappa : S_R \to S_S$, and define the gauge-neutral, formally self-adjoint zeroth-order coupling

$$R_\kappa = \begin{pmatrix} 0 & \kappa \\ \kappa^\dagger & 0 \end{pmatrix}, \qquad D_{\lambda,\kappa} = D_0 + \lambda R_\kappa, \qquad \lambda\in\mathbb{R}. \tag{68}$$

Its support $K = \operatorname{supp}\kappa$ is the coupling region. Both fields carry the same global charge, so (68) preserves the diagonal $U(1)$ action and fermion parity.

Let $G_0^\pm$ and $G_{\lambda,\kappa}^\pm$ be the advanced/retarded Green operators of D_0 and $D_{\lambda,\kappa}$, with a fixed common sign convention. Because R_κ is compactly supported, multiplication by R_κ sends every smooth solution to a compactly supported source. Hence the Møller maps on spacelike-compact solution spaces are the explicit operators

$$r_{\lambda,\kappa}^\pm = \mathbf{1} - \lambda G_{\lambda,\kappa}^\pm R_\kappa : \mathrm{Sol}_{\mathrm{sc}}(D_0) \longrightarrow \mathrm{Sol}_{\mathrm{sc}}(D_{\lambda,\kappa}), \tag{69}$$

$$(r_{\lambda,\kappa}^\pm)^{-1} = \mathbf{1} + \lambda G_0^\pm R_\kappa. \tag{70}$$

Indeed, the resolvent identities

$$G_{\lambda,\kappa}^\pm - G_0^\pm = -\lambda G_{\lambda,\kappa}^\pm R_\kappa G_0^\pm = -\lambda G_0^\pm R_\kappa G_{\lambda,\kappa}^\pm \tag{71}$$

prove both the field equations and the inverse relations. They also show that $r^\pm\phi = \phi$ outside $J^\pm(K)$, in the corresponding past/future identification region.

Theorem 7.1 (Covariant relative-Cauchy Dirac coupling)**.** *For every* $(\mathbf{M}, \kappa)$ *and* $\lambda \in \mathbb{R}$*:*

(i) $D_{\lambda,\kappa}$ *is Green-hyperbolic and has unique advanced and retarded Green operators;*

(ii) the past and future Møller identifications induce a unitary scattering map on the free solution space and hence a gauge-preserving CAR automorphism

$$\Theta_{\lambda,\kappa} = \mathrm{CAR}\Big[(r_{\lambda,\kappa}^+)^{-1} r_{\lambda,\kappa}^-\Big]\,; \tag{72}$$

this is the relative-Cauchy scattering morphism comparing the coupled and uncoupled theories through the time-slice isomorphisms;

(iii) the assignment $(\mathbf{M}, \kappa) \mapsto \Theta_{\lambda,\kappa}$ *is natural under causality-preserving spin embeddings that carry* κ *to the corresponding background coupling;*

(iv) $\Theta_{\lambda,\kappa}$ *fixes pointwise every local algebra in the causal complement* $K^\perp = M \setminus [J^+(K) \cup J^-(K)]$*;*

(v) for two couplings with causally ordered supports, the scattering morphism factorizes in that causal order. For causally disjoint supports the two orders coincide.

All statements descend from the field CAR algebra to the even gauge-invariant observable algebra.

Proof. The principal symbol of $D_{\lambda,\kappa}$ is that of the direct-sum Dirac operator because R_κ is zeroth order. Stability of Green hyperbolicity under smooth formally self-adjoint zeroth-order perturbations gives unique $G_{\lambda,\kappa}^\pm$ with $\operatorname{supp} G_{\lambda,\kappa}^\pm f \subset J^\pm(\operatorname{supp} f)$ [16]. Equations (69)–(71) then construct the two identifications rather than merely invoking them.

For any Cauchy surface Σ, formally self-adjoint Dirac evolution preserves

$$(\phi, \psi)_\Sigma = \int_\Sigma \overline{\phi}\, \gamma(n_\Sigma)\psi \, d\Sigma.$$

The Møller maps agree with the identity on Cauchy data before or after the interaction region, so they preserve this form. Therefore $s_{\lambda,\kappa} = (r_{\lambda,\kappa}^+)^{-1} r_{\lambda,\kappa}^-$ is unitary on the completed one-particle CAR space. Functorial CAR quantization gives (72); its commutation with the diagonal charge and parity actions follows directly from the block-off-diagonal, charge-neutral form of R_κ.

Let $\chi : (\mathbf{M}, \kappa) \to (\mathbf{N}, \kappa')$ be an orientation-, time-orientation-, metric-, and spin-preserving causally convex embedding carrying κ to κ'. Uniqueness of Green operators gives $\chi_* G_{\lambda,\kappa}^\pm = G_{\lambda,\kappa'}^\pm \chi_*$ on sources whose causal hull lies in the image. Substitution in (69) yields

$$\chi_* s_{\lambda,\kappa} = s_{\lambda,\kappa'} \chi_*, \qquad \mathcal{A}(\chi)\Theta_{\lambda,\kappa} = \Theta_{\lambda,\kappa'}\mathcal{A}(\chi), \tag{73}$$

which is the claimed naturality, with the usual time-slice restriction when the image does not contain all of $J(K)$.

If $O \subset K^{\perp}$, finite propagation makes both Møller identifications equal on solutions generated from test spinors in O; hence $s_{\lambda,\kappa}$ and $\Theta_{\lambda,\kappa}$ fix the corresponding local subspaces and algebras pointwise. Finally suppose K_1 is later than K_2 and choose Cauchy surfaces $\Sigma_- \prec K_2 \prec \Sigma_0 \prec K_1 \prec \Sigma_+$. Applying the time-slice identification successively across these surfaces gives

$$s_{\lambda,\kappa_1+\kappa_2} = s_{\lambda,\kappa_1} s_{\lambda,\kappa_2}, \qquad \Theta_{\lambda,\kappa_1+\kappa_2} = \Theta_{\lambda,\kappa_1} \circ \Theta_{\lambda,\kappa_2}. \tag{74}$$

For causally disjoint supports either ordering is admissible, so the two scattering morphisms commute. This is the concrete Green-operator realization of Fewster–Verch causal factorization [22]. ☐

The preceding theorem is fully background covariant, but it is not secretly the rank-one model of Section 6. The difference is forced by locality.

Proposition 7.2 (Exact rank-one locality obstruction)**.** *Let Σ be a smooth spacelike Cauchy surface meeting the interior of K. If the restriction of R_κ is nonzero on a subset of Σ of positive measure, then multiplication by $R_\kappa|_\Sigma$ on the one-particle Cauchy space has infinite rank. Consequently it cannot equal a nonzero finite-rank exchange*

$$|f\rangle\langle g| + |g\rangle\langle f|, \tag{75}$$

and its CAR quantization cannot be exactly $a^*B + B^*a$ for only two selected modes.

Moreover, the free Dirac Hamiltonian on Minkowski space has no nonzero compactly supported stationary one-particle mode. More generally, a stationary Dirac bound state for a smooth local potential cannot have compact spatial support under the standard unique-continuation hypothesis.

Proof. On a positive-measure subset where $R_\kappa \neq 0$, choose infinitely many mutually orthogonal smooth Cauchy data with disjoint sufficiently small supports and values outside the pointwise kernel. Their images are nonzero and mutually orthogonal, so the multiplication operator has infinite rank. Equation (75) has rank at most two. This proves the first claim.

For the free stationary claim, Fourier transformation diagonalizes the Minkowski Dirac Hamiltonian into the positive and negative mass shells, which have no L^2 point spectrum. With a smooth local stationary potential, the spatial Dirac equation is elliptic; unique continuation forbids an eigenfunction that vanishes on the open complement of a compact set [17]. ☐

Remark 7.3 (Why the obstruction is useful)*.* The selected-mode interaction of Sections 5–6 is a bounded operator localized in a laboratory algebra, but it is not the integral of a strictly local bilinear field density. Conversely, (68) is a strictly local covariant field equation, but it necessarily couples infinitely many one-particle directions. Proposition 7.2 rules out identifying the two models without a compression or approximation.

There is nevertheless an exact bridge. Specialize to a stationary slab of $\mathbf{M}$, allow a covariantly specified static external potential in D_S, and suppose the uncoupled one-particle Hamiltonian h_0 has a normalized simple bound mode f of energy Ω, separated from the remaining spectrum. Put

$$P_f = |f\rangle\langle f|, \qquad Q_f = \mathbf{1} - P_f, \qquad g = Q_f R_\kappa f, \qquad h_{Q,\mu} = Q_f(h_0 + \mu R_\kappa)Q_f. \tag{76}$$

Because R_κ is off diagonal and f lies in the system species, $P_f R_\kappa P_f = 0$.

Lemma 7.4 (Exact stationary Feshbach equation). *Within a stationary plateau of the coupling, the survival amplitude* $u_\lambda(t) = \langle f, e^{-ih_\lambda t} f\rangle$ *obeys*

$$\dot{u}_\lambda(t) + i\Omega u_\lambda(t) + \lambda^2 \int_0^t K_\lambda(t-s) u_\lambda(s)\, ds = 0, \qquad K_\mu(s) = \langle g, e^{-ih_{Q,\mu} s} g\rangle. \tag{77}$$

Prepare the complementary modes, for each μ, *in the gauge-invariant quasi-free* β*-KMS state of* $h_{Q,\mu}$. *The selected bound-mode CAR channel is then gauge-covariant and quasi-free, and its thermal injection coefficient satisfies the analogous stationary equation with occupied correlation kernel* $K_{+,\mu}$.

Proof. Project the one-particle Schrödinger equation for $e^{-ih_\lambda t} f$ with P_f and Q_f. Duhamel elimination of the Q_f-component gives (77) exactly. CAR second quantization of the resulting linear evolution and a gauge-invariant quasi-free state on the complementary modes give the channel statement and the occupied kernel. $\square$

The dependence of K_μ on μ records the complementary field modes that are absent in the rank-one reservoir. Introduce the uniform, directly checkable quantities

$$\mathsf{M}_j = \sup_{|\mu|\le\lambda_0} \int_0^\infty s^j \left(|K_\mu(s)| + |K_{+,\mu}(s)|\right) ds, \qquad j = 0, 1, \tag{78}$$

$$\mathsf{L}_0 = \sup_{0<|\mu|\le\lambda_0} \frac{1}{|\mu|} \int_0^\infty \left(|K_\mu(s) - K_0(s)| + |K_{+,\mu}(s) - K_{+,0}(s)|\right) ds. \tag{79}$$

Theorem 7.5 (Covariant-to-Davies quantitative bridge). *Assume* (78)–(79) *are finite and the on-shell Fermi-golden-rule rate at* Ω *is positive. Let* $\widetilde{\mathcal{R}}_\lambda^{\mathrm{loc}}$ *be the exact selected-bound-mode channel obtained from the locally coupled Dirac doublet, and let* $e^{\tau\mathcal{L}_D^{(0)}}$ *be the Davies channel computed from* $K_0, K_{+,0}$. *Then, for* $|\lambda| \le \lambda_0$,

$$\boxed{\sup_{0\le\tau\le T} \|\widetilde{\mathcal{R}}_\lambda^{\mathrm{loc}}(\tau) - e^{\tau\mathcal{L}_D^{(0)}}\|_{\mathrm{cb}} \le 16\left[\lambda^2 \mathsf{M}_1(1+\mathsf{M}_0 T)^2 + |\lambda|\mathsf{L}_0 T(1+\mathsf{M}_0 T)\right].} \tag{80}$$

If $Q_f R_\kappa Q_f = 0$, *then the kernels are independent of* λ, $\mathsf{L}_0 = 0$, *and* (80) *reduces to the* $O(\lambda^2)$ *bound of Theorem 6.2. For a nonzero strictly local multiplier this finite-rank condition is excluded by Proposition 7.2.*

Proof. In van Hove time, Lemma 7.4 is the exact Volterra equation (46) with K replaced by K_λ. Repeat the delay–tail estimate of Theorem 6.2 uniformly using $\mathsf{M}_0, \mathsf{M}_1$. Comparing the instantaneous integrals of K_λ and K_0 contributes at most $|\lambda|\mathsf{L}_0 T$ to the amplitude on $[0, T]$. The occupied kernel contributes the same way to the injection coefficient. The amplified M_2-basis estimate (54) then gives the second term in (80); the delay and tail terms give the first. If $Q_f R_\kappa Q_f = 0$, then $h_{Q,\lambda} = Q_f h_0 Q_f$ and both kernels are λ-independent. $\square$

Corollary 7.6 (Maintenance regime for the strictly local theory). *Let the right side of* (80) *be* $\epsilon_{\lambda,T}^{\mathrm{loc}} \le 1$. *Equations* (55)–(56) *hold with* $\epsilon_{\lambda,T}$ *replaced by* $\epsilon_{\lambda,T}^{\mathrm{loc}}$. *In the generic local case* $\mathsf{L}_0 > 0$, *the sufficient joint regime becomes*

$$|\lambda| \log(1/|\lambda|) = o(\tau \log(1/\tau))\,. \tag{81}$$

Thus the logarithmic coefficient survives strict local covariance, but the admissible approach to the rapid limit is narrower than in the exactly rank-one model.

Remark 7.7 (What "covariant" means here). The construction is a natural theory on the category whose objects include the compact coupling κ (and, in the stationary reduction, the external potential selecting the bound mode). This is background covariance in the standard locally covariant sense, not a claim that a nonzero compact coupling can be selected naturally from the metric and spin structure alone. The stationary plateau is used only to identify the Davies scaling limit; the scattering morphism (72) itself is defined on arbitrary globally hyperbolic spin spacetimes.

8 Soft support and the logarithmic cutoff law

The exact rank boundary $\rho Q = 0$ is an idealization. We now give a nonasymptotic identity for a softened target and then take the softening parameter to zero.

Let $P + Q = \mathbf{1}$. Let $\rho = P\rho P$ be faithful on $P\mathcal{H}$, let $\tau = Q\tau Q$ be faithful on $Q\mathcal{H}$, and define

$$\rho_\varepsilon = (1-\varepsilon)\rho + \varepsilon\tau, \qquad 0 < \varepsilon < 1. \tag{82}$$

Thus ρ_ε is faithful on the full finite-dimensional Hilbert space.

Theorem 8.1 (Exact soft-boundary entropy identity)**.** *Let $\mathcal{L}$ be a bounded trace-preserving Hermiticity-preserving generator for which $e^{t\mathcal{L}}\rho_\varepsilon$ is differentiable at $t = 0$. Put*

$$X_\varepsilon = \mathcal{L}(\rho_\varepsilon), \qquad b_\varepsilon = \operatorname{Tr}(QX_\varepsilon).$$

Then

$$\frac{d}{dt}S(e^{t\mathcal{L}}\rho_\varepsilon)\Big|_{t=0} = b_\varepsilon \log\frac{1-\varepsilon}{\varepsilon} - \operatorname{Tr}(PX_\varepsilon P\log\rho) - \operatorname{Tr}(QX_\varepsilon Q\log\tau). \tag{83}$$

If $e^{t\mathcal{L}}$ is positive for $t \geq 0$, then $a = \operatorname{Tr}[Q\mathcal{L}(\rho)] \geq 0$. For fixed $\rho, \tau, \mathcal{L}$,

$$\frac{d}{dt}S(e^{t\mathcal{L}}\rho_\varepsilon)\Big|_{t=0} = a\log(1/\varepsilon) + O(1), \qquad \varepsilon \downarrow 0. \tag{84}$$

For a fixed Hamiltonian H, the instantaneous negative free-energy derivative satisfies

$$-\frac{d}{dt}F(e^{t\mathcal{L}}\rho_\varepsilon)\Big|_{t=0} = k_{\mathrm{B}}T\,a\log(1/\varepsilon) + O(1). \tag{85}$$

Proof. For a faithful differentiable state path σ_t, trace preservation gives

$$\frac{d}{dt}S(\sigma_t)\Big|_{t=0} = -\operatorname{Tr}(\dot\sigma_0\log\sigma_0).$$

The block functional calculus of (82) gives

$$\log\rho_\varepsilon = P\big(\log(1-\varepsilon) + \log\rho\big)P + Q\big(\log\varepsilon + \log\tau\big)Q.$$

Off-diagonal blocks of X_ε do not contribute to the trace. Since $\operatorname{Tr}X_\varepsilon = 0$, one has $\operatorname{Tr}(PX_\varepsilon) = -b_\varepsilon$. Substitution proves (83).

Linearity gives

$$b_\varepsilon = (1-\varepsilon)a + \varepsilon\operatorname{Tr}[Q\mathcal{L}(\tau)] = a + O(\varepsilon).$$

The two block-log traces in (83) remain bounded because the block states are fixed and faithful. Also $\varepsilon\log(1/\varepsilon) \to 0$, proving (84). Positivity implies $a \geq 0$ by Proposition 2.1. Finally,

$$-\frac{dF}{dt} = -\operatorname{Tr}(HX_\varepsilon) + k_{\mathrm{B}}T\frac{dS}{dt},$$

and the energy derivative is bounded as $\varepsilon \downarrow 0$. □

Corollary 8.2 (Logarithmic cutoff law)**.** *If $a > 0$, exact rank deficiency and soft rank deficiency produce the same leakage coefficient but different logarithmic cutoffs:*

$$\begin{aligned} \rho \text{ rank deficient:} \quad & \frac{F(\rho) - F(e^{t\mathcal{L}}\rho)}{t} = k_{\mathrm{B}} T a \log(1/t) + O(1), \\ \rho_\varepsilon \text{ faithful:} \quad & -\frac{d}{dt} F(e^{t\mathcal{L}}\rho_\varepsilon)\Big|_0 = k_{\mathrm{B}} T a \log(1/\varepsilon) + O(1). \end{aligned}$$

Therefore a small full-rank floor regularizes the rapid-time divergence at fixed ε, but the regularized rate diverges as the floor is removed.

Remark 8.3 (Hamiltonian motion). For the block-diagonal softened state, adding $-i[K, \cdot]$ changes neither b_ε nor the entropy derivative: cyclicity gives $\mathrm{Tr}(Q[K, \rho_\varepsilon]) = 0$ and $\mathrm{Tr}([K, \rho_\varepsilon] \log \rho_\varepsilon) = 0$. A Hamiltonian may rotate a support projection and create population relative to the fixed initial P at order t^2, but it preserves rank. It cannot generate the first-order kernel spectrum responsible for either logarithm in Corollary 8.2.

8.1 An exactly solvable crossover

Example 8.4 (One-way two-level leakage)**.** Let $J = |1\rangle\langle 0|$, $\mathcal{L} = \gamma \mathcal{D}[J]$, $H = 0$, and $\rho_\varepsilon = \mathrm{diag}(1-\varepsilon, \varepsilon)$. The excited population is exactly

$$p_\varepsilon(t) = 1 - (1 - \varepsilon) e^{-\gamma t},$$

and the entropy contribution to exact-refresh power is

$$\Pi_\varepsilon(t) = \frac{k_{\mathrm{B}} T}{t} \left[h_2(p_\varepsilon(t)) - h_2(\varepsilon) \right], \tag{86}$$

where $h_2(x) = -x \log x - (1 - x) \log(1 - x)$. At fixed $\varepsilon > 0$,

$$\lim_{t \downarrow 0} \Pi_\varepsilon(t) = k_{\mathrm{B}} T \gamma (1 - \varepsilon) \log \frac{1 - \varepsilon}{\varepsilon}.$$

At $\varepsilon = 0$, $\Pi_0(t) = k_{\mathrm{B}} T \gamma \log(1/t) + O(1)$. Hence the same exact model displays both branches of the cutoff law.

The fixed-state no-go still permits a singular *family* of faithful preparations. The next theorem gives the exact crossover exponent and also holds, with Araki relative entropy in place of matrix entropy, for the binary face constructed in Theorem 3.1.

Theorem 8.5 (Faithful-family realization and exponent saturation)**.** *Let $\varphi_0, \varphi_1, \eta_r$ and (Φ_t) be as in Theorem 3.1, put*

$$\Lambda = \kappa_+ + \kappa_-, \qquad \pi = \frac{\kappa_+}{\Lambda}, \qquad \varphi_\varepsilon = (1 - \varepsilon)\varphi_0 + \varepsilon \varphi_1.$$

Every φ_ε, $0 < \varepsilon < 1$, is faithful, and its exact orbit is

$$\Phi_{t,*}\varphi_\varepsilon = (1 - p_\varepsilon(t))\varphi_0 + p_\varepsilon(t)\varphi_1, \qquad p_\varepsilon(t) = \pi + (\varepsilon - \pi) e^{-\Lambda t}. \tag{87}$$

Define the Araki-relative-free-energy restoration power

$$\mathcal{P}_\varepsilon(t) = \frac{k_{\mathrm{B}} T}{t} \left[D(\varphi_\varepsilon \| \eta_r) - D(\Phi_{t,*}\varphi_\varepsilon \| \eta_r) \right]. \tag{88}$$

For every $\alpha > 0$, the faithful choice $\varepsilon(t) = t^{\alpha}$ satisfies

$$\boxed{\mathcal{P}_{t^{\alpha}}(t) = k_{\mathrm{B}} T \kappa_{+} \min\{\alpha, 1\} \log(1/t) + O(1), \qquad t \downarrow 0.} \tag{89}$$

Thus the full rank-boundary coefficient is recovered by faithful states exactly when their floor closes at least linearly, $\alpha \geq 1$. At every nonzero t the prepared state is faithful; no single fixed faithful state is claimed to be singular.

Proof. The two corner states have orthogonal supports summing to $\mathbf{1}$, and each is faithful on its corner, so positive weights make φ_{ε} faithful. The rate equation gives (87). The two data-processing equalities in the proof of Theorem 3.1 give

$$D(\varphi_p \| \eta_r) = (1-p) \log \frac{1-p}{1-\pi} + p \log \frac{p}{\pi}.$$

Moreover $p_{\varepsilon}(t) - \varepsilon = \kappa_{+} t + O(\varepsilon t + t^2)$. If $0 < \alpha < 1$, this increment is $o(\varepsilon)$, so the mean-value theorem evaluates the logarithmic derivative at $\varepsilon = t^{\alpha}$ and gives $\alpha \log(1/t)$. If $\alpha = 1$, every intermediate population is of order t. If $\alpha > 1$, integrate the binary logarithmic derivative from t^{α} to $\kappa_{+} t + o(t)$; the lower endpoint contributes $o(t \log(1/t))$. The latter two cases both give $\log(1/t)$. Division by t proves (89). □

Corollary 8.6 (Faithful physical ambient state, conditional boundary)**.** *Let ω be a faithful normal vacuum or KMS restriction and let P be a nontrivial local projection. The filtered preparation*

$$\omega_P(x) = \frac{\omega(PxP)}{\omega(P)}$$

is normal but has support P, hence is not faithful on the ambient algebra. Adding any corner-faithful state χ_Q on $Q\mathcal{M}Q$,

$$\omega_{P,\varepsilon} = (1-\varepsilon)\omega_P + \varepsilon \chi_Q, \qquad Q = \mathbf{1} - P,$$

produces a faithful normal local preparation for every $\varepsilon > 0$. If χ_Q is chosen as the transported corner state and the binary QMS is implemented as in Theorem 3.1, Theorem 8.5 realizes the singular law inside the folium of a faithful physical ambient state. The singular $\varepsilon = 0$ conditional target itself remains nonfaithful. The probability and work cost of the filtering operation are additional ledger data.

9 Uniform regulator families and noncommuting limits

Pointwise asymptotics at each regulator size do not justify interchanging the regulator and rapid-control limits. We state the needed uniformity explicitly.

For each n, let $(\rho_n, H_n, \mathcal{L}_n)$ be a finite-dimensional maintenance datum, $P_n = \operatorname{supp} \rho_n$, and

$$a_n = \mathrm{Tr}[(\mathbf{1} - P_n)\mathcal{L}_n(\rho_n)].$$

Let

$$\Pi_n(t) = \frac{F_n(\rho_n) - F_n(e^{t\mathcal{L}_n}\rho_n)}{t}$$

be exact-refresh power in the same ledger used to derive (2).

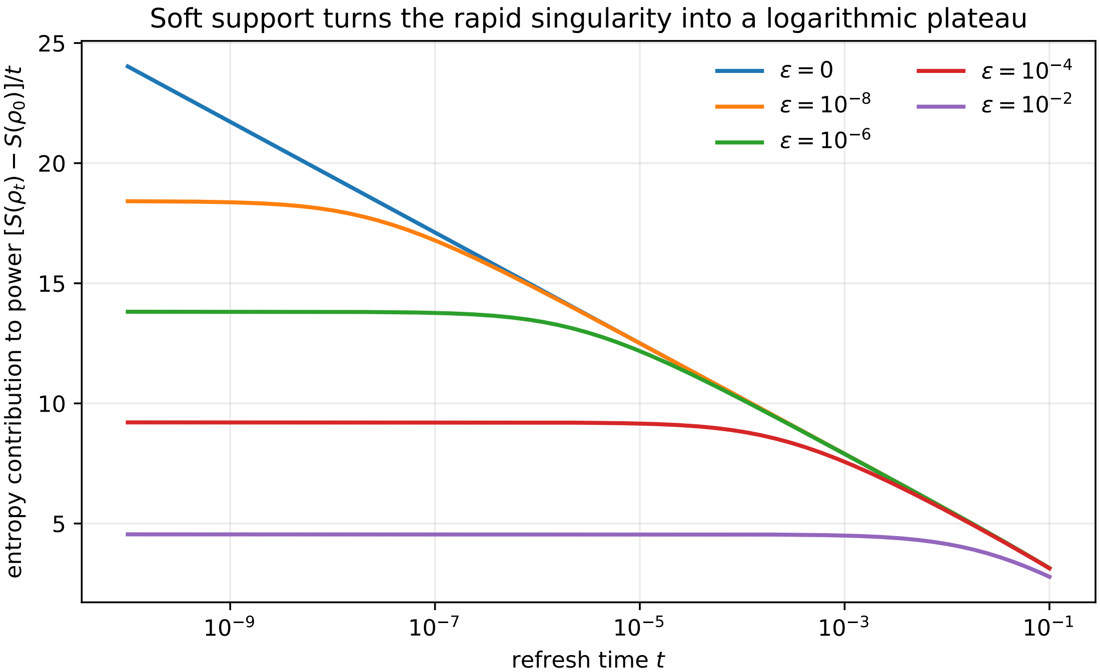

Figure 1: Exact two-level crossover from Example 8.4, in units $k_\mathrm{B}T = \gamma = 1$. A full-rank floor ε replaces the $\log(1/t)$ divergence by a plateau of height approximately $\log(1/\varepsilon)$. The $\varepsilon = 0$ curve retains the rank-boundary singularity.

Hypothesis 9.1 (Uniform boundary expansion)**.** *There exist $t_0 > 0$, a bounded real sequence (c_n), and a function $r(t) \downarrow 0$ such that, for all n and $0 < t < t_0$,*

$$|\Pi_n(t) - k_\mathrm{B}Ta_n \log(1/t) - c_n| \le r(t). \tag{90}$$

This hypothesis is stronger than applying the boundary expansion separately at every n. It can fail if target spectral gaps close, generator norms grow, the first-order leakage spectrum develops an unbounded logarithmic moment, or the thermodynamic Hamiltonians are not uniformly controlled. A typical sufficient package is: uniform lower bounds on the nonzero spectrum of ρ_n, uniform first and second path bounds, uniform $\|H_n\|$, and a uniform bound on $-\operatorname{Tr}(A_n \log A_n)$ for $A_n = Q_n\mathcal{L}_n(\rho_n)Q_n$. We retain (90) as the audit-friendly statement because it is precisely what the limit theorem uses.

Theorem 9.2 (Joint regulator–refresh criterion)**.** *Assume Hypothesis 9.1 and let $t_n \downarrow 0$. Then*

$$\Pi_n(t_n) = k_\mathrm{B}Ta_n \log(1/t_n) + O(1). \tag{91}$$

Consequently:

(i) if $a_n \log(1/t_n) \to \infty$, then $\Pi_n(t_n) \to \infty$;

(ii) if $a_n \log(1/t_n) \to 0$, the singular contribution disappears, although the bounded regular part need not vanish;

(iii) if $a_n \log(1/t_n) \to \lambda \in (0, \infty)$, the singular contribution approaches $k_\mathrm{B}T\lambda$, while convergence of the full power additionally requires convergence of c_n.

Proof. Evaluate (90) at $t = t_n$. Boundedness of c_n and $r(t_n) \to 0$ gives (91); the three cases follow. $\square$

Corollary 9.3 (Refresh exponent subtraction)**.** *Suppose*

$$a_n = e^{-\alpha n+o(n)}, \qquad \log(1/t_n) = e^{\gamma n+o(n)}, \qquad \alpha > 0,\ \gamma \geq 0.$$

Then the singular term in (91) *is*

$$k_{\mathrm{B}}T \exp[-(\alpha-\gamma)n + o(n)]. \tag{92}$$

Thus the growth exponent is $\gamma - \alpha$, equivalently the decay exponent is $\alpha - \gamma$. A superexponential refresh schedule can erase or overcome an exponential geometric leakage suppression according as $\gamma < \alpha$ or $\gamma > \alpha$; when $\gamma = \alpha$, the displayed $o(n)$ terms decide the marginal behavior. By contrast, if $t_n = e^{-\beta n+o(n)}$, then $\log(1/t_n) = \beta n + o(n)$ and the exponential decay rate remains α, with only a polynomial factor.

Remark 9.4 (Why this differs from fixed-latency exponent transmutation)*.* The SSH results in Refs. [29, 31] compare distinct microscopic leakage channels at a fixed refresh period and obtain minima of geometric and filter exponents. Corollary 9.3 concerns a different operation: the refresh time itself varies with regulator size. The two mechanisms compose but should not be identified.

The same theorem applies to soft target floors. If the remainder in Theorem 8.1 is uniform over n, then replacing t_n by ε_n shows that the instantaneous softened loss is controlled by $a_n \log(1/\varepsilon_n)$.

10 The split-regulated AQFT interface

10.1 Split data are structural, not dynamical

Let $O_1 \Subset O_2$ be local regions in a Haag–Kastler net. The split property provides a type-I factor

$$\mathcal{A}(O_1) \subset \mathcal{N} \subset \mathcal{A}(O_2). \tag{93}$$

It is a strong statistical-independence property and follows from nuclearity in broad classes of models [23, 24]. For standard split inclusions with specified standard data, a canonical intermediate type-I factor can be selected; this is the basis of rigorous split entropies in QFT [25]. However, a collar-dependent sequence, its finite ultraviolet corners, and the induced open dynamics are not fixed by the word “split” alone.

Because $\mathcal{N}$ contains the type-III factor $\mathcal{A}(O_1)$, the intermediate type-I factor is necessarily type I_∞, not a matrix algebra. To import the finite-dimensional expansion without hidden trace-class assumptions, an explicit finite corner is mandatory.

Definition 10.1 (Split-corner maintenance datum)**.** At regulator label n, a split-corner maintenance datum consists of:

(a) regions $O_{1,n} \Subset O_{2,n}$ and a chosen type-I factor $\mathcal{A}(O_{1,n}) \subset \mathcal{N}_n \subset \mathcal{A}(O_{2,n})$;

(b) a declared type-I identification $\mathcal{N}_n \simeq B(\mathcal{K}_n)$ and a finite-rank projection $e_n \in \mathcal{N}_n$, giving the matrix corner $\mathcal{C}_n = e_n \mathcal{N}_n e_n$;

(c) a target density matrix $\rho_n \in \mathcal{C}_n$, its support P_n, a Hamiltonian $H_n \in \mathcal{C}_n$, and a positive trace-preserving semigroup with generator $\mathcal{L}_n$ on $\mathcal{C}_n$;

(d) a thermodynamic ledger specifying why exact restoration is priced by the free-energy difference.

All four items are part of the regulator. In particular, P_n is not determined by $\mathcal{A}(O_{1,n}) \subset \mathcal{A}(O_{2,n})$.

Theorem 10.2 (Conditional split-regulated survival theorem)**.** *Let a family of split-corner data satisfy Hypothesis 9.1, and let $t_n \downarrow 0$. Then the diagonal rank-boundary contribution to exact-refresh power is governed by the joint parameter*

$$\Lambda_n = a_n \log(1/t_n), \qquad a_n = \mathrm{Tr}_{\mathcal{C}_n}[(e_n - P_n)\mathcal{L}_n(\rho_n)]. \tag{94}$$

The conclusions of Theorem 9.2 hold. In particular, divergence is equivalent to $\Lambda_n \to \infty$ under the uniform expansion. A finite nonzero limit of Λ_n is not identifiable from the full power without separate control of the regular term c_n. No conclusion about Λ_n follows from the split property, finite split entropy, or static recoverability alone.

Proof. Each corner is a matrix algebra, so Theorem 9.2 applies. The last sentence follows because the split property constrains the inclusion (93), whereas a_n also depends on e_n, P_n, ρ_n, and $\mathcal{L}_n$. None of those dynamical/support data is fixed by split existence. □

Corollary 10.3 (Faithful continuum targets)**.** *Let φ be a faithful normal state on the limiting local algebra. Its intrinsic support-leakage coefficient from Proposition 2.1 is zero. Therefore any nonzero limit inferred solely from rank-deficient split-corner approximants is not established as an intrinsic property of φ unless an additional theorem relates the corner coefficients a_n to a different, well-defined continuum observable or relative-entropy derivative.*

Remark 10.4 (Vacuum and KMS states). In standard AQFT representations, Reeh–Schlieder properties (cyclicity and separation) make local vacuum restrictions faithful, and thermal KMS reference states are also normally taken faithful. Corollary 10.3 applies to the rank-boundary mechanism only. It does not imply zero entropy production, zero Dirichlet form, or reversible dynamics. Quantum Markov semigroups with faithful normal invariant states on arbitrary von Neumann algebras have a substantial modular and spectral-gap theory [27]; those rates are different objects from (3).

10.2 What can be intrinsic in type III

There are two legitimate continuum routes.

A nonfaithful target on the type-III algebra. If a normal target φ has proper support $p < \mathbf{1}$, the leakage derivative (3) is intrinsic. Theorem 3.1 supplies a complete regulator-free realization. Theorems 4.1 and 5.1 then derive it first from thermal probe collisions and, more strongly, from one continuous relativistic Dirac KMS reservoir without resets. For arbitrary nonfaithful targets and arbitrary semigroups, however, support conditions for finiteness and domains of relative modular operators still cannot be omitted [11, 26].

A faithful target with a different operational corridor. For a faithful state there is no complement of its support, but one may choose a code projection, spectral window, conditional expectation, or error observable. Leakage from that chosen object can be operationally meaningful, yet it is no longer the support leakage of the state. It needs its own stability and ledger theorem.

This distinction closes a gap in the existing programme. Split-regularized recoverability is a static reconstruction statement [32]; rapid-maintenance singularity is a dynamical rank-boundary statement [29]. Neither implies the other without a typed bridge.

11 Counterexamples and audit rules

The following failures delimit the theorem.

1. **Pointwise is not uniform.** An expansion with $O_n(1)$ at every n does not imply (91). A growing regular term may dominate $a_n \log(1/t_n)$.
2. **State convergence is insufficient.** Norm convergence of regulated states does not control a_n if the support projections, compressed generators, or corners vary. Generator/corner compatibility must be proved.
3. **Type I is not finite dimension.** A split factor is usually $B(\mathcal{K})$ with $\dim \mathcal{K} = \infty$. Infinite-dimensional entropy derivatives can be infinite for reasons unrelated to a finite-rank kernel. Definition 10.1 makes the extra cutoff visible.
4. **Faithful does not mean regular.** Full support removes the finite-dimensional rank-boundary coefficient, but eigenvalues may accumulate at zero in type I_∞, and relative modular operators are unbounded in type III. Differentiability still requires domain control.
5. **Recoverability does not set a rate floor.** Exponentially small conditional mutual information can yield a recovery channel, but it does not determine the outward GKLS coefficient a_n. Static Markovness and dynamic leakage require a separate bridge.
6. **A thermodynamic ledger is part of the claim.** Free-energy difference is exactly achievable with fresh preloaded target cells in the replacement model. It is not automatically the minimum work of an arbitrary autonomous controller.

These are not merely disclaimers. They provide falsifiers: a proposed AQFT maintenance theorem must state a support or corridor, a normal dynamics, a relative-free-energy domain, a split/corner selection rule when used, uniform remainder control, and an operational ledger.

12 Numerical audit and reproducibility

The proofs above are analytic. A deterministic script accompanies this draft to test the identities most vulnerable to sign, parity, or block-convention errors.

First, forty random two-rate instances of the two-sector generator (12) were evaluated in its matrix representative. Invariance of η_r, the exact orbit (7), and the binary relative-entropy identity (8) held to respective absolute errors 1.3×10^{-15}, 1.6×10^{-16}, and 6.7×10^{-16}. This checks the rate and sign conventions; the proof of the general-factor result remains the two-channel data-processing argument.

Second, forty random four-level GKLS examples were generated with a two-dimensional target support, a faithful two-dimensional softening block, a random Hamiltonian, and a random jump operator. The positive square-sum identity (5) agreed with the direct Q-trace to 1.8×10^{-15}. The exact entropy-slope identity (83) agreed with direct matrix functional calculus to 8.9×10^{-14}. A Richardson-extrapolated semigroup finite difference agreed within relative error 7.5×10^{-7}.

Third, the graded-CAR collision construction was represented by a Jordan–Wigner transform and tested for forty random choices of (β, Ω, g). The cb-matrix representative of $(T_\delta - \mathrm{id})/\delta$ agreed with (21) to absolute error below 3.7×10^{-4}. The error of $T^n_{t/n}$ at $n = 512$ was below 1.2×10^{-4}, and doubling n reduced it by a factor 0.5001, numerically confirming the predicted first-order product convergence and, crucially, the fermionic parity signs.

Fourth, forty random massive and massless on-shell Dirac-bath instances were used to audit (33). The fermionic KMS ratio agreed with $e^{-\beta\Omega}$ to 5.6×10^{-17}; stationarity of (39) and the exact orbit (40) held to 3.2×10^{-16} and 1.4×10^{-17}, while the absorption/leakage coefficient agreed exactly in floating-point arithmetic. This checks the thermal signs and stationary weights, not the analytic Davies limit itself.

Fifth, the explicit compact datum (58) was evaluated at $\Omega = 1$, $R = 1/2$. Direct radial quadrature gave $\widehat{b}_R(\Omega) = 0.05437047915 > 0$, while $\sin(1/2)/(1/2) = 0.9588510772$ is a rigorous positive lower bound for the sinc factor throughout the support. On eighty logarithmically spaced frequencies, $J(\omega)/\omega^{10}$ stayed between 2.956×10^{-3} and 3.040×10^{-3}. This audits the on-shell and threshold arithmetic of Lemma 6.5; Paley–Wiener analyticity remains an analytic proof.

Sixth, the finite-coupling amplitude estimate (52) was tested on one hundred exponential kernels $K(s) = \gamma e^{-\alpha s}$. For this family the Volterra equation embeds into an exactly solved two-dimensional linear ODE, so no quadrature or time-stepping error enters. Across 401 times per instance, the largest ratio of the exact error to the analytic bound was 0.9964, and the bound margin remained positive. The near saturation is useful: it shows that the first-moment estimate is not merely dimensionally correct but can be sharp.

Seventh, the all-time scaling was tested on sixty analytic exponential kernels over 901 physical times extending to twenty-four Davies lifetimes. The largest uniform amplitude error divided by λ^2 was 1.245; at the terminal time the largest absolute remainder was 2.94×10^{-12}. This does not prove spectral deformation, but it stress-tests the pole expansion and the conditional uniform $O(\lambda^2)$ scaling on an exactly soluble family.

Eighth, the faithful-family exponent law (89) was tested for $\alpha = 0.25, 0.5, 1, 1.5, 2$ at three decreasing times. All five normalized errors decreased toward $\min\{\alpha, 1\}$; at $t = 10^{-10}$ the largest absolute normalized error was 4.35×10^{-2}, consistent with the expected logarithmically slow $O(1/\log(1/t))$ convergence.

Ninth, Equation (86) was evaluated on a logarithmic grid to generate Figure 1. At the smallest plotted time, the softened curves agreed with their analytic plateaux to relative error below 2.8×10^{-4}.

12.1 Reproducibility contract

The companion archive `AQFT_rapid_maintenance_v8_reproducibility.zip` contains:

- the final PDF and complete LaTeX source;
- `verify_soft_boundary.py`, `parameters.json`, and `verification_report.json`;
- `dirac_form_factor_samples.csv` and both figure formats; and
- pinned requirements, a rerun README, and `SHA256SUMS.txt`.

From the extracted directory,

```
python verify_soft_boundary.py
```

regenerates the JSON report, CSV table, and both figure files. The fixed seed is 20260805. Every numerical claim in this section is addressable by a named JSON field; the archive hash manifest detects accidental replacement or omission.

13 Discussion and novelty boundary

The strongest conclusion is now algebraic, microscopic, concrete all-time, and background covariant. On every sigma-finite properly infinite local factor, one can choose a nonfaithful normal target and a bounded commutant-fixing local QMS for which both average restoration power and instantaneous Araki-relative-entropy production diverge with the exact coefficient $k_{\mathrm{B}}T\kappa_+$. In the free charged Dirac model, the same QMS is obtained both from unitary thermal-probe collisions and from one continuous Dirac KMS reservoir. In the autonomous realization the coefficient is the on-shell absorption rate $\Gamma_+ = 2\pi J_g(\Omega)n_\beta(\Omega)$. The direct moment estimate controls finite van Hove windows. Theorem 6.6 removes the window for the explicit compact datum (58) with a proved $O(|\lambda|)$ cb error; Theorem 6.9 states the sharper $O(\lambda^2)$ order only under its additional reduced-coefficient hypotheses. Proposition 6.10 proves that either estimate is optimal in kind: exact bounded Hamiltonian leakage is quadratic at physical time zero. For a strictly local two-Dirac coupling, Theorem 7.1 supplies relative-Cauchy scattering on curved spacetime, Proposition 7.2 proves why exact rank one is impossible, and Theorem 7.5 quantifies the resulting $O(|\lambda|)$ correction. The binary formula (8) is not an analogy: two data-processing inequalities identify the full modular relative entropy with a classical two-point divergence.

This does not make the effect universal. Proposition 3.3 turns the faithful-vacuum/KMS caveat into a theorem: a fixed faithful state has support $\mathbf{1}$ and cannot carry this rank-boundary mechanism. Theorem 8.5 supplies the optimal replacement. A family that is faithful at every nonzero floor $\varepsilon = t^\alpha$ recovers the fraction $\min\{\alpha, 1\}$ of the singular coefficient; a conditioned code state realizes the full boundary inside the folium of a faithful physical state, but the zero-floor conditional target is explicitly nonfaithful. If a different code boundary is imposed through split corners, its continuum significance remains the stability problem encoded by Λ_n in (94).

The positive advance is therefore a pair of saturation laws plus a two-tier all-time result. The explicit compact Dirac datum gives unconditional $O(|\lambda|)$ control; reduced even remainders would sharpen it to $O(\lambda^2)$. The Zeno proposition identifies the exact microscopic endpoint neither approximation can cross. The faithful-family theorem saturates the state-space regularization at $\min\{\alpha, 1\}$, while the faithful-state proposition identifies the exact support endpoint it cannot cross.

Relative to the preceding corpus [28, 29, 30, 31, 32], the new content is:

1. the exact regulator-free type-III construction (12), including a faithful invariant reference and a causally local QMS;
2. the exact reduction of Araki relative entropy to the binary law (8), and the resulting power and entropy-production singularities (9)–(10);
3. the free-Dirac collision realization (17)–(23), deriving the QMS and its thermal detailed-balance ratio from unitary local probe interactions;
4. the autonomous finite-bandwidth realization (29)–(36), deriving the same QMS from one relativistic Dirac KMS reservoir with a fixed smooth spatial form factor and no resets;
5. the on-shell singular law (41), whose coefficient is the bath absorption rate $2\pi J_g(\Omega)n_\beta(\Omega)$;
6. the exact scalar memory equations (46)–(47), the finite-coupling cb estimate (49), and the joint stability regime (57);
7. the compact form factor (58), its discharged Araki–Wyss and FGR estimates (59)–(60), the concrete all-time bound (61), the conditional quadratic refinement (64), and the fixed-coupling Zeno obstruction (66);
8. the relative-Cauchy Dirac-doublet construction (68)–(72), including naturality, time-slice propagation, spacelike triviality, and causal factorization;

9. the exact locality obstruction of Proposition 7.2, the Feshbach memory equation (77), and the corrected local-field bound (80);
10. the sharp fixed-faithful no-go of Proposition 3.3, the faithful-family saturation law (89), and the conditional physical-state corridor of Corollary 8.6;
11. the exact finite-dimensional softened-target identity (83) and state-space logarithmic cutoff (85);
12. the uniform joint-limit criterion (91) and refresh-exponent subtraction (92); and
13. the split-corner contract identifying when a separately imposed cutoff code still needs regulator stability.

The remaining boundary is one of model scope and operational accounting, not an unclassified singular corner. The concrete all-time theorem applies to the autonomous selected-mode CAR model with (58); extending the same deformation uniformly to the strictly local two-Dirac Feshbach family would require control of the complementary resonances. The $O(\lambda^2)$ order remains conditional, but the existence of an all-time Dirac approximation no longer is. Polynomially decaying, nonanalytic form factors may retain a weaker all-time coupling power, as in Ref. [8]. Background covariance does not select a nonzero compact κ, a stationary potential, a code projection, or a filtering apparatus from the metric alone; they remain typed physical data. Nor does the conditional-state corridor price the success probability and work cost of preparation. A fixed faithful vacuum/KMS support singularity is ruled out by Proposition 3.3, and fixed-coupling first-order short-time leakage by Proposition 6.10.

Acknowledgement of proof audit

The central identities were stress-tested during drafting and then re-derived directly in the displayed conventions. In particular, the type-III entropy formula was checked without density matrices by applying data processing in both directions between $\mathcal{M}$ and $\mathbb{C}^2$. Adversarial Claude Fable 5 High audits, run through the explicitly selected `masterythief` profile, led to four qualifications incorporated above: stationary matching is independent, exact resonance residues are not automatically completely positive, the second Araki–Wyss weight is $e^{-\beta u/2}$, and the negative-energy branch requires an explicit charge-conjugate identification. The applicability of Ref. [7] was then checked against its displayed thermal free-Fermi model, and the Klein reduction for the present even exchange was written into the proof. Numerical audits are checks, not evidence replacing proof.

References

[1] M. R. Grace and S. Guha, "Perturbation Theory for Quantum Information," *IEEE Information Theory Workshop*, 500–505 (2022), doi:10.1109/ITW54588.2022.9965836; expanded version arXiv:2106.05533.

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

[3] G. Lindblad, "On the generators of quantum dynamical semigroups," *Commun. Math. Phys.* **48**, 119–130 (1976), doi:10.1007/BF01608499.

[4] E. B. Davies, "Markovian master equations," *Commun. Math. Phys.* **39**, 91–110 (1974), doi:10.1007/BF01608389.

[5] E. B. Davies, "Markovian master equations. II," *Math. Ann.* **219**, 147–158 (1976), doi:10.1007/BF01351898.

[6] M. Merkli, "Quantum Markovian master equations: resonance theory shows validity for all time scales," *Ann. Phys.* **412**, 167996 (2020), doi:10.1016/j.aop.2019.167996; arXiv:1908.01984.

[7] M. Merkli and S. Starr, "A resonance theory for open quantum systems with time-dependent dynamics," *J. Stat. Phys.* **134**, 871–898 (2009), doi:10.1007/s10955-008-9645-5; arXiv:0810.3540.

[8] M. Merkli, "Dynamics of Open Quantum Systems II, Markovian Approximation," *Quantum* **6**, 616 (2022), doi:10.22331/q-2022-01-03-616.

[9] H. Spohn, "Entropy production for quantum dynamical semigroups," *J. Math. Phys.* **19**, 1227–1230 (1978), doi:10.1063/1.523789.

[10] K. M. R. Audenaert, "A sharp continuity estimate for the von Neumann entropy," *J. Phys. A* **40**, 8127–8136 (2007), doi:10.1088/1751-8113/40/28/S18.

[11] H. Araki, "Relative entropy of states of von Neumann algebras," *Publ. Res. Inst. Math. Sci.* **11**, 809–833 (1976), doi:10.2977/prims/1195191148.

[12] F. Hiai, "Quantum f-divergences in von Neumann algebras I: Standard f-divergences," *J. Math. Phys.* **59**, 102202 (2018), doi:10.1063/1.5039973.

[13] Y. Kitajima, "Local Operations and Completely Positive Maps in Algebraic Quantum Field Theory," in *Reality and Measurement in Algebraic Quantum Theory*, 83–95 (Springer, 2018), doi:10.1007/978-981-13-2487-1_3.

[14] K. Sanders, "The locally covariant Dirac field," *Rev. Math. Phys.* **22**, 381–430 (2010), doi:10.1142/S0129055X10003990.

[15] R. Brunetti, K. Fredenhagen, and R. Verch, "The generally covariant locality principle—a new paradigm for local quantum field theory," *Commun. Math. Phys.* **237**, 31–68 (2003), doi:10.1007/s00220-003-0815-7.

[16] C. Bär, "Green-hyperbolic operators on globally hyperbolic spacetimes," *Commun. Math. Phys.* **333**, 1585–1615 (2015), doi:10.1007/s00220-014-2097-7.

[17] C. Bär and A. Strohmaier, "Semi-bounded restrictions of Dirac type operators and the unique continuation property," *Differ. Geom. Appl.* **15**, 175–182 (2001), doi:10.1016/S0926-2245(01)00056-0.

[18] H. Araki, "On quasifree states of CAR and Bogoliubov automorphisms," *Publ. Res. Inst. Math. Sci.* **6**, 385–442 (1970), doi:10.2977/PRIMS/1195193913.

[19] C. Dappiaggi, G. Ruzzi, and E. Vasselli, "Aharonov–Bohm superselection sectors," *Lett. Math. Phys.* **110**, 3243–3278 (2020), doi:10.1007/s11005-020-01335-4.

[20] S. Attal and Y. Pautrat, "From repeated to continuous quantum interactions," *Ann. Henri Poincaré* **7**, 59–104 (2006), doi:10.1007/s00023-005-0242-8.

[21] S. Andréys, "Repeated interaction processes in the continuous-time limit, applied to quadratic fermionic systems," *Ann. Henri Poincaré* **21**, 115–154 (2020), doi:10.1007/s00023-019-00852-w.

[22] C. J. Fewster and R. Verch, "Quantum fields and local measurements," *Commun. Math. Phys.* **378**, 851–889 (2020), doi:10.1007/s00220-020-03800-6.

[23] S. Doplicher and R. Longo, "Standard and split inclusions of von Neumann algebras," *Invent. Math.* **75**, 493–536 (1984), doi:10.1007/BF01388641.

[24] C. J. Fewster, "The split property for quantum field theories in flat and curved spacetimes," *Abh. Math. Semin. Univ. Hambg.* **86**, 153–175 (2016), doi:10.1007/s12188-016-0130-9.

[25] R. Longo and F. Xu, "Von Neumann entropy in QFT," *Commun. Math. Phys.* **381**, 1031–1054 (2021), doi:10.1007/s00220-020-03702-7.

[26] B. M. Reible, "Monotonicity of the Relative Entropy and the Two-sided Bogoliubov Inequality in von Neumann Algebras," arXiv:2501.04564 (2025).

[27] M. Wirth, "The KMS and GNS Spectral Gap of Quantum Markov Semigroups," arXiv:2604.21630 (2026).

[28] L. Eriksson, "Information Recovery, Coherence Maintenance, and the Quantum–Classical Boundary — Master Block v3," 40-paper archive (2026), doi:10.5281/zenodo.18800919.

[29] L. Eriksson, "Support Leakage Makes Rapid Quantum Maintenance Singular: Soft Spectral Filters and Exponent Transmutation in Dual-Rail SSH Codes," ai.viXra:2607.0028 (2026).

[30] L. Eriksson, "Exact Stabilization by Additive GKLS Control Is Impossible at a Leaky Boundary," ai.viXra:2607.0029 (2026).

[31] L. Eriksson, "Many-Body Filling Turns Soft Spectral Leakage into a Maintenance Floor: Exact Filled-SSH Sum Rules and an Interaction-Stable Obstruction," ai.viXra:2607.0038 (2026).

[32] L. Eriksson, "Split-Regularized Recoverability in Type III AQFT: Conditional Expectations, Split-Dependent CMI, and an Audit-Friendly Recoverability Contract," ai.viXra:2601.0065 v2 (2026).