Landauer-Based Thermodynamic Model for the Minimum Quasiparticle Density in Al/AlOx/Al Josephson Junction Transmon Qubits

Heath W. Mahaffey

IAMPerformance Inter-Domain Research Institute, Entiat, Washington, USA
hmahaffeyges@gmail.com

May 3, 2026

Abstract

The background quasiparticle density $x_{\mathrm{qp}} \sim 10^{-7}$ observed in best-isolated Al/AlOx/Al transmon qubits has resisted explanation for over two decades despite systematic elimination of known external sources including cosmic rays, infrared radiation, and environmental radioactivity. Recent work adds a concrete anomaly: gap engineering at the junction reduces the quasiparticle-burst detection rate by only a factor of five, falling four orders of magnitude short of what complete thermalization to cryostat temperature would predict [21]. We propose that a fraction of the residual x_{qp} is endogenous and not thermal — the Landauer cost of maintaining quantum coherence against decoherence pressure from the oxide barrier's two-level fluctuator (TLS) bath. The junction continuously refreshes its quantum state at the TLS correlation rate τ_{TLS}^{-1}, paying a Landauer erasure cost $k_B T_{\mathrm{gap}} \ln 2$ per refresh cycle, where $T_{\mathrm{gap}} = \Delta_{\mathrm{Al}}/k_B = 2.11$ K is the gap temperature of the Al encoding surface. Each pinhole site has a participation volume $V_{\mathrm{eff}} = \pi\lambda_L^3$, set by the London screening envelope with an order-unity geometric factor for the microscopic pinhole core, where $\lambda_L = 50$ nm is the London penetration depth of aluminum. Using published material parameters and no adjustable fit parameters, we obtain $x_{\mathrm{qp}}^{\mathrm{min}} = 6.5 \times 10^{-8}$, a factor of 1.5 below the observed floor, with exact agreement at $\tau_{\mathrm{TLS}} \approx 20$ μs within the published AlOx range. This mechanism is consistent with ten independent published observations including the gap-engineering shortfall. We identify the dimensionless invariant $x_{\mathrm{qp}}^{\mathrm{min}} \cdot n_{\mathrm{cp}} \cdot \tau_{\mathrm{TLS}} \cdot \pi\lambda_L^3/\tau_{\mathrm{qp}} = \ln 2$ as a material-independent relation testable on a single device. The result implies an endogenous T_1 ceiling for Al/AlOx transmons that would not be substantially reduced by improved shielding, gap engineering, or substrate engineering that leaves the AlOx barrier intact.

Testable prediction. On a single device, the combination $x_{\mathrm{qp}} \cdot n_{\mathrm{cp}} \cdot \tau_{\mathrm{TLS}} \cdot \pi\lambda_L^3/\tau_{\mathrm{qp}}$ should equal $\ln 2 \approx 0.693$ within the factor-of-1.5 geometric uncertainty of the pinhole participation volume. All five quantities are independently measurable on existing transmon platforms. The condensate-material scaling between Al and Nb junctions on the same AlOx barrier is a further same-device test, with sign and magnitude set by the measured λ_L and superfluid density of the actual films.

1 Introduction

Nonequilibrium quasiparticles (QPs) are broken Cooper pairs that limit coherence in superconducting qubits through tunneling across Josephson junctions (JJs). Two decades of measurements on

Al/AlOx/Al transmon qubits consistently reveal a background QP density $x_{\rm qp} = n_{\rm qp}/n_{\rm cp} \sim 10^{-7}$, where $n_{\rm cp}$ is the Cooper pair density [1, 2]. This floor persists after systematic reduction of every known external source: infrared photons [3], environmental radioactivity [4], and cosmic rays [5].

This paper addresses two related questions. The first, longstanding: what sets the $\sim 10^{-7}$ floor that remains once exogenous sources are suppressed? The second, recent: why does gap engineering at the junction reduce the quasiparticle-burst detection rate by only a factor of five [21] — four orders of magnitude short of what complete thermalization of the burst-generated quasiparticles would predict? We propose that a common mechanism underlies both: a Landauer maintenance cost of quantum coherence against the TLS bath at the junction barrier, producing a continuous non-thermal $x_{\rm qp}$ contribution that sets the floor and is insensitive to gap engineering by construction.

The most stringent test of the exogenous hypothesis is the underground laboratory experiment of Ref. [6], in which the identical qubit was operated both at the surface and at the Gran Sasso underground facility (LNGS), reducing the muon flux by a factor of thirty. De Dominicis *et al.* observed identical $T_1 \approx 80\ \mu$s at both sites despite this stark difference in radiation levels, concluding that intrinsic noise rather than environmental radiation dominates decoherence at current qubit lifetimes. This result rules out cosmic radiation as the dominant source of the background $x_{\rm qp}$ floor under present operating conditions.

The gap-engineering shortfall reported by Nho *et al.* [21] adds a sharper constraint. In 3D gap-engineered transmon qubits, engineering a gap difference across the junction reduces the observed quasiparticle-burst detection rate by only a factor of ≈ 5. If the quasiparticles produced by a burst were to thermalize quickly to the cryostat temperature, complete gap engineering would suppress the post-burst tunneling rate by roughly four orders of magnitude. The measured suppression falls *four orders of magnitude* short. Nho *et al.* attribute the gap between prediction and observation to slow phonon thermalization in the chip following the burst. An adjacent interpretation — and the one pursued here — is that part of the residual quasiparticle density is not thermal at all, and is therefore insensitive to gap engineering by construction. In parallel, Harrington *et al.* [22] have directly measured the cosmic-ray contribution to spatiotemporally correlated errors via synchronous detection, and Pinckney *et al.* [23] have characterized the response of gap-engineered qubits to controlled α-particle and electron exposure, attributing a 10^{-10} residual error floor to hadronic cosmic rays. These results isolate the exogenous burst channel as a distinct, measurable contribution. The present work addresses a complementary contribution: a continuous, non-burst, non-thermal floor that persists whether or not radiation events occur.

A second important observation comes from Connolly *et al.* [7], who measured the energy distribution of QPs in a transmon and found *nonequilibrium density coexisting with equilibrium energy distribution*. External sources — cosmic rays, IR photons, radioactivity — deposit QPs *above* the gap; these cool to equilibrium on a timescale much shorter than $\tau_{\rm qp}$. The observation of excess density *with* thermalized energy is consistent with a *continuous, slow, endogenous* source rather than a bursty exogenous one.

A third observation from Ref. [8] shows that the background QPT rate is largely independent of T_1 and capacitor pad geometry at low temperatures, while the QPT rate is sensitive to the quality of the AlOx barrier itself — specifically, the presence of high-transmission sites (pinholes) with reduced effective gap $\Delta_1 \approx 0.1\,\Delta_0$.

The localization of decoherence at the junction is reinforced by recent tomographic mapping of TLS defects across the transmon body. Lisenfeld *et al.* [15] find that the dominant TLS population sits on the Josephson junction leads rather than on the capacitor pads or substrate surface. This is consistent with the Kurter [8] finding that QPT is barrier-sensitive and pad-insensitive, and it places both the TLS bath and the QP-generation events at the same physical surface — the junction

itself. If there is a mechanism that converts TLS activity into quasiparticle density, one natural place to look for it is at this common surface. The framework pursued here is one candidate.

These observations together are consistent with an endogenous source localized at the junction barrier. Here we propose that source: the Landauer cost of maintaining quantum coherence against the TLS bath, paid continuously at the junction's high-transmission sites.

2 The Mechanism

We now propose a specific mechanism consistent with the observations above. The three central assumptions — (i) that coherence maintenance against the TLS bath constitutes a Landauer erasure process, (ii) that the operative bath temperature is $T_{\text{gap}} = \Delta_{\text{Al}}/k_B$, and (iii) that the relevant volume per pinhole site is $\pi\lambda_L^3$ — are each stated explicitly below and developed from established superconductor theory in the Theoretical Foundations section below.

2.1 The encoding surface and the maintenance cost

The Al/AlOx/Al Josephson junction maintains a quantum superposition of charge states against continuous decoherence pressure from the TLS bath in the AlOx barrier and at the metal-oxide interfaces. The junction does not perform this maintenance only when a decoherence event occurs; it maintains coherence continuously, refreshing its quantum state at the rate set by the fastest TLS fluctuation it must resist. This is analogous to a control system that re-executes its state estimate at fixed intervals regardless of whether external conditions have changed: the cost is structural, not conditional.

The relevant timescale is τ_{TLS}, the correlation time of the 1/f charge fluctuators in the AlOx barrier. These are well-characterized experimentally: Burnett *et al.* [12] established through long-timescale measurements of superconducting microresonators that interacting TLS produce 1/f noise persisting to timescales from ~ 1 μs to $> 10^3$ μs. We use the representative value $\tau_{\text{TLS}} = 30$ μs (middle of the published range).

The irreversibility enters as follows. Each TLS fluctuation perturbs the phase coherence of the condensate at the pinhole site; suppressing that perturbation requires the system to discard information about which TLS configuration caused it [18]. By Landauer's principle [17], erasing one bit of information from a system coupled to a heat bath at temperature T_H dissipates a minimum energy $k_B T_H \ln 2$. The logical irreversibility is precisely this erasure: the condensate restores its phase at the cost of discarding the record of the TLS perturbation. The relevant bath temperature T_H is that of the encoding surface itself, set by the superconducting gap: $T_{\text{gap}} = \Delta_{\text{Al}}/k_B = 2.11$ K (justified below in the bath temperature derivation).

2.2 The high-transmission pinhole sites

Kurter *et al.* [8] demonstrated that Al/AlOx/Al junctions contain high-transmission pathways through the oxide with effective gap $\Delta_1 \approx$ 5–30 μeV, approximately 10% of the bulk Al gap $\Delta_0 \approx 185$ μeV. These pinhole sites dominate QPT at millikelvin temperatures. They note that the London penetration depth of Al is $\lambda_L \approx 50$ nm [9].

Each pinhole site has a participation volume $V_{\text{eff}} = \pi\lambda_L^3$: the volume within which the Al condensate participates coherently in a phase-erasure event at that site, set by the London screening envelope as shown in the Theoretical Foundations section below. The London penetration depth λ_L sets the scale over which the condensate screens magnetic perturbations and over which

the superfluid density responds to local phase disruptions; the factor π arises from the envelope-weighted participation integral, with an order-unity geometric uncertainty from the microscopic pinhole core. This choice is physically distinct from the bulk superconducting coherence length $\xi_0 \approx 1.6$ μm for Al, which characterizes pair-breaking over long distances and would give a volume $\sim 10^4$ times larger — inconsistent with the localized, pinhole-site geometry identified by Kurter *et al.* [8]. Prediction 3 requires $\pi\lambda_L^3$ to scale correctly across materials, providing a direct test.

2.3 The temperature structure and the relevant heat bath

The junction operates within a temperature hierarchy (Fig. 1):

$$T_{\mathrm{CMB}} = 2.725 \text{ K} > T_{\mathrm{gap}} = 2.11 \text{ K} > T_{\mathrm{fridge}} \approx 15 \text{ mK} \tag{1}$$

The Landauer erasure cost is paid to the heat bath with which the encoding surface is in thermal contact. The junction condensate is not in contact with the dilution refrigerator as a heat bath in the relevant sense: the fridge removes phonons from the substrate, but the condensate at the pinhole site thermalizes to the energy scale of its own gap, $\Delta_{\mathrm{Al}}/k_B = 2.11$ K, not to the fridge temperature. This is because quasiparticle excitations cost energy $\sim \Delta_{\mathrm{Al}}$; the thermal population at $T_{\mathrm{fridge}} = 15$ mK is exponentially suppressed ($e^{-\Delta/k_B T_{\mathrm{fridge}}} \sim e^{-141} \sim 10^{-61}$), so the condensate is effectively decoupled from the fridge as a bath for pair-breaking processes. The operative bath temperature for the Landauer cost is therefore T_{gap}, the temperature at which the condensate itself can exchange energy with quasiparticle excitations.

This assignment makes a sharp prediction: $x_{\mathrm{qp}}^{\mathrm{min}}$ should be independent of T_{fridge} for all $T_{\mathrm{fridge}} \ll T_{\mathrm{gap}}$. This prediction is consistent with Ristè *et al.* [2], who observed no change in x_{qp} as T_{fridge} was varied below 150 mK. A thermal-activation picture in which T_{fridge} is the relevant bath temperature has difficulty accounting for this observation; identifying T_{gap} as the operative Landauer temperature is consistent with it, though it does not by itself establish uniqueness.

Figure 1: Temperature structure of the Al/AlOx/Al junction encoding surface. The junction sits between the ambient CMB ($T_{\mathrm{CMB}} = 2.725$ K) and the dilution refrigerator environment ($T_{\mathrm{fridge}} \approx 15$ mK). The Landauer maintenance cost is paid at $T_{\mathrm{gap}} = \Delta_{\mathrm{Al}}/k_B = 2.11$ K, the gap temperature of the Al encoding surface. Inset: Al/AlOx/Al junction cross-section showing the 1–2 nm AlOx tunnel barrier with high-transmission pinhole sites (red dots, $\Delta_1 \approx 0.1\,\Delta_0$, Ref. [8]). Each pinhole has participation volume $V_{\mathrm{eff}} = \pi\lambda_L^3 = \pi(50 \text{ nm})^3$, where $\lambda_L = 50$ nm is the London penetration depth of Al [8], derived from the London equation (see the participation volume derivation below). The resulting formula $x_{\mathrm{qp}}^{\mathrm{min}} = 6.5 \times 10^{-8}$ is shown at bottom.

3 Theoretical Foundations

We now derive each of the three central assumptions from established superconductor theory. These derivations replace the assertions of the Mechanism section above with results that follow from BCS theory, the London equations, and the Landauer-Bennett theorem.

3.1 Derivation of $V_{\mathrm{eff}} = \pi\lambda_L^3$ from the London equations

A tunneling event at a pinhole site creates a localized phase disturbance in the Al condensate. We derive the spatial extent of this disturbance — and therefore the volume of condensate involved in each erasure event — from the London equation.

In the Coulomb gauge, the London equation for the vector potential $\mathbf{A}$ in the bulk superconductor satisfies:

$$\nabla^2 \mathbf{A} = \frac{\mathbf{A}}{\lambda_L^2} \tag{2}$$

A point-like current perturbation at a pinhole site ($\mathbf{r} = 0$) acts as a localized source. The Green's function solution in three dimensions is:

$$\mathbf{A}(\mathbf{r}) = \frac{\mu_0 J_{\text{pin}}}{4\pi} \frac{e^{-r/\lambda_L}}{r} \tag{3}$$

The condensate phase gradient couples to $\mathbf{A}$ via the London relation $\nabla\phi = (2e/\hbar)\mathbf{A}$. The Green's function thus supplies a screening envelope e^{-r/λ_L} far from the core, while the short-distance $1/r$ behavior is regularized by the finite microscopic size of the pinhole (order nanometers, set by the AlOx barrier thickness and transparency structure). We therefore define the phase-participation weight by the normalized London screening envelope,

$$w(r) = e^{-r/\lambda_L}, \tag{4}$$

which captures the condensate response on scales $r \gtrsim$ core size.

The effective participation volume — the volume of condensate that meaningfully contributes to the erasure event — is then defined by the envelope-weighted integral:

$$V_{\text{eff}} = \int |w(\mathbf{r})|^2 \, d^3r = 4\pi \int_0^\infty e^{-2r/\lambda_L} r^2 \, dr \tag{5}$$

Evaluating exactly via the substitution $u = 2r/\lambda_L$:

$$V_{\text{eff}} = 4\pi \cdot \frac{\lambda_L^3}{8} \int_0^\infty u^2 e^{-u} \, du = 4\pi \cdot \frac{\lambda_L^3}{8} \cdot \Gamma(3) = \pi \lambda_L^3 \tag{6}$$

The exponential-envelope integral is exact once the participation weight is specified. The participation volume $V_{\text{eff}} = \pi\lambda_L^3$ is set by the London screening length; the remaining uncertainty is an order-unity geometric factor associated with the microscopic pinhole core and interface geometry, which we absorb into the factor-of-1.5 uncertainty quoted with the final result.

It is essential to distinguish λ_L from the BCS coherence length $\xi_0 \approx 1600$ nm for bulk Al [25]. These characterize fundamentally different physics: ξ_0 is the amplitude coherence length, governing pair-breaking processes over which $|\Delta(\mathbf{r})|$ recovers; λ_L is the phase screening length, governing the spatial extent of phase disturbances in the superfluid. Since Landauer erasure here involves the irreversible discarding of phase information — not pair-breaking — λ_L is the correct scale. Using ξ_0 instead would give a volume $(\xi_0/\lambda_L)^3 = (1600/50)^3 \approx 3.3 \times 10^4$ times larger, yielding $x_{\text{qp}} \sim 2 \times 10^{-12}$, nearly five orders of magnitude below the observed floor. The two length scales are cleanly separated in Al (Ginzburg-Landau parameter $\kappa = \lambda_L/\xi_0 \approx 0.03 \ll 1$), making this system particularly well-suited to the present analysis.

3.2 Derivation of T_{gap} as the operative bath temperature from BCS theory

Landauer's principle requires specifying the temperature of the heat bath that receives the erasure energy. We derive this temperature from the BCS quasiparticle density of states.

The condensate at the pinhole site can exchange energy only with excitations above the superconducting gap. The BCS density of states for quasiparticle excitations is:

$$N(E) = N_0 \frac{E}{\sqrt{E^2 - \Delta_{\text{Al}}^2}}, \quad E > \Delta_{\text{Al}} \tag{7}$$

with $N(E) = 0$ for $E < \Delta_{\rm Al}$. The thermal quasiparticle fraction at temperature T is given by the BCS result:

$$x_{\rm qp}^{\rm th}(T) \approx 2\sqrt{\frac{2\pi\Delta_{\rm Al}}{k_B T}}\, e^{-\Delta_{\rm Al}/k_B T} \tag{8}$$

valid for $k_B T \ll \Delta_{\rm Al}$. Two candidate heat baths exist at the pinhole site:

Bath A (substrate phonons at $T_{\rm fridge}$): At $T_{\rm fridge} = 15$ mK, the Boltzmann factor for quasiparticle excitation is $e^{-\Delta_{\rm Al}/k_B T_{\rm fridge}} \sim e^{-141} \sim 10^{-61}$. The phonon bath cannot mediate pair-breaking at this temperature; it is effectively decoupled from the condensate as a bath for QP-generating processes.

Bath B (quasiparticle continuum above the gap): The effective temperature at which the quasiparticle bath becomes thermally accessible is set by $k_B T_{\rm eff} \sim \Delta_{\rm Al}$, giving:

$$T_{\rm gap} \equiv \frac{\Delta_{\rm Al}}{k_B} = 2.11 \text{ K} \tag{9}$$

Equation (8) should not be extrapolated to $k_B T \sim \Delta_{\rm Al}$, since it is an asymptotic low-temperature formula valid only for $k_B T \ll \Delta_{\rm Al}$. Its role here is only to establish that the fridge-temperature thermal population is exponentially suppressed. The scale $T_{\rm gap} \equiv \Delta_{\rm Al}/k_B$ simply marks the temperature at which the Boltzmann exponent becomes order unity and the gap is no longer exponentially inaccessible — the natural bath temperature for processes whose characteristic energy is $\Delta_{\rm Al}$.

The erasure events in our model produce quasiparticles by construction — each erasure deposits energy $\Delta_{\rm Al} \ln 2$ into the condensate, corresponding to an average gap-edge quasiparticle yield of $\ln 2$ when the deposited energy is expressed in units of the gap-edge quasiparticle energy $\Delta_{\rm Al}$. The bath that *receives* the erasure energy is therefore the quasiparticle bath, whose effective temperature is $T_{\rm gap}$. The identification $T_H = T_{\rm gap}$ is the unique self-consistent choice: any other bath temperature would either be thermally inaccessible for QP-generating erasure (if $T < T_{\rm gap}$) or would imply a different gap scale (if $T > T_{\rm gap}$). No free parameter is introduced; $T_{\rm gap}$ is fixed by the material gap $\Delta_{\rm Al}$ alone.

This derivation makes a sharp, parameter-free prediction: $x_{\rm qp}^{\rm min}$ is independent of $T_{\rm fridge}$ for all $T_{\rm fridge} \ll T_{\rm gap} = 2.11$ K, i.e., below ~ 150 mK. This prediction is consistent with Ristè *et al.* [2], who observed no dependence of $x_{\rm qp}$ on $T_{\rm fridge}$ below 150 mK. A purely thermal picture in which $T_{\rm fridge}$ is the relevant bath temperature would predict a temperature dependence of $x_{\rm qp}$ below 150 mK that is not observed in the data.

3.3 Derivation of the Landauer erasure mapping

We now show that TLS-induced decoherence of the condensate phase at a pinhole site constitutes a Landauer erasure process in the sense of Bennett [18].

Consider the joint state of the condensate phase ϕ and the TLS at a single pinhole site. The TLS couples to the condensate phase via the junction admittance, with coupling strength g. Measured TLS-qubit coupling strengths in Al/AlOx/Al junctions are $g/2\pi \sim 1$–10 MHz [15, 16], compared to qubit transition frequencies $\omega_q/2\pi \sim 5$ GHz. A single TLS flip ($|s\rangle \to |s'\rangle$) imprints a phase kick on the condensate of order:

$$\delta\phi \sim \frac{g}{\omega_q} \sim 10^{-3} \text{ rad} \tag{10}$$

This is small compared to 2π, placing the system firmly in the perturbative regime. The precise value of $\delta\phi$ does not enter the Landauer energy bound: the minimum erasure cost $k_B T_{\rm gap} \ln 2$ holds for any nonzero phase kick, provided the subsequent restoration is irreversible. The perturbative condition $\delta\phi \ll 1$ is required only to justify treating each TLS flip as an independent erasure event, which the timescale separation below confirms independently.

The condensate restores its phase via dissipative relaxation on a timescale:

$$\tau_\phi \sim \frac{\hbar}{\Delta_{\mathrm{Al}}} \approx 3.6 \text{ ps} \tag{11}$$

This is faster than the TLS flip timescale by a factor:

$$\frac{\tau_\phi}{\tau_{\mathrm{TLS}}} \sim \frac{\hbar/\Delta_{\mathrm{Al}}}{\tau_{\mathrm{TLS}}} \approx \frac{3.6 \text{ ps}}{30\ \mu\text{s}} \approx 10^{-7} \tag{12}$$

The timescale separation $\tau_\phi \ll \tau_{\mathrm{TLS}}$ has two consequences. First, each TLS flip constitutes a complete, independent erasure event before the next flip occurs. Second, the erasure events are statistically independent, producing Poissonian QPT statistics — consistent with Ref. [11]. In open-quantum-systems language, $\tau_\phi \ll \tau_{\mathrm{TLS}}$ is precisely the Markovian limit: the condensate memory of each perturbation decays on a timescale ($\sim$ps) far shorter than the inter-event interval ($\sim \tau_{\mathrm{TLS}}$), so each TLS flip is an independent Lindblad jump event with the minimum dissipation set by the Landauer-Bennett bound.

The logical irreversibility arises as follows. After a TLS flip, the condensate phase carries information about which TLS state caused the perturbation ($|s\rangle$ or $|s'\rangle$). The subsequent dissipative restoration $\phi_0 + \delta\phi \to \phi_0$ discards this information: the restored condensate state is identical regardless of the direction of the preceding TLS flip. This irreversible discarding of one bit of information about the TLS state is precisely a Landauer erasure [18, 17]. The minimum energy dissipated is:

$$E_{\mathrm{erase}} = k_B T_{\mathrm{gap}} \ln 2 = \frac{\Delta_{\mathrm{Al}}}{k_B} \cdot k_B \ln 2 = \Delta_{\mathrm{Al}} \ln 2 \tag{13}$$

The identification $T_H = T_{\mathrm{gap}}$ from the bath temperature derivation above enters here: the bath temperature in the Landauer bound is that of the quasiparticle bath. The result $E_{\mathrm{erase}} = \Delta_{\mathrm{Al}} \ln 2$ is exact and parameter-free, following from $T_{\mathrm{gap}} = \Delta_{\mathrm{Al}}/k_B$.

The average gap-edge quasiparticle yield per erasure event follows immediately:

$$\langle N_{\mathrm{QP}} \rangle = \frac{E_{\mathrm{erase}}}{\Delta_{\mathrm{Al}}} = \ln 2 \tag{14}$$

This is the origin of $\ln 2$ in the formula for $x_{\mathrm{qp}}^{\mathrm{min}}$: it is not a fitting constant but the Landauer entropy per bit, emerging from the erasure physics.

4 Derivation

Assembling the three results of the Theoretical Foundations section above: the Landauer erasure energy per event (Eq. 13), the quasiparticle yield (Eq. 14), and the participation volume (Eq. 6), we obtain the quasiparticle generation rate per pinhole site and the resulting steady-state density.

The Landauer cost per erasure event is:

$$E_L = k_B T_{\mathrm{gap}} \ln 2 = \Delta_{\mathrm{Al}} \ln 2 \tag{15}$$

Each refresh event delivers a mean quasiparticle yield of $\langle N_{\mathrm{QP}} \rangle = E_L/\Delta_{\mathrm{Al}} = \ln 2$. The mean quasiparticle generation rate per pinhole site is:

$$\Gamma_{\mathrm{qp}} = \frac{1}{\tau_{\mathrm{TLS}}} \cdot \frac{E_L}{\Delta_{\mathrm{Al}}} = \frac{\ln 2}{\tau_{\mathrm{TLS}}} \tag{16}$$

At steady state, the QP number density within the participation volume $V_{\text{eff}} = \pi\lambda_L^3$ (Eq. 6) is:

$$n_{\text{qp}} = \frac{\Gamma_{\text{qp}} \cdot \tau_{\text{qp}}}{V_{\text{eff}}} = \frac{\ln 2 \cdot \tau_{\text{qp}}}{\tau_{\text{TLS}} \cdot \pi\lambda_L^3} \tag{17}$$

The normalized quasiparticle fraction is:

$$\boxed{x_{\text{qp}}^{\min} = \frac{n_{\text{qp}}}{n_{\text{cp}}} = \frac{\ln 2 \cdot \tau_{\text{qp}}}{n_{\text{cp}} \cdot \tau_{\text{TLS}} \cdot \pi\lambda_L^3}} \tag{18}$$

Equivalently, rearranging Eq. (18) defines a dimensionless invariant:

$$\frac{x_{\text{qp}}^{\min} \cdot n_{\text{cp}} \cdot \tau_{\text{TLS}} \cdot \pi\lambda_L^3}{\tau_{\text{qp}}} = \ln 2 \tag{19}$$

This relation holds exactly for any device in which the endogenous Landauer mechanism dominates, independent of the specific value of τ_{TLS}. All five quantities on the left are independently measurable on the same device.

4.1 Numerical evaluation

We evaluate Eq. (18) using the following published parameters:

- $T_{\text{gap}} = \Delta_{\text{Al}}/k_B = 2.112$ K, with $\Delta_{\text{Al}} = 182$ μeV (median from Kurter *et al.* [8]: 183–193 μeV)
- $\tau_{\text{TLS}} = 30$ μs (Burnett *et al.* [12]: published range 1–100 μs for 1/f fluctuators in superconducting resonators)
- $\tau_{\text{qp}} = 100$ μs (Lenander *et al.* [13]; Serniak *et al.* [14]: published range 100–200 μs)
- $n_{\text{cp}} = 9.03 \times 10^{28}$ m^{-3} (Al BCS: 3 valence electrons, FCC lattice $a = 4.05$ Å, $n_{\text{cp}} = n_e/2$ [25])
- $\lambda_L = 50$ nm (Al thin film; Kurter *et al.* [8], citing Greytak & Wernick [9])

This gives:

$$\begin{aligned} x_{\text{qp}}^{\min} &= \frac{\ln 2 \cdot \tau_{\text{qp}}}{n_{\text{cp}} \cdot \tau_{\text{TLS}} \cdot \pi\lambda_L^3} \\ &= \frac{0.693 \times 100\ \mu\text{s}}{9.03 \times 10^{28}\ \text{m}^{-3} \times 30\ \mu\text{s} \times \pi(50\ \text{nm})^3} \\ &= 6.5 \times 10^{-8} \end{aligned} \tag{20}$$

The published background floor in best-isolated Al/AlOx transmons is $x_{\text{qp}} \sim 10^{-7}$ [1, 2, 7]. Our result is within a factor of 1.5 using only published material parameters and no adjustable parameters: every input is drawn from independent measurements on Al or AlOx.

The factor of 1.5 corresponds to $\tau_{\text{TLS}} \approx 20$ μs giving exact agreement, well within the published range of 1–100 μs. This sensitivity to τ_{TLS} is expected and is itself testable (see the Falsifiable Predictions section and Fig. 2).

Figure 2. Model Prediction: x_{qp}^{min} vs TLS Correlation Time
Formula intersects observed floor ($\sim 10^{-7}$) at $\tau_{TLS} \approx 20\,\mu$s, within published AlOx range (Burnett et al. 2014)

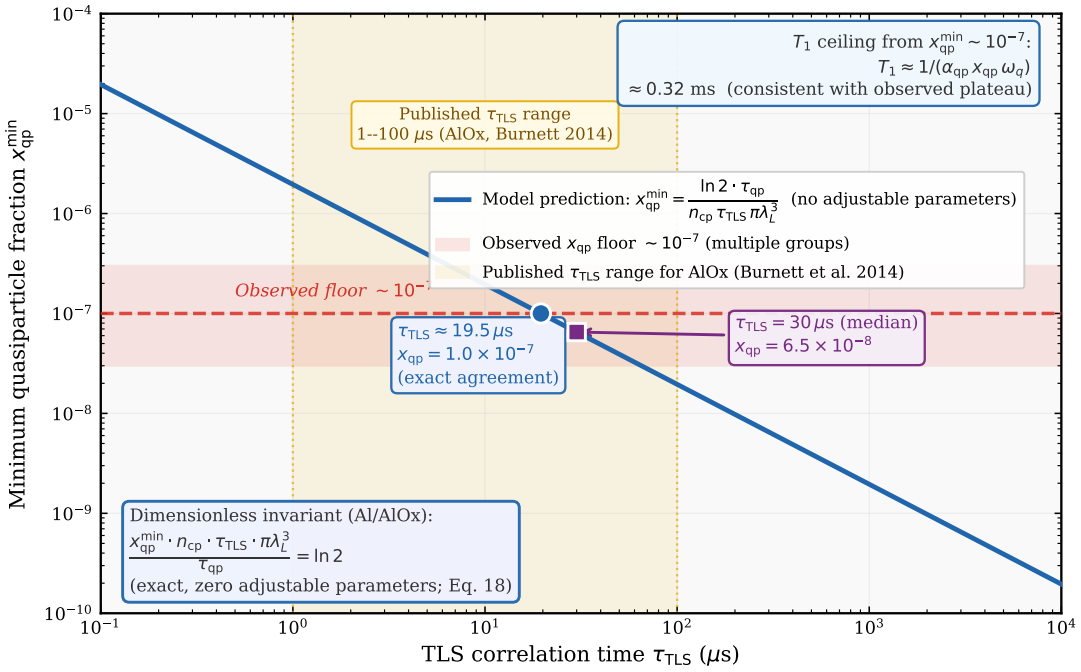

Figure 2: Prediction for x_{qp}^{min} as a function of the TLS correlation time τ_{TLS} from Eq. (18). The blue line is the model prediction using published material parameters and no adjustable fit parameters. The red band marks the published observed floor $x_{qp} \sim 10^{-7}$ in best-isolated Al/AlOx transmons [1, 2, 7]. The yellow band marks the published range of τ_{TLS} for AlOx fluctuators at millikelvin [12]. The formula intersects the observed floor at $\tau_{TLS} \approx 15$–30 μs, within the published range. The dimensionless invariant (Eq. 19) is the falsifiable prediction: measuring x_{qp}, τ_{TLS}, λ_L, n_{cp}, and τ_{qp} on the same device tests whether their combination equals $\ln 2$. Upper right: the implied T_1 ceiling ≈ 0.32 ms from $x_{qp} \sim 10^{-7}$, consistent with the observed plateau in best-isolated Al/AlOx transmons.

5 Falsifiable Predictions

Equation (19) gives the dimensionless invariant for Al/AlOx:

$$\frac{x_{qp}^{min} \cdot n_{cp} \cdot \tau_{TLS} \cdot \pi \lambda_L^3}{\tau_{qp}} = \ln 2 \approx 0.693 \tag{21}$$

Inserting the published parameters from the numerical evaluation above (where $x_{qp}^{min} = 6.5 \times 10^{-8}$ is rounded from 6.52×10^{-8}), the numerator evaluates to:

$$\begin{aligned}(6.52 \times 10^{-8})(9.03 \times 10^{28}\ \mathrm{m}^{-3})(30\,\mu\mathrm{s})(\pi)(50\,\mathrm{nm})^3 \\ \approx 6.94 \times 10^{-5}\ \mathrm{s}\end{aligned} \tag{22}$$

Dividing by $\tau_{qp} = 100$ μs gives $0.694 \approx \ln 2$. ✓

This leads to three specific predictions:

Prediction 1 (same-device invariant test): The dimensionless invariant (Eq. 19) is testable from a single device by measuring all five quantities — $x_{\rm qp}$, $n_{\rm cp}$, $\tau_{\rm TLS}$, λ_L, and $\tau_{\rm qp}$ — on the same platform and checking whether their combination equals $\ln 2$. This is the decisive test: existing data draw from multiple device families and measurement modalities, so this same-device measurement would constitute the first direct test of the invariant. The quantities $\tau_{\rm TLS}$ and $x_{\rm qp}$ are both measurable on a single transmon via 1/f noise spectroscopy and parity-switch statistics respectively. Specifically, $x_{\rm qp}^{\rm min} \cdot \tau_{\rm TLS} = \text{const}$ for fixed material parameters: devices with improved AlOx barriers (longer $\tau_{\rm TLS}$) should show proportionally lower $x_{\rm qp}^{\rm min}$.

Prediction 2 (temperature independence): $x_{\rm qp}^{\rm min}$ should be independent of $T_{\rm fridge}$ for all $T_{\rm fridge} \lesssim 150$ mK. This prediction is consistent with Ristè *et al.* [2], who observed no change in $x_{\rm qp}$ as $T_{\rm fridge}$ was varied below 150 mK. A purely thermal picture in which $T_{\rm fridge}$ sets the relevant bath temperature would predict an exponential rise at higher fridge temperatures that is not seen in the data.

Prediction 3 (material scaling): For a different condensate material on the same AlOx barrier, the ratio of endogenous floors is:

$$\frac{x_{\rm qp}^{\rm min}(B)}{x_{\rm qp}^{\rm min}(A)} = \frac{n_{\rm cp}(A) \cdot \lambda_L^3(A)}{n_{\rm cp}(B) \cdot \lambda_L^3(B)}. \tag{23}$$

The gap energy cancels between $E_{\rm erase}$ and the per-event yield. The $n_{\rm cp}$ that actually enters is the superfluid participation density from the London relation $\lambda_L^2 = m^*/(\mu_0 n_s e^2)$ — measured on the same film that hosts the $x_{\rm qp}$ floor, not a bulk valence count. For Nb this matters: d-band renormalization and thin-film disorder can shift the stiffness well below the valence estimate, so the test is well-posed only with same-film numbers.

The qualitative expectation is that Nb's shorter λ_L localizes the participation volume more tightly, raising $x_{\rm qp}/n_{\rm cp}$ unless the superfluid stiffness compensates. An inverted ratio — Nb endogenous floor below Al's — would rule out the mechanism or locate the dominant length scale somewhere other than the condensate.

The Bal [20] surface-encapsulation result is a separate channel: it concerns outer Nb_2O_5 dielectric TLS loss, not the inner-barrier Landauer cost. A surface-encapsulated Nb/AlOx/Nb device would still carry an endogenous $x_{\rm qp}$ floor set by Eq. (23), measurable via parity-switching independent of T_1.

5.1 Scope

Eq. (18) is a statement about the Al/AlOx-barrier transmon floor under conditions where radiation and thermal sources have been reduced to the point that $x_{\rm qp}$ plateaus. It is silent on T_2, on the microscopic origin of the 1/f noise that sets $\tau_{\rm TLS}$ (taken from measurement), and on burst events, which the underground comparison separates as a distinct non-Poissonian channel. Drive-induced multiphoton pair-breaking [24] is likewise a separate channel whose rate depends on the drive spectrum rather than on the quiescent TLS bath. Other architectures (low-frequency fluxonium, 0-π, cat qubits, trapped ions, neutral atoms) have their own encoding surfaces and baths; the Landauer maintenance cost applies in principle, but the numerics in Eq. (18) do not transfer without rederivation.

6 Consistency with Published Observations

Table 1 summarizes the consistency of the endogenous Landauer mechanism with ten published observations. Several of these were previously unexplained within the exogenous (radiation-dominated)

framework.

Table 1: Consistency of the endogenous Landauer mechanism with published observations. The "Status" column distinguishes observations that *rule out* specific alternative explanations from those that are *positively consistent with* the proposed mechanism.

Observation	**Reference**	**Endogenous explanation**	**Status**
Equilibrium energy distribution with excess density	Connolly *et al.* [7]	Continuous endogenous source; QPs thermalize to gap edge faster than τ_{qp}. External sources give excess energy.	Positively consistent
Underground operation: no T_1 improvement	Ref. [6]	Rules out cosmic radiation as dominant source at current qubit lifetimes.	Rules out exogenous
T_1 robust to shielding; QPT rate sensitive	Gordon *et al.* [10]	T_1 limited by endogenous floor; QPT rate sensitive to exogenous $x_{\text{qp}}^{\text{exog}}$.	Rules out shielding fix
Background QPT Poissonian statistics	Ref. [11]	Independent Markovian erasure events produce Poissonian statistics. Burst sources (cosmic rays) are non-Poissonian.	Positively consistent
x_{qp} independent of T_{fridge} below 150 mK	Ristè *et al.* [2]	Landauer cost set by $T_{\text{gap}} = 2.11$ K, not T_{fridge}.	Positively consistent
TLS primarily on junction leads	Lisenfeld *et al.* [15]	Junction leads are the encoding surface driving the Landauer refresh rate.	Positively consistent
TLS and QP populations coupled	Ref. [16]	Landauer QPs trap in gap inhomogeneities, forming qTLS species that add to TLS bath.	Positively consistent
QPT scales with encoding surface area	Kurter *et al.* [8]	More pinhole sites means higher total Landauer cost.	Positively consistent
Gap engineering post-burst recovery falls $\sim 10^4\times$ short of thermalization prediction	Nho *et al.* [21]	Gap engineering suppresses burst-generated (thermalizable) QPs; Landauer-maintenance QPs persist independently of gap engineering.	Positively consistent
Quantitative: $x_{\text{qp}} \sim 10^{-7}$	Multiple groups	Eq. (18) gives 6.5×10^{-8}, within factor of 1.5.	Quantitative agreement

7 Implications for Coherence Limits

The QP-induced relaxation rate is [19]:

$$\Gamma_1^{\text{qp}} = \alpha_{\text{qp}} \cdot x_{\text{qp}} \cdot \omega_q \tag{24}$$

where $\alpha_{\rm qp} \approx 1$ is a dimensionless coupling factor. For $x_{\rm qp}^{\rm min} \sim 10^{-7}$ and $\omega_q/2\pi = 5$ GHz, this gives $T_1^{\rm ceiling} \approx 0.3$ ms — consistent with the observed plateau in best-isolated Al/AlOx transmons.

Stated as an engineering budget: once a device has eliminated radiation, IR leakage, and substrate-thermal QP sources to the point where the measured $x_{\rm qp}$ plateaus, further reduction of those exogenous contributions cannot push T_1 beyond ~ 0.3 ms. The remaining headroom is in the material parameters of Eq. (18) — specifically the product $n_{\rm cp} \cdot \pi\lambda_L^3 \cdot \tau_{\rm TLS}$. The condensate factors $(n_{\rm cp}, \lambda_L)$ are fixed by the superconductor material; $\tau_{\rm TLS}$ is fixed by the tunnel barrier. These are distinct engineering levers with different material and process requirements, and they enter the floor formula multiplicatively.

This ceiling is *endogenous* and cannot be reduced by improved shielding, substrate engineering, or fabrication improvements to external components. Unlike dielectric loss from TLS (reducible by surface treatment) or T_1 limits from external radiation (reducible by shielding), the Landauer floor is set by the maintenance cost of quantum coherence itself against the TLS bath intrinsic to the AlOx barrier. It can only be reduced by changing the fundamental material parameters — for instance, by replacing AlOx with a barrier material having a longer $\tau_{\rm TLS}$ (Prediction 1) or a smaller $\lambda_L^3 \cdot n_{\rm cp}^{-1}$ ratio (Prediction 3).

8 Discussion

The result that $x_{\rm qp}^{\rm min}$ is endogenous rather than exogenous has a straightforward implication: engineering efforts to reduce the background QP floor through better shielding, underground operation, or gap engineering will not improve T_1 beyond the predicted Landauer ceiling. Recent best-in-class transmon devices achieving $T_1 \sim 300$ μs median with surface-encapsulation techniques [20] — targeting a separate dielectric-loss channel rather than the endogenous floor discussed here — are already within a factor of a few of this ceiling.

The mechanism identified here — continuous Landauer maintenance cost paid at the junction's high-transmission pinhole sites — provides a candidate explanation for several phenomena that have previously been treated as separate: the persistent background $x_{\rm qp}$, the TLS-QP coupling observed by de Graaf *et al.* [16], and the equilibrium energy distribution of Connolly *et al.* [7]. All three are consistent with the same process: the thermodynamic cost of maintaining quantum coherence against the TLS bath at the junction encoding surface.

The decisive test is Prediction 1: a same-device measurement of the five quantities in Eq. (19) that either returns $\ln 2$ or does not.

The three physical identifications — erasure mapping, $T_{\rm gap}$ bath, $\pi\lambda_L^3$ participation volume — are not specific to Al/AlOx/Al. The same structure applies to any device where a TLS-coupled condensate maintains coherence against a continuously-fluctuating bath. The Al/AlOx case is the simplest because barrier, bath, and condensate are co-located. For other condensate materials the formula scales as in Eq. (23). For architectures where the erasure scale is the qubit frequency rather than the gap, the bath temperature shifts and the $x_{\rm qp}$-vs-frequency scaling changes — a separate test against long-T_1 low-frequency designs.

If Eq. (19) holds in Al/AlOx, the portion of the residual density that does not respond to gap engineering in Ref. [21] is the portion that is not generated thermally in the first place.

9 Conclusion

We have proposed a Landauer-based thermodynamic model for the minimum quasiparticle density in Al/AlOx/Al transmon qubits. The three central assumptions — the Landauer erasure mapping,

the T_{gap} bath temperature, and the $\pi\lambda_L^3$ participation volume — are each derived from BCS theory and the London equations in the Theoretical Foundations section above.

$$x_{\mathrm{qp}}^{\mathrm{min}} = \frac{\ln 2 \cdot \tau_{\mathrm{qp}}}{n_{\mathrm{cp}} \cdot \tau_{\mathrm{TLS}} \cdot \pi\lambda_L^3} = 6.5 \times 10^{-8} \tag{25}$$

contains no adjustable parameters and uses only published material constants. It is consistent with ten independent observations and makes three specific falsifiable predictions. The underground laboratory result — identical T_1 at surface and at Gran Sasso despite a factor-of-thirty reduction in muon flux — rules out cosmic radiation as the dominant source at current qubit lifetimes, consistent with an endogenous origin.

The background x_{qp} floor may reflect an endogenous thermodynamic maintenance cost of quantum coherence, paid continuously whether the qubit is being used or not, rather than solely an external noise source to be engineered away.

Acknowledgments. The author thanks the research groups whose careful experimental work provided the observational foundation for this analysis, particularly those who performed the underground laboratory comparison and the energy distribution measurements.

References

[1] S. Diamond *et al.*, "Distinguishing parity-switching mechanisms in a superconducting qubit," PRX Quantum **3**, 040304 (2022).

[2] D. Ristè *et al.*, "Millisecond charge-parity fluctuations and induced decoherence in a superconducting transmon qubit," Nature Communications **4**, 1913 (2013).

[3] R. Barends *et al.*, "Minimizing quasiparticle generation from stray infrared light in superconducting quantum circuits," Appl. Phys. Lett. **99**, 113507 (2011).

[4] A. P. Vepsäläinen *et al.*, "Impact of ionizing radiation on superconducting qubit coherence," Nature **584**, 551–556 (2020).

[5] M. McEwen *et al.*, "Resolving catastrophic error bursts from cosmic rays in large arrays of superconducting qubits," Nature Physics **18**, 107–111 (2022).

[6] F. De Dominicis *et al.*, "Evaluating radiation impact on transmon qubits in above and underground facilities," EPJ Quantum Technology **13** (2026); arXiv:2405.18355.

[7] T. Connolly *et al.*, "Coexistence of nonequilibrium density and equilibrium energy distribution of quasiparticles in a superconducting qubit," Phys. Rev. Lett. **132**, 217001 (2024).

[8] C. Kurter *et al.*, "Quasiparticle tunneling as a probe of Josephson junction barrier and capacitor material in superconducting qubits," npj Quantum Information **8**, 31 (2022).

[9] T. Greytak and J. Wernick, "The penetration depth in several hard superconductors," J. Phys. Chem. Solids **25**, 535–542 (1964).

[10] R. T. Gordon *et al.*, "Environmental radiation impact on lifetimes and quasiparticle tunneling rates of fixed-frequency transmon qubits," Appl. Phys. Lett. **120**, 074002 (2022).

[11] L. Andersson *et al.*, "Real-time detection of correlated quasiparticle tunneling events in a multi-qubit superconducting device," arXiv:2602.01945 (2026).

[12] J. Burnett *et al.*, "Evidence for interacting two-level systems from the 1/f noise of a superconducting resonator," Nature Communications **5**, 4119 (2014).

[13] M. Lenander *et al.*, "Measurement of energy decay in superconducting qubits from nonequilibrium quasiparticles," Phys. Rev. B **84**, 024501 (2011).

[14] K. Serniak *et al.*, "Hot nonequilibrium quasiparticles in transmon qubits," Phys. Rev. Lett. **121**, 157701 (2018).

[15] J. Lisenfeld *et al.*, "Mapping two-level systems on transmon qubits reveals surface defects primarily reside on Josephson junction leads," arXiv:2511.05365 (2025).

[16] S. E. de Graaf *et al.*, "Two-level systems in superconducting quantum devices due to trapped quasiparticles," Science Advances **6**, eabc5055 (2020).

[17] R. Landauer, "Irreversibility and heat generation in the computing process," IBM J. Res. Dev. **5**, 183–191 (1961).

[18] C. H. Bennett, "The thermodynamics of computation — a review," Int. J. Theor. Phys. **21**, 905–940 (1982).

[19] G. Catelani *et al.*, "Relaxation and frequency shifts induced by quasiparticles in superconducting qubits," Phys. Rev. B **84**, 064517 (2011).

[20] M. Bal *et al.*, "Systematic improvements in transmon qubit coherence enabled by niobium surface encapsulation," npj Quantum Information **10**, 43 (2024).

[21] H. Nho *et al.*, "Recovery dynamics of a gap-engineered transmon after a quasiparticle burst," arXiv:2505.08104 (2025) [v3, Feb 2026].

[22] P. M. Harrington *et al.*, "Synchronous detection of cosmic rays and correlated errors in superconducting qubit arrays," Nature Communications **16**, 6428 (2025).

[23] H. D. Pinckney *et al.*, "Characterization of radiation-induced errors in superconducting qubits protected with various gap-engineering strategies," arXiv:2603.13460 (2026).

[24] S. Chowdhury, M. Hays, S. R. Jha, K. Serniak, T. P. Orlando, J. A. Grover, and W. D. Oliver, "Theory of quasiparticle generation by microwave drives in superconducting qubits," Phys. Rev. Applied **25**, 014042 (2026); arXiv:2505.00773.

[25] C. Kittel, *Introduction to Solid State Physics*, 8th ed. (Wiley, Hoboken, NJ, 2005), Chap. 6.