
Calorimetric Signature of Quantum Measurement: A Record-Formation Heat Bound and Differential Microcalorimetry Test

[Moses Rahnama](#)*

Posted Date: 5 March 2026

doi: 10.20944/preprints202602.1739.v2

Keywords: quantum measurement; Landauer principle; thermodynamics of information; objective classical record; record formation; mutual information; decoherence; quantum Darwinism; circuit QED; superconducting qubits; differential microcalorimetry; nanocalorimetry



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

Calorimetric Signature of Quantum Measurement: A Record-Formation Heat Bound and Differential Microcalorimetry Test

Moses Rahnama

Mina Analytics, New York, USA; mooses@minaanalytics.com

Abstract

We propose that quantum measurement is a *boundary event*: a physically identifiable, irreversible transition in which a reversible system/pointer correlation is forced across an operational irreversibility threshold into objective classical record stabilization. We formulate a three-stage taxonomy separating reversible premeasurement (Stage 1), irreversible record formation (Stage 2, the boundary event), and memory reset (Stage 3), and identify the stage at which a Landauer-scale heat bound applies. Under explicit operational conditions (C1 to C6) in the uncontrolled-decoherence regime, the record-formation channel must dissipate at least $k_B T \ln 2$ of heat per bit of classical mutual information $I(X; Y) = I(X; Y)$. We propose a circuit-QED differential microcalorimetry experiment with matched ON/OFF branches that share identical premeasurement pulses and routing losses, differing only in whether an objective record is stabilized. The measurand is the differential deposited energy $\Delta Q \equiv Q_{\text{ON}} - Q_{\text{OFF}}$, which isolates record-formation dissipation from common-mode backgrounds. The primary deep-quantum demonstration targets the *temporal coincidence* of heat onset and reversibility loss via a reversal-delay sweep (Control 3), providing a distinctive boundary-event signature even when $\Delta Q \gg k_B T \ln 2$. Near-floor residual tests ($r \equiv \Delta Q - k_B T \ln 2 \cdot I(X; Y)$) require lower-energy pointer implementations or elevated operating temperatures and are presented as a concrete roadmap. The bound is falsified if r is negative at high statistical significance under verified conditions.

Keywords: quantum measurement; Landauer principle; thermodynamics of information; objective classical record; record formation; mutual information; decoherence; quantum Darwinism; circuit QED; superconducting qubits; differential microcalorimetry; nanocalorimetry

1. Introduction

The quantum measurement problem concerns how definite outcomes emerge from quantum superpositions. The standard formalism treats measurement as instantaneous and thermodynamically neutral; we instead treat it as a physical process with temporal structure and a dissipative signature. We propose an operational thermodynamic diagnostic: the formation of an *objective classical record* is a physical event (a **boundary event**) that converts quantum uncertainty into classical certainty with a thermodynamic resource cost that becomes a dissipative floor under explicit operational conditions.

Definition (Measurement as a Boundary Event). A measurement is a *boundary event*: the first time in a protocol run when the system/pointer correlation is forced across an operational irreversibility boundary into an uncontrolled decoherence channel that (i) creates a stable classical register Y (objective record) and (ii) prevents practical reversal without microscopic control of bath degrees of freedom. The boundary event is certified by loss of reversibility ($\mathcal{F} < \mathcal{F}_*$), not by the calorimetric signal; the heat bound is a predicted consequence. The ON/OFF architecture moves the experiment across (ON) or keeps it before (OFF) this boundary while holding Stage 1 control work common-mode.

This cost is not merely conceptual. Landauer's principle establishes that any logically irreversible operation erasing one bit of information must dissipate at least $k_B T \ln 2$ of heat. [1] We hypothesize that irreversible record formation in quantum measurement is such an operation: stabilizing a classical record irreversibly distinguishes quantum alternatives into a stable classical register and exports entropy to the environment. This paper formulates the hypothesis as a conditional, falsifiable bound and designs an experiment to test it.

However, applying Landauer's principle to quantum measurement has generated controversy. Critics argue that: (i) the principle applies to erasure, not measurement; [3] (ii) measurement can be performed reversibly with no heat cost; [2] and (iii) any "Landauer cost" can be deferred indefinitely to memory reset. [2]

We resolve these tensions by providing a precise taxonomy of measurement stages and identifying exactly where the Landauer bound applies. Our claim is not that "measurement" in general costs $k_B T \ln 2$ per bit. Rather, the bound is a conditional statement about *record formation* in a thermal environment with a classical register and a cyclic apparatus.

Recent theoretical work supports this connection and provides substantial precedent for thermodynamic stage decompositions. Latune and Elouard independently developed a three-stage decomposition of measurement (premeasurement, objectification, reset) with quantitative thermodynamic bounds for each stage, [8] showing that measurements *can* approach reversibility when no stable record is created. Mohammady and Buscemi identified a "thermodynamic trilemma" constraining efficient measurements. [9] Touil *et al.* formalized objective record formation through quantum Darwinism, where classical objectivity emerges via redundant environmental encoding. [7]

Our contribution operationalizes and extends prior stage decompositions (in particular the Latune-Elouard framework [8]) into a falsifiable experimental package:

1. A three-stage taxonomy cleanly separating reversible correlation, record formation (the boundary event), and reset, with six explicit operational conditions (C1 to C6) specifying when the Landauer bound applies and when it does not
2. An experimentally testable specialization of information-thermodynamic second-law inequalities to the record-stabilization channel, yielding a **record-formation heat bound** expressed in terms of measurable classical mutual information $I(X; Y)$
3. The first, to our knowledge, matched ON/OFF differential microcalorimetry experiment designed to isolate record-formation heat, with quantitative sensitivity analysis and a complete calibration/systematics program




Figure 1. Predicted differential heat ΔQ (solid blue, left axis) and Landauer bound $k_B T \ln 2 \cdot I(X; Y)$ (dashed green, left axis) versus pointer photon number \bar{n} , at $T = 10$ mK with uniform prior $p(0) = p(1) = 1/2$. The shaded region shows the positive residual $\Delta Q - k_B T \ln 2 \cdot I(X; Y) > 0$ at all tested measurement strengths. The dotted red curve (right axis) shows the classical mutual information $I(X; Y)$ rising from 0 to 1 bit as \bar{n} increases. In the deep-quantum regime ($\bar{n} \sim 1$), $\Delta Q \sim h\nu$ exceeds the bound by $\sim 48\times$; near-saturation tests require pointer energies closer to $k_B T \ln 2$.

2. Three-Stage Taxonomy and Operational Definitions

2.1. Stages: Premeasurement, Record Formation, Reset

We decompose quantum measurement into three analytically distinct stages. This is a conceptual framework for identifying where the Landauer cost is paid; the measurement event itself is Stage 2.

1. **Stage 1: Reversible premeasurement (unitary correlation).** The system S becomes correlated/entangled with a pointer P via a unitary U_{SP} . No entropy production occurs. No Landauer cost. The pointer states $|P_i\rangle$ are distinguishable but the process can be unitarily reversed.
Thermodynamic status: Thermodynamically reversible in principle; no Landauer cost in the ideal limit. (Real hardware produces residual dissipation from control pulses, but this is common-mode and cancels in the differential measurement.)
Information status: Quantum correlations between S and P increase (the reduced states become mixed), but no classical record Y exists yet. The classical mutual information $I(X; Y)$ is not defined at this stage because no decohered outcome register has been created.
2. **Stage 2: Boundary Event, irreversible record formation (objective stabilization).** The pointer couples to a thermal bath, creating a macroscopic record via redundant environmental encoding ($N \gg 1$ environmental degrees of freedom). The record is stabilized against reversal. This stage is logically irreversible and is where Landauer-scale heat is paid.
Thermodynamic status: Entropy flows to the bath. The bound applies here when conditions C1 to C6 hold.
Information status: Classical mutual information $I(X; Y)$ is created. The outcome Y becomes a definite classical variable.
3. **Stage 3: Reset (optional reusability).** Erasing the record to reuse the apparatus. Standard Landauer erasure cost. Reset can occur long after the measurement; it is conceptually distinct from record formation.
Thermodynamic status: Entropy flows to the bath during reset.
Information status: Classical information is destroyed.

This taxonomy resolves apparent contradictions: “reversible measurements” address Stage 1 (and possibly engineered avoidance of Stage 2), while Landauer bounds constrain Stage 2 when an objective record is actually created. The “cost can be deferred to reset” objection addresses Stage 3 rather than Stage 2.

2.2. Operational Definition: Objective Record Formation

Let X be a classical random variable labeling the prepared input and Y be the classical outcome stored in a record register. We call Y an **objective record** if, operationally:

1. **Stability:** Y persists for time $t \gg \tau_c$, where τ_c is the reversibility time beyond which a coherent reversal operation fails to recover the initial state fidelity (defined precisely in Section 12).
2. **Classical definiteness:** The record is well-approximated by a classical register with negligible coherence between distinct outcomes in the record basis.
3. **Redundancy (strengthening criterion):** Following quantum Darwinism, [6,7] the outcome information is redundantly encoded in environmental degrees of freedom. Quantitatively, the accessible information I_{acc}^F between the system variable X and an environmental fragment F satisfies

$$\frac{I_{\text{acc}}^F}{H(X)} \geq 1 - \delta \quad \text{for } f \leq f_0, \quad (1)$$

where I_{acc}^F is the accessible information between X and fragment F , $H(X)$ is the Shannon entropy of the prepared input, $f = |F|/|E|$ is the fragment fraction (ratio of fragment size to total environment), f_0 is the smallest fragment fraction achieving near-complete information, and δ is a small tolerance (e.g., $\delta \sim 0.1$), giving redundancy $R \approx 1/f_0$. (This criterion provides theoretical grounding from quantum Darwinism; [6,7] the proposed experiment does not directly measure

the fragment-information curve or certify Darwinist redundancy. Instead, we use the reversibility witness \mathcal{F} and record persistence $\tau > \tau_c$ as operational proxies for an objective-record-compatible regime, and test the predicted thermodynamic consequence. Fragment-information probing is a natural future-work extension.)

Separating definition from prediction. Record formation (the boundary event) is defined *independently* of the calorimetric signal: it is certified by the failure of the OFF-branch reversibility witness ($\mathcal{F} < \mathcal{F}_*$) and the persistence of the record beyond the reversibility time τ_c (Section 12). The heat bound $\langle Q_{\text{rec}} \rangle \geq k_B T \ln 2 \cdot I(X; Y)$ is the *predicted thermodynamic consequence* of that boundary crossing under conditions C1 to C6, not part of its definition. The experiment tests this prediction: it asks whether the onset of irreversible heat in the ON/OFF differential coincides with, and exceeds the Landauer floor for, the independently certified boundary event. This separation prevents circularity: a system that crosses the irreversibility boundary without producing the predicted heat would *falsify* the bound, not be reclassified as “not having formed a record.”

Because the fragment-information (redundancy) plateau is not directly measured in the proposed experiment, we do not claim to certify Darwinist redundancy itself. Rather, an *objective-record-compatible* regime is inferred operationally when the record is stable ($\tau > \tau_c$) and the reversibility witness confirms that reversal fails, with null protocols excluding reversible premeasurement-only explanations.

3. Record-Formation Heat Bound

3.1. General Inequality

We model record formation as a completely positive trace-preserving (CPTP) map \mathcal{E} that (i) produces a classical register Y with statistics $p(y|x)$ and (ii) couples irreversibly to a thermal bath at temperature T . From the generalized second law for information processing, [5] the entropy production Σ satisfies

$$\Sigma = \beta \langle Q_{\text{rec}} \rangle + \Delta S_{\text{sys}} + \Delta S_{\text{mem}} - I_{S:M}^{\text{gain}} \geq 0, \quad (2)$$

where ΔS_{sys} is the entropy change of the measured system S , ΔS_{mem} is the entropy change of *reusable* apparatus degrees of freedom kept inside the Stage 2 system boundary (pointer/control modes reset between trials), and $I_{S:M}^{\text{gain}}$ is the mutual information gained between S and the memory M . In this bookkeeping, the stabilized macroscopic record in the absorber/environment is treated on the bath side of the channel, so its irreversible entropy production contributes to $\beta \langle Q_{\text{rec}} \rangle$ rather than to ΔS_{mem} . In the experiment, the measured classical quantity $I(X; Y) = I(X; Y)$ is used as a conservative lower bound on this gain. We take an initially thermal bath so that the bath-entropy term is $\beta \langle Q_{\text{rec}} \rangle$. Throughout, entropy terms (ΔS and $I_{S:M}^{\text{gain}}$) are dimensionless (nats, i.e., in units of k_B), and information quantities ($I(X; Y)$) are in bits. The factor $\ln 2$ converts between them: 1 bit = $\ln 2$ nats. In natural units,

$$\beta \langle Q_{\text{rec}} \rangle \geq (\ln 2) I(X; Y) - \Delta S_{\text{sys}} - \Delta S_{\text{mem}}, \quad (3)$$

with $\beta = 1/(k_B T)$. This is a channel-level inequality for the Stage 2 map under the stated bookkeeping boundary, not a universal standalone theorem for arbitrary measurement implementations.

3.2. Conditional Bound (Corollary)

If, during Stage 2, a classical record is created via irreversible entropy export to the bath (redundancy generation), and if the reusable apparatus degrees included in ΔS_{mem} are cycled back to a reference state over each ON/OFF trial (so net $\Delta S_{\text{mem}} \approx 0$), with $\Delta S_{\text{sys}} \approx 0$ and $W \approx 0$ for the record-formation channel, then: (These are *design constraints* on the protocol, not universal physical facts. The condition $\Delta S_{\text{sys}} \approx 0$ holds because Stage 1 premeasurement has already brought the system's reduced state to near-maximal entropy; the Lindblad simulation confirms $S_q \approx 0.99$ bit throughout

Stage 2 (Section 5.5). If either entropy term is non-negligible, the full bound including ΔS_{sys} , ΔS_{mem} , and W [Equation below] must be used.)

$$\langle Q_{\text{rec}} \rangle \geq k_{\text{B}} T \ln 2 \cdot I(X; Y) \quad (4)$$

where

$$\begin{aligned} I(X; Y) &= H(Y) - H(Y|X) \\ &= \sum_{x,y} p(x) p(y|x) \log_2 \frac{p(y|x)}{\sum_{x'} p(x') p(y|x')} \end{aligned} \quad (5)$$

For an ideal binary projective measurement with uniform prior $p(X=0) = p(X=1) = 1/2$, this gives $I(X; Y) = 1$ bit. Here Y denotes the classical record variable created in Stage 2: the stabilized classical information encoded in the $N \gg 1$ environmental degrees of freedom of the absorber/bath node via irreversible pointer thermalization. Because the pointer states in dispersive readout ($|\alpha\rangle$ vs. $|\alpha\rangle$) carry identical mean photon number, the record is not an energy-discriminating absorber macrostate; rather, it resides in the microstate correlations of the bath, precisely the redundant environmental encoding described by quantum Darwinism (Section 2.2). The absorber serves as the *monitored bath node* whose total entropy export is calorimetrically measured; it need not resolve individual outcomes shot-by-shot for the channel-level bound to apply.

In production calorimetry we do not perform shot-resolved outcome readout of Y ; instead $I(X; Y)$ is obtained from *separate* dispersive calibration blocks (not interleaved with calorimetric acquisition) using identical premeasurement pulses to infer $p(x)$ and $p(y|x)$. We denote this calibration-derived quantity I_{cal} and use it as the operational estimator of $I(X; Y)$ entering the bound. The identification $I_{\text{cal}} \approx I(X; Y)$ rests on three explicit assumptions: (i) the dispersive calibration and the absorber thermalization both respond to the same pointer displacement $|\alpha_{\sigma}\rangle$ (same physical degree of freedom, different readout method); (ii) the absorber, which captures the pointer field with near-unit efficiency (capture efficiency $> 99.8\%$), retains at least as much information about X in its environmental microstate as the noise-limited dispersive chain extracts; and (iii) acceptance criteria 1 to 4 below verify that calibration and calorimetric operating points are matched within stated uncertainty. The tested inequality $\langle Q_{\text{rec}} \rangle \geq k_{\text{B}} T \ln 2 \cdot I_{\text{cal}}$ is therefore conservative: the manuscript's general bound (Section 3) is written in terms of the information gain $I_{S:M}^{\text{gain}} \geq I(X; Y) \geq I_{\text{cal}}$, so any shortfall in I_{cal} relative to the true environmental record information only strengthens the inequality. This substitution is accepted only after the following acceptance criteria are met; otherwise the run is rejected for bound testing:

1. Branch-matching injections yield $|\epsilon_{\text{bal}}| < 10^{-4}$ as defined in Section 11.4.
2. OFF-branch reversibility satisfies the design target $\mathcal{F}(100 \text{ ns}) > 0.99$ (and the data-quality gate $\mathcal{F} > 0.9$), without state-dependent acceptance bias.
3. State-conditioned absorber response histograms (conditioned on prepared x) are consistent with the calibrated confusion matrix $p(y|x)$ within stated uncertainty. Concretely, for each preparation x we compare the empirical outcome counts in the absorber record register to the predicted outcome probabilities using a χ^2 goodness-of-fit test (with Yates correction for binary outcomes when appropriate) and require $p > 0.01$ after Holm-Bonferroni correction over the tested x values. Runs failing this criterion are rejected for bound testing. This establishes that the calorimetric record register and the dispersive calibration probe the same effective pointer distinguishability within uncertainty.
4. Null protocols (toggle-only and pointer-disabled; Section 13) yield ΔQ consistent with zero within 2σ on the modulation timescale.

The dispersive calibration blocks may themselves generate classical records in the calibration readout chain; this does not enter the Stage 2 calorimetric channel accounting being tested, but is used only

to estimate $I(X;Y)$ for the chosen operating point. If ΔS_{sys} or ΔS_{mem} is non-negligible, or if external work W is supplied, the full bound reads

$$\langle Q_{\text{rec}} \rangle \geq k_B T \ln 2 \cdot I(X;Y) - k_B T \Delta S_{\text{sys}} - k_B T \Delta S_{\text{mem}} - W, \quad (6)$$

where $W \geq 0$ is work supplied to the apparatus by an external agent during Stage 2 (if any). We design the protocol to minimize these correction terms (cyclic apparatus, no external work). The quantity $I(X;Y)$ is a conservative (lower-bound) estimator of the thermodynamic information created because the actual record may contain additional correlations with unmonitored environmental degrees of freedom, so the true information gain $I_{S:M}^{\text{gain}} \geq I(X;Y)$. Equation (4) is therefore a conditional operational statement for the Stage 2 channel under C1 to C6, not a universal claim for arbitrary measurement decompositions. Protocols that keep the pointer coherent and defer entropy export fall into the cost-shifting classes of Section 4, where dissipation can move to later readout/reset.

Interpretation of Stage 2 versus Stage 3 bookkeeping. The generalized second law bounds entropy production for each irreversible channel. Equation (4) is written for the Stage 2 channel under our boundary choice: reusable pointer/control modes are internal and cycled ($\Delta S_{\text{mem}} \approx 0$ under C3), while the stabilized macroscopic record is bath-side. Under that bookkeeping, entropy exported to the bath during record stabilization contributes to $\beta \langle Q_{\text{rec}} \rangle$ in Stage 2. If Stage 3 later erases a stored record, it typically incurs further dissipation; however, the exact Stage 2/Stage 3 split in a full-cycle budget is protocol-dependent and must be evaluated with explicit state accounting.

Why this bookkeeping matches the proposed ON/OFF experiment. In the ON branch, the absorber node is simultaneously (i) the physical medium that stabilizes the macroscopic record via rapid thermalization into many degrees of freedom and (ii) the calorimetric node whose deposited energy is measured. Once the pointer excitation has thermalized into the absorber's microscopic degrees of freedom, recovering it as usable work would require microscopic control of those degrees of freedom and their correlations, which is operationally excluded (C4). In the OFF branch, the protocol is engineered to keep the pointer in a protected mode and to apply an explicit inverse operation before uncontrolled bath coupling occurs. Under these conditions, the ON/OFF differential isolates the Stage 2 entropy-export event rather than a later reset cost.

The temperature T is the effective temperature of the dissipative degrees of freedom that stabilize the record (Stage 2), calibrated in situ on the absorber/bath node.

3.3. Operational Conditions (C1 to C6)

The bound is asserted as a *conditional* statement:

- C1** A thermal bath at temperature T is present during Stage 2.
- C2** A classical register Y is created (decohered pointer with redundant environmental encoding).
- C3** The apparatus begins in a standard state and is cyclic over Stage 2, or its entropy change is explicitly accounted.
- C4** Irreversibility arises from bath coupling: no reversal operation that includes the bath degrees of freedom is performed on the experiment timescale. This is the operational definition of "objective": the record persists because reversing it would require controlling the bath.
- C5** No unaccounted work reservoir supplies free energy; if work is supplied, the inequality includes the work term.
- C6** $I(X;Y)$ is the actual classical mutual information for the chosen measurement strength.

What the bound does not imply:

- Measurement *in general* always costs $k_B T \ln 2$ per bit.
- Heat dissipation alone proves objectivity.
- All entropy production in a measurement protocol is due to Landauer cost.

Condition verification. Table 1 maps each condition to its experimental diagnostic and the failure signature that would invalidate the bound.

Table 1. Experimental diagnostics for conditions C1 to C6.

Cond.	Diagnostic	Failure signature
C1	In-situ thermometry on absorber	T undefined or drifting
C2	OFF-branch \mathcal{F} + stability $\tau > \tau_c$	Reversal succeeds at all delays
C3	Periodic null checks (toggle/pointer-disabled)	Net $\Delta S_{\text{mem}} \neq 0$
C4	Control $3\tau_c$ sweep	Record reversible on expt. timescale
C5	Uncompute-only null	Asymmetric drive heat detected
C6	Dispersive calibration + acceptance criteria	I_{cal} inconsistent with $p(y x)$

Relation to coherence-based thermodynamic bounds. Independent results in quantum thermodynamics quantify the free-energy cost of destroying coherence in a specified basis via the *relative entropy of coherence* $C_{\text{rel}}(\rho) \equiv S(\rho_{\text{diag}}) - S(\rho)$, where ρ_{diag} is ρ dephased in the relevant measurement/energy basis. In uncontrolled dephasing, this implies a minimal dissipation scale of order $Q \gtrsim k_B T C_{\text{rel}}(\rho)$ for the coherence that is irreversibly discarded to the bath. For pure-state projective measurements in the measurement basis, $C_{\text{rel}}/\ln 2$ reduces to the Shannon outcome entropy and reproduces the familiar Landauer scale. Our experiment instead targets the *record-formation* channel and uses the directly calibratable classical mutual information $I(X;Y)$ of the stabilized record as the operational metric; the coherence diagnostics shown in Figure 2 serve as witnesses of decoherence/stabilization rather than as the tested inequality.

4. When the Bound Does Not Apply (or Is Shifted)

We distinguish four non-applicability and cost-shifting classes:

- A. Premeasurement only.** If the pointer is captured before environmental coupling (Stage 1 only), no record is formed and no Landauer cost is incurred. This is the regime of reversible measurement protocols. [10] Coherence can be recovered by applying the inverse unitary.
- B. Deferred record in quantum memory.** If the pointer state is stored coherently without thermalization, no classical record exists and the cost is deferred until readout or reset. This is not free: maintaining coherence over time is a resource cost that scales with storage duration.
- C. Fresh memory or work-assisted recording.** If a low-entropy memory is consumed or external work supplies free energy, the heat released to the bath can be reduced. The cost is shifted into memory entropy or work, not eliminated.
- D. Weak or zero-information regimes.** When $I(X;Y) < 1$, the bound scales proportionally. For constant preparation with $I(X;Y) = 0$, the predicted heat is zero within experimental noise.

4.1. Reconciliation with Reversible Measurement Protocols

Latune and Elouard [8] analyze thermodynamically optimal measurement protocols and show that measurement can approach reversibility. This is **compatible** with our framework:

1. Their “reversible” protocols operate in Stage 1 (premeasurement).
2. No objective classical record is created in the reversible limit.
3. If an objective record is eventually created, the Landauer cost is paid.

The key distinction: **correlation** (Stage 1) is reversible; **record formation** (Stage 2) is not.

This mapping is also consistent with Bennett’s reversible computation arguments: [2] premeasurement is analogous to reversible computation; record formation is analogous to writing output to a permanent register. The Landauer cost is paid when you commit to a permanent record.

4.2. Addressing Norton's "Waiting for Landauer"

Norton [3] argues that Landauer's principle is not universally applicable. We agree with Norton's caution, and respond (following the defense by Ladyman and Robertson [4]) by specifying the operational conditions under which it applies:

- **Does apply:** Record formation with entropy flow to the bath (Stage 2), when conditions C1 to C6 hold.
- **Does not apply:** Reversible premeasurement (Stage 1).
- **Does not apply:** Quantum correlations that are never stabilized as classical records.

4.3. Work vs. Heat Classification of Pointer Absorption

A natural objection is that the pointer photon arrives as a coherent microwave pulse, directed energy that might be classified as *work* rather than heat, trivially satisfying the bound. We address this as follows.

The work term W in condition C5 refers to free energy supplied by an external agent to *drive* the record-formation process (e.g., a work reservoir that lowers the thermodynamic barrier to stabilization). In our protocol, no such external work is supplied: the pointer-absorber coupling is passive, and thermalization occurs spontaneously. The pointer photon's energy is the *input* to Stage 2, the physical carrier of information, not work done *on* the record-formation channel.

Upon absorption, the photon's energy is distributed among $N \gg 1$ electron-phonon degrees of freedom in the absorber at temperature T . After thermalization, the energy cannot be extracted as work without additional information about the absorber's microstate. This irreversible thermalization *is* the record-formation process: it converts a low-entropy coherent state (the pointer) into a high-entropy thermal state (the absorber macrostate), creating the classical record. The deposited energy therefore satisfies the thermodynamic definition of heat: stochastic energy transfer from the system to a thermal reservoir, increasing the reservoir's entropy.

Operational criterion used in this work. In the accounting of Section 3, energy transfer is treated as "heat" for the Stage 2 channel once the pointer excitation has irreversibly thermalized into many uncontrolled absorber degrees of freedom at temperature T , producing an entropy increase that cannot be reversed on the experimental timescale without microscopic bath control (C4). The bound is asserted for the irreversible absorber-coupling map (Stage 2), not for the coherent control work used to prepare Stage 1 correlations.

The experimental controls provide an independent check: the toggle-only and pointer-disabled null protocols (Section 13) bound any residual work-like contributions from control lines and switching operations. Any ON/OFF asymmetry from these sources enters the systematic budget (Table 3) and is subtracted before residual analysis.

5. Explicit Model: Where the Landauer Cost Is Paid

Having specified the operational regime and its boundaries, we now construct a concrete system/pointer/bath model to identify the precise microscopic origin of the heat flow.

5.1. The Model: System + Pointer + Bath

Consider a spin-1/2 system in superposition:

$$|\psi_S\rangle = \alpha|\uparrow\rangle + \beta|\downarrow\rangle, \quad |\alpha|^2 + |\beta|^2 = 1. \quad (7)$$

The apparatus consists of a pointer degree of freedom P (initially in state $|P_0\rangle$), e.g., a resonator mode) and a thermal environment at temperature T with many degrees of freedom E , initially in $\rho_E^{\text{th}} \propto e^{-H_E/(k_B T)}$. For compactness, the ket notation below uses $|E_0\rangle$ as a representative purification/microstate of this thermal ensemble.

5.2. Stage 1: Premeasurement (Reversible)

Unitary system/pointer coupling:

$$(\alpha|\uparrow\rangle + \beta|\downarrow\rangle)|P_0\rangle \xrightarrow{U_{SP}} \alpha|\uparrow\rangle|P_\uparrow\rangle + \beta|\downarrow\rangle|P_\downarrow\rangle. \quad (8)$$

This process is thermodynamically reversible in principle. The pointer states are correlated but no entropy has been exported to the environment. This is *not* record formation.

5.3. Stage 2: Record Formation (Irreversible; Landauer Cost Paid)

The pointer couples to the thermal bath, creating redundant environmental copies:

$$\begin{aligned} & \alpha|\uparrow\rangle|P_\uparrow\rangle|E_0\rangle + \beta|\downarrow\rangle|P_\downarrow\rangle|E_0\rangle \\ & \xrightarrow{\text{spread}} \alpha|\uparrow\rangle|P_\uparrow\rangle|E_\uparrow^{(N)}\rangle + \beta|\downarrow\rangle|P_\downarrow\rangle|E_\downarrow^{(N)}\rangle, \end{aligned} \quad (9)$$

where $|E_\sigma^{(N)}\rangle$ represents $N \gg 1$ environmental degrees of freedom encoding outcome σ . The environmental states satisfy:

$$\langle E_\uparrow^{(N)} | E_\downarrow^{(N)} \rangle \approx e^{-N/N_c} \rightarrow 0 \quad \text{for } N > N_c. \quad (10)$$

In the operational regime C1 to C6, we hypothesize that this environmental coupling is where a Landauer-scale thermodynamic cost is incurred. The physical basis: the redundant environmental encoding is an irreversible entropy export; reversing it would require controlled access to all N bath degrees of freedom, which is operationally excluded by C4. Under the generalized second law (Section 3), this irreversible entropy production implies $Q \geq k_B T \ln 2$ per bit of classical information stabilized. The proposed experiment tests whether this cost is detectable as a calorimetric signature.

Equation (4) is therefore a lower bound for the first logical bit in this open-system regime; in this model family, additional redundancy generally increases total entropy export to the environment, though exact scaling is protocol-dependent.

5.4. Identification of the Record-Formation Information

After Stage 2, tracing over the environment gives the reduced state:

$$\rho_{SP}^{\text{final}} = |\alpha|^2 |\uparrow\rangle\langle\uparrow| \otimes |P_\uparrow\rangle\langle P_\uparrow| + |\beta|^2 |\downarrow\rangle\langle\downarrow| \otimes |P_\downarrow\rangle\langle P_\downarrow|. \quad (11)$$

The classical information created in the record is (in bits):

$$H(Y) = H(|\alpha|^2) = -|\alpha|^2 \log_2 |\alpha|^2 - |\beta|^2 \log_2 |\beta|^2. \quad (12)$$

For equal superposition ($|\alpha|^2 = |\beta|^2 = 1/2$), $H(Y) = 1$ bit. This single-state expression is the record-outcome entropy. The experimental bound variable remains $I(X; Y) = I(X; Y)$, where X is the prepared classical label (Section 2.2).

The quantitative heat bound $Q \geq k_B T \ln 2 \cdot I(X; Y)$ follows from the generalized second law (Section 3), applied to this model with $I_{S:M}^{\text{gain}} \geq I(X; Y)$. The model identifies the *physical mechanism*: record formation occurs when the pointer-environment overlap $\langle E_\uparrow^{(N)} | E_\downarrow^{(N)} \rangle$ drops below the reversal criterion, and the resulting entropy flow to the environment carries the Landauer cost. For 1 bit: $Q \geq k_B T \ln 2$.

5.5. Lindblad Master Equation Validation

To bridge the gap between the analytical model and the proposed experiment, we solve the Lindblad master equation for a qubit-resonator system with parameters matching the experimental architecture. [17]

Model. A transmon qubit (two-level, $T_1 = T_2 = 50 \mu\text{s}$) dispersively coupled ($\chi/2\pi = 2 \text{ MHz}$) to a resonator ($\omega_r/2\pi = 7 \text{ GHz}$). The post-premeasurement state is

$$|\psi_{\text{PM}}\rangle = \frac{1}{\sqrt{2}}(|0\rangle|\alpha\rangle + |1\rangle|-\alpha\rangle), \quad (13)$$

with $|\alpha|^2 = \bar{n} = 1$, representing Stage 1 completion. Two Stage 2 branches are simulated independently:

- **ON branch:** the resonator decays into the absorber with $\kappa_{\text{ON}}/2\pi = 5 \text{ MHz}$ (fast thermalization).
- **OFF branch:** the pointer is held in a storage cavity with $Q_{\text{storage}} = 5 \times 10^7$ ($\kappa_{\text{OFF}}/2\pi \approx 140 \text{ Hz}$), and an ideal reversal (controlled displacement) is applied after delay τ_d .

ON branch results (Figure 2). The resonator photon number decays from $\bar{n} = 1$ to zero within $\sim 0.5 \mu\text{s}$. The cumulative heat deposited in the bath reaches $Q_{\text{bath}} = \hbar\omega_r \approx 4.64 \times 10^{-3} \text{ zJ}$, equal to the pointer photon energy $h\nu$ and exceeding $k_{\text{B}}T \ln 2$ by a factor of ~ 48 . Because the simulation begins from the post-premeasurement entangled state $|\psi_{\text{PM}}\rangle$, the reduced qubit state is already nearly maximally mixed at $t = 0$ (for $\bar{n} = 1$, $S_q(t=0) \approx 0.99 \text{ bit}$ in our QuTiP model) and remains $\sim 1 \text{ bit}$ as resonator decay transfers which-path information irreversibly to the bath. The qubit off-diagonal element $|\rho_{01}|$ simultaneously decays, reflecting the loss of quantum coherence during record stabilization. Here C_{rel} is computed for the reduced qubit state in the measurement basis and is included as a coherence witness rather than as the tested inequality. These dynamics confirm that the Landauer bound $Q \geq k_{\text{B}}T \ln 2 \cdot I(X;Y)$ is satisfied (with large positive residual) in the deep-quantum operating regime.

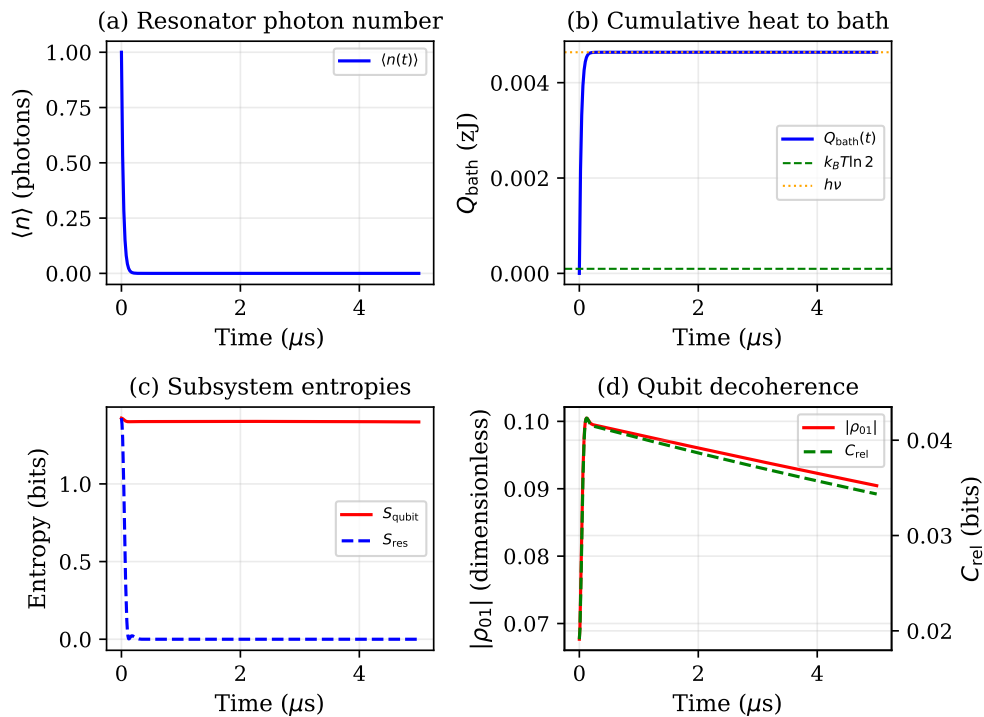


Figure 2. Lindblad simulation of the ON branch (record formation). (a) Resonator photon number decaying into the absorber ($\kappa_{\text{ON}}/2\pi = 5 \text{ MHz}$). (b) Cumulative heat $Q_{\text{bath}} = \hbar\omega_r(\bar{n}_0 - \bar{n}(t))$ reaching $h\nu \approx 4.64 \times 10^{-3} \text{ zJ}$, far above the Landauer floor $k_{\text{B}}T \ln 2 \approx 9.57 \times 10^{-5} \text{ zJ}$. (c) Subsystem entropies: the qubit entropy is already $\sim 1 \text{ bit}$ after premeasurement and remains near that value; resonator entropy peaks during decay and returns to zero. (d) Qubit off-diagonal $|\rho_{01}|$ and relative entropy of coherence C_{rel} during decoherence (left axis: $|\rho_{01}|$; right axis: C_{rel} in bits). The qubit-marginal C_{rel} can increase transiently because the qubit starts nearly maximally mixed after premeasurement; non-monotonicity in the reduced-state coherence does not contradict the channel-level entropy production and is expected for entangled bipartite states. Parameters: $\bar{n} = 1$, $\chi/2\pi = 2 \text{ MHz}$, $T = 10 \text{ mK}$, $T_1 = T_2 = 50 \mu\text{s}$.

OFF branch and τ_c extraction (Figure 3). A reversal delay sweep determines the reversibility time τ_c . At each delay τ_d , the premeasurement state evolves with the storage-cavity decay rate, then a time-dependent controlled displacement (accounting for the dispersive phase rotation) returns the resonator to vacuum. The recovery fidelity $\mathcal{F}(\tau_d) = \langle \psi_0 | \rho_{\text{rev}} | \psi_0 \rangle$ is computed against the initial (pre-premeasurement) state.

The fidelity remains above 0.99 for $\tau_d < 100$ ns, providing $\approx 78\times$ margin over the catch-and-release reversal window. Fidelity drops below 0.9 at $\tau_c \approx 8 \mu\text{s}$, driven primarily by qubit T_2 decoherence and slow photon leakage from the storage cavity. This is consistent with the order-of-magnitude estimate $\tau_c \sim 8$ to $50 \mu\text{s}$ in Section 12; varying the fidelity threshold from 0.85 to 0.95 shifts τ_c by $\lesssim 30\%$, confirming robustness of the extraction.

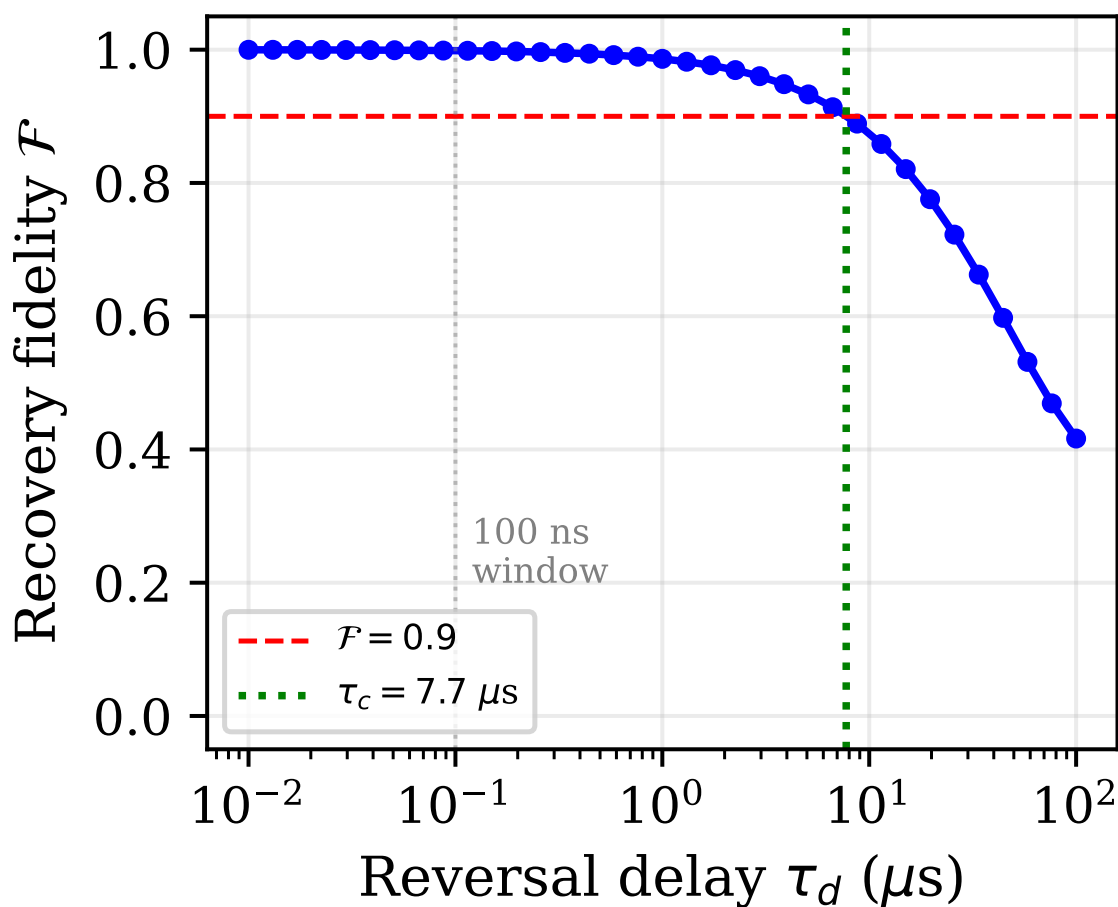


Figure 3. OFF-branch reversal fidelity versus delay τ_d . The fidelity remains near unity for $\tau_d \lesssim 1 \mu\text{s}$ and drops below 0.9 at $\tau_c \approx 8 \mu\text{s}$ (green dashed line), providing $\approx 78\times$ margin over the 100 ns reversal window (gray dotted line). Storage cavity $\mathcal{Q} = 5 \times 10^7$; qubit $T_1 = T_2 = 50 \mu\text{s}$.

Model limitations. The simulation uses a Markovian bath (Lindblad), finite Fock-space truncation ($N = 10$), and an idealized controlled-displacement reversal. The premeasurement state is prepared directly rather than simulated through an explicit drive pulse. Crucially, the simulated reversal applies an instantaneous unitary rather than modeling the Hamiltonian dynamics of the microwave drive lines. Consequently, the simulation validates qualitative system dynamics and establishes an upper bound on τ_c , but it does not track physical heat dissipated by the control electronics. In the proposed experiment, any asymmetric drive heat is bounded empirically by the uncompute-only null protocol, not by the simulation.

6. Experimental Architecture

6.1. Platform and Setup

- **System:** Superconducting transmon qubit ($\omega_q/2\pi \sim 5$ GHz)
- **Readout:** Dispersive measurement via coupled resonator
- **Calorimeter:** On-chip nanocalorimeter (TES or SNS nanobolometer class) in thermal contact with the record-formation channel (not the amplifier chain)
- **Environment:** Dilution refrigerator; $T \approx 10$ mK is the measured effective temperature of the record-formation absorber, calibrated in situ

Thermal-detector qubit readout has been demonstrated with single-shot fidelity and microsecond durations, [11] establishing platform feasibility.

6.2. The ON/OFF Toggle Architecture

The core design element is a **hardware toggle** that routes the pointer mode to different fates while keeping all other dissipation channels identical.

6.2.1. Common Path (Both Branches)

1. A classical random number generator selects a preparation label $x \in \{0, 1, \dots\}$ according to the prior $p(x)$. The qubit is then prepared in the corresponding state $|\psi_x\rangle$ (e.g., $|\psi_0\rangle = |0\rangle$, $|\psi_1\rangle = |+\rangle$). The label x is the classical variable X that enters $I(X; Y)$.
2. Apply dispersive readout pulse (identical timing, amplitude, frequency).
3. Qubit-resonator entanglement creates pointer state (premeasurement).

6.2.2. Branch Point: Three-Port Tunable Coupler

A flux-biased three-port tunable coupler (e.g., transformer SQUID device [15,16]) routes the pointer mode to different fates while keeping upstream conditions identical. The coupler must satisfy:

- **Amplitude/phase matching:** identical transfer function in both states
- **Isolation:** ≥ 40 dB between ON and OFF paths to limit leakage
- **Low added dissipation:** routing loss $\ll \Delta Q$ per shot; switching within a few ns. The switch is thermally anchored to the bath and spatially isolated from the calorimeter to ensure control-pulse dissipation does not couple to the sensor.
- **Characterization:** verified by injected calibration pulses upstream

6.2.3. ON Branch: Record Formation

The on-chip absorber (AuPd thin-film resistor or SNS junction) has a dual role: it is both the *record register* (where the pointer photon is thermalized into $N \gg 1$ electron and phonon degrees of freedom, creating the objective record) and the *calorimetric sensor* (whose temperature rise is the measured signal). The calorimeter therefore directly couples to the physical location where the record is formed, not to a downstream amplifier chain.

1. Pointer mode absorbed by the on-chip absorber.
2. Thermalization creates the macroscopic record in the absorber's internal degrees of freedom ($N \gg 1$ environmental DOF).
3. Heat deposited in calorimeter: $Q_{\text{ON}} = Q_{\text{cm}} + Q_{\text{rec}}$.

6.2.4. OFF Branch: Coherent Reversal

The coupler must route the pointer into the storage cavity before the pointer has time to couple to uncontrolled environmental modes (stray radiation, substrate phonons). The routing delay must satisfy $\tau_{\text{route}} \ll \tau_c$, where τ_c is the record-formation timescale. For typical dispersive readout, the pointer is a coherent state in a well-isolated resonator mode, and the routing can complete in a few ns, well before environmental coupling becomes significant.

1. The pointer mode is adiabatically captured in a high- Q storage cavity ($Q \sim 5 \times 10^7$, with minimum acceptable $Q > 10^7$) via the tunable coupler. Transfer leakage must satisfy $\delta Q_{\text{transfer}} < 0.1 \cdot k_B T \ln 2$, i.e., $< 9.57 \times 10^{-6}$ zJ at 10 mK. For a single-photon pointer ($E_\gamma \approx 4.64 \times 10^{-3}$ zJ at 7 GHz), this requires capture efficiency $> 99.8\%$, consistent with demonstrated catch-and-release protocols. [15]
2. A measurement reversal pulse sequence is applied: [10] by reversing the dispersive interaction (through qubit echo or opposite-phase drive), the qubit-pointer entanglement is erased. This uncomputation must complete within the pointer's coherence time. For transmon qubits with $T_2 \gtrsim 50 \mu\text{s}$, a < 100 ns echo sequence is consistent with demonstrated weak-measurement reversal timescales. [10]
3. The qubit's coherence is fully restored with high fidelity ($\mathcal{F} > 0.9$), and the storage cavity returns to vacuum.
4. No classical outcome is recorded; ideally no net heat is generated: $Q_{\text{OFF}} = Q_{\text{cm}} + \delta Q_{\text{leak}}$, with $\delta Q_{\text{leak}} \ll Q_{\text{rec}}$.

6.3. Measurand and Differential Estimator

The primary observable is the **differential heat**:

$$\Delta Q \equiv \langle Q \rangle_{\text{ON}} - \langle Q \rangle_{\text{OFF}}. \quad (14)$$

A minimal branch decomposition gives:

$$Q_{\text{ON}} = Q_{\text{cm}} + Q_{\text{rec}}, \quad (15)$$

$$Q_{\text{OFF}} = Q_{\text{cm}} + \delta Q_{\text{leak}}, \quad (16)$$

where Q_{cm} is common-mode energy deposited in the absorber/calorimeter node (control pulses, routing losses, switch-drive dissipation) shared by both branches; room-temperature amplifier-chain dissipation is outside the Stage 2 system boundary and does not enter this accounting. δQ_{leak} is parasitic heat in the OFF branch (Section 8). Under successful matching and low leakage, $\langle \Delta Q \rangle \approx \langle Q_{\text{rec}} \rangle - \langle \delta Q_{\text{leak}} \rangle \approx \langle Q_{\text{rec}} \rangle$.

Switch-drive dissipation (flux-pulse heating in bias lines, 50Ω termination losses) is common-mode because both branches are toggled with identical control pulses. The toggle-only null (Section 13) verifies that any residual ON/OFF asymmetry from the switch drive is below the detection criterion.

This design suppresses common-mode backgrounds (drive dissipation, routing losses) that are identical in both branches. The OFF branch is not merely a different input state but an explicit *no-objective-record* implementation, isolating the thermodynamic cost of record formation from premeasurement correlation.

6.4. Shot Timing and Experimental Sequence

1. **Initialize:** Reset qubit and resonator to ground state.
2. **Premeasurement:** Apply dispersive readout pulse to create pointer.
3. **Route:** Switch pointer to ON (absorber) or OFF (catch-and-release).
4. **Absorb/Capture:** ON branch thermalizes; OFF branch stores coherently.
5. **Reversal check (OFF):** Uncompute and verify recovered-state fidelity.
6. **Thermal wait:** Allow calorimeter to relax or deconvolve overlapping pulses.

To avoid pile-up, choose $f_{\text{rep}} \ll 1/\tau_{\text{th}}$ or validate deconvolution for partially overlapping responses. Here f_{rep} denotes the ON/OFF *pair* repetition rate (pairs per second).

7. Thermal Circuit Model

7.1. Single-Node Model

We model the calorimeter as a single thermal node with heat capacity C and thermal conductance G to the bath. The thermal time constant is $\tau_{\text{th}} = C/G$. For an impulse energy deposition Q at t_0 , the temperature response is

$$\Delta T(t) = \frac{Q}{C} e^{-(t-t_0)/\tau_{\text{th}}} \Theta(t-t_0), \quad (17)$$

with impulse response $h(t) = (1/C) e^{-t/\tau_{\text{th}}} \Theta(t)$. (A single-node model is used as a first-order description; electron-phonon two-temperature dynamics may modify the impulse response shape at early times but do not change the integrated energy $Q = \int P dt$, which is the thermodynamic quantity entering the bound.)

Using representative parameters at 10 mK:

$$C \approx 10^{-18} \text{ J/K} \quad (\text{AuPd absorber}), \quad (18)$$

$$G \approx 10^{-12} \text{ W/K} \quad (\text{electron-phonon}), \quad (19)$$

giving $\tau_{\text{th}} \approx 1 \mu\text{s}$ (design target; demonstrated SNS devices achieve $\tau_{\text{th}} \sim 30 \mu\text{s}$ at comparable NEP [12]). The single-shot temperature rise for $Q_{\text{rec}} \approx 9.57 \times 10^{-5} \text{ zJ}$ is

$$\Delta T \approx \frac{Q_{\text{rec}}}{C} \approx 0.1 \mu\text{K} \quad (95.7 \text{ nK}). \quad (20)$$

To avoid attenuation of the lock-in response, choose $f_{\text{mod}} \ll 1/(2\pi\tau_{\text{th}})$.

7.2. Per-Shot Energy Estimator

Define a linear estimator using a weight function matched to the impulse response:

$$\hat{Q}_i = \int w(t-t_i) V_i(t) dt, \quad (21)$$

where $w(t)$ is the matched filter (optimal for stationary Gaussian noise) and $V_i(t)$ is the calorimeter voltage trace for shot i .

8. OFF-Branch Leakage Budget

Even in the OFF branch (no intentional dissipation), parasitic heat flows can occur. We enumerate all known sources and bound the total leakage δQ_{leak} .

Cavity photon decay.

The storage cavity has finite Q , so pointer photons leak as heat during storage. For a 7 GHz photon ($E_\gamma \approx 4.64 \times 10^{-3} \text{ zJ}$) in a cavity with $Q = 5 \times 10^7$, the photon lifetime is $\tau_{\text{cav}} = Q/\omega \approx 1.1 \text{ ms}$. During a 100 ns reversal window, the fractional loss is $\sim 10^{-4}$, giving $\delta Q_{\text{cav}} \approx 4.1 \times 10^{-7} \text{ zJ}$ per photon. This is the dominant leakage source but remains $\ll 9.57 \times 10^{-5} \text{ zJ}$.

Coupler isolation leakage.

The tunable coupler provides $\geq 40 \text{ dB}$ isolation (10^{-4} power transmission). For a single-photon pointer ($E_\gamma \approx 4.64 \times 10^{-3} \text{ zJ}$), the leakage is $\delta Q_{\text{coupler}} \approx 4.64 \times 10^{-7} \text{ zJ}$.

Transfer leakage.

Imperfect capture into the storage cavity contributes $\delta Q_{\text{transfer}}$. The requirement from Section 6 is $\delta Q_{\text{transfer}} < 0.1 \cdot k_B T \ln 2 \approx 9.57 \times 10^{-6} \text{ zJ}$, corresponding to $> 99.8\%$ capture efficiency.

Quasiparticle generation.

Fast flux pulses can break Cooper pairs. Without mitigation, quasiparticle heating could reach $\sim 10^{-3}$ zJ per cycle. With well-shaped pulses and filtered bias lines (standard practice at 10 to 20 mK), quasiparticle generation is suppressed to $\delta Q_{qp} < 10^{-6}$ zJ per cycle.

Total leakage budget.

Summing the mitigated contributions:

$$\begin{aligned}\delta Q_{\text{leak}} &\approx \delta Q_{\text{cav}} + \delta Q_{\text{coupler}} + \delta Q_{\text{qp}} + \delta Q_{\text{transfer}} \\ &\lesssim 4.1 \times 10^{-7} + 4.64 \times 10^{-7} + 10^{-6} + 9.57 \times 10^{-6} \\ &< 1.12 \times 10^{-5} \text{ zJ},\end{aligned}\tag{22}$$

satisfying the requirement $\delta Q_{\text{leak}} < 0.15 \cdot \Delta Q_{\text{target}}$ for $\Delta Q_{\text{target}} \approx 9.57 \times 10^{-5}$ zJ. This budget assumes $Q \sim 5 \times 10^7$ for the storage cavity (minimum acceptable $Q > 10^7$), ≥ 40 dB coupler isolation, reversal completion within ~ 100 ns, and filtered bias lines for quasiparticle suppression. If coupler isolation degrades to 30 dB, the coupler leakage rises by $10\times$, reducing the safety margin. In-situ verification of isolation is therefore a go/no-go gate: the experiment should not proceed until the blocked-branch test (Section 11) confirms the required isolation. Transfer leakage ($\delta Q_{\text{transfer}}$) dominates the budget at $\sim 83\%$ of total leakage; capture efficiency degradation below 99.5% would breach the 15% margin, making this the most fragile link in the leakage chain.

To keep OFF leakage below the Landauer scale, we require $(1 - \mathcal{F}) Q_{\text{pointer}} \ll k_B T \ln 2$. For a single-photon pointer ($Q_{\text{pointer}} \approx 4.64 \times 10^{-3}$ zJ), this requires $\mathcal{F} \gtrsim 0.98$. For $\bar{n} = 10$ photons ($Q_{\text{pointer}} \approx 4.64 \times 10^{-2}$ zJ), the requirement tightens to $\mathcal{F} \gtrsim 0.998$. Imperfect reversal ($\mathcal{F} < 1$) contributes an $\mathcal{O}(1 - \mathcal{F})$ systematic at the Landauer scale, calibrated experimentally via Control 3 (Section 13).

9. Sensitivity Analysis

9.1. Energy Scales and the Quantum-Thermal Hierarchy

At $T = 10$ mK,

$$k_B T \ln 2 \approx 9.57 \times 10^{-26} \text{ J} \approx 9.57 \times 10^{-5} \text{ zJ}.\tag{23}$$

For comparison, a single 7 GHz readout photon carries $h\nu \approx 4.64 \times 10^{-24}$ J $\approx 48.5 \times k_B T \ln 2$.

This separation of scales ($h\nu \gg k_B T$) is advantageous for detection but requires careful interpretation. We distinguish

$$\Delta Q_{\text{signal}} \equiv \langle Q_{\text{ON}} - Q_{\text{OFF}} \rangle,\tag{24}$$

$$\Delta Q_{\text{bound}} \equiv k_B T \ln 2 \cdot I(X; Y).\tag{25}$$

In the deep-quantum operating point used here (typically $\bar{n} \sim 1$ at GHz frequencies), ΔQ_{signal} is expected to be dominated by pointer absorption ($\sim \bar{n}h\nu$) and therefore to lie well above ΔQ_{bound} . The primary falsification statistic is the residual $r = \Delta Q_{\text{signal}} - \Delta Q_{\text{bound}}$ after full leakage/systematics propagation; negative r at high significance would falsify the bound. Near-saturation tests (where r is not a priori large) require different operating points that reduce pointer-energy scale toward the Landauer scale (e.g., lower-frequency pointers or other low-energy platforms). In numerical validation, finite-sample Monte Carlo estimates of r can fluctuate below zero when per-shot noise dominates at low effective SNR; this does not constitute falsification. The experimental falsification criterion is a statistically significant negative residual after full uncertainty propagation (Section 13).

Control 4 (prior variation) is interpreted as a mechanism check at fixed pulse energy. In the symmetric fixed-energy implementation described in Section 13, ΔQ_{signal} is expected to be approximately prior-independent; any observed dependence on the preparation prior provides a sensitive diagnostic

of state-dependent absorption, routing imbalance, or drift in the effective confusion matrix. By itself this control does not isolate the Landauer floor; it constrains the systematics model used in the residual analysis.

9.2. Detector Anchor

We anchor the feasibility estimate in the SNS nanobolometer of Kokkonen *et al.* [12] That platform has response times down to $\sim 30 \mu\text{s}$ at $\text{NEP} \sim 60 \text{ zW}/\sqrt{\text{Hz}}$ and a predicted calorimetric energy resolution $\varepsilon \approx 0.32 \text{ zJ}$ when integrating up to the thermal cutoff frequency. Related calorimetric qubit-readout demonstrations motivate using this detector class as a realistic baseline. [11,13]

9.3. Per-Shot Energy Uncertainty

For a calorimetric readout with effective bandwidth limited by τ_{th} , a conservative per-shot energy uncertainty may be taken as the reported calorimetric resolution:

$$\sigma_Q \sim \varepsilon \approx 3.2 \times 10^{-22} \text{ J} \quad (0.32 \text{ zJ}). \quad (26)$$

A more conservative estimate uses $\sigma_Q = \text{NEP} \cdot \sqrt{\tau_m}$ with a shorter integration window τ_m , giving $\sigma_Q \approx 0.6 \text{ zJ}$. The per-shot resolution ε assumes non-overlapping shots ($f_{\text{rep}} \ll 1/\tau_{\text{th}}$); for partially overlapping pulses, matched-filter deconvolution is required and the effective σ_Q increases. The accompanying simulation models the readout-error probability via a simplified homodyne discrimination ($p_{\text{err}} = \frac{1}{2} \text{erfc}(\sqrt{\eta})$); a near-floor test would require modeling finite measurement time, qubit relaxation during readout (T_1 errors), and state-dependent readout fidelity. We use $\sigma_Q = 0.32 \text{ zJ}$ for the primary shot-count estimate below; Table 2 presents the $\sigma_Q = 0.6 \text{ zJ}$ scenario to bound feasibility under less favorable conditions.

9.4. Required Shot Count

Each differential measurement $\Delta Q_i = Q_{\text{ON},i} - Q_{\text{OFF},i}$ has per-shot variance

$$\sigma_{\Delta Q}^2 = \sigma_{\text{ON}}^2 + \sigma_{\text{OFF}}^2 - 2 \text{Cov}(Q_{\text{ON}}, Q_{\text{OFF}}), \quad (27)$$

which reduces to $\sigma_{\Delta Q}^2 = 2\sigma_Q^2$ when ON and OFF noise are independent and have equal variance. In practice, lock-in demodulation and a shared thermal environment can introduce nonzero covariance. We therefore estimate $\sigma_{\Delta Q}$ empirically from interleaved calibration blocks and propagate the measured covariance as part of the systematic/uncertainty budget. For N independent ON/OFF pairs, the standard error of the mean is $\sigma_{\Delta Q} = \sigma_{\Delta Q} / \sqrt{N}$. Targeting $\text{SNR} \equiv \langle \Delta Q \rangle / \sigma_{\Delta Q}$ gives

$$N \gtrsim \left(\frac{\text{SNR} \sigma_{\Delta Q}}{\Delta Q} \right)^2, \quad (28)$$

which reduces to $N \gtrsim 2 \left(\frac{\text{SNR} \sigma_Q}{\Delta Q} \right)^2$ when $\sigma_{\Delta Q} = \sqrt{2} \sigma_Q$ (independent equal-variance ON/OFF noise). with $\Delta Q \sim h\nu \approx 4.64 \times 10^{-3} \text{ zJ}$ (using the actual signal scale in the quantum regime, which is $\sim 48 \times$ larger than the Landauer floor). At $T = 10 \text{ mK}$ and $\sigma_Q \approx 0.32 \text{ zJ}$:

$$N_{\text{SNR}=10} \sim 2 \left(\frac{10 \times 0.32}{4.64 \times 10^{-3}} \right)^2 \approx 9.5 \times 10^5. \quad (29)$$

This is drastically more feasible than detecting the bare Landauer limit. However, to rigorously test the pointwise inequality near the much smaller $k_B T \ln 2$ floor (at 10 mK , $\sim 9.57 \times 10^{-5} \text{ zJ}$), the required averaging is much larger: $N \sim 2.24 \times 10^9$ for $\sigma_Q = 0.32 \text{ zJ}$ at $\text{SNR} = 10$, or $N \sim 7.86 \times 10^9$ for $\sigma_Q = 0.6 \text{ zJ}$ (Table 2).

9.5. Integration Time Estimates

For $\tau_{\text{th}} \sim 1 \mu\text{s}$ (a design target not yet demonstrated), a repetition rate of $\sim 10^5$ pairs/s gives two distinct regimes: **(i)** for the full quantum-scale signal ($\Delta Q \sim h\nu$), $N_{\text{SNR}=10} \sim 9.5 \times 10^5$ integrates in ~ 10 s; and **(ii)** for Landauer-scale resolution ($\Delta Q \sim 9.57 \times 10^{-5}$ zJ), integration is ~ 1.5 to 6 hours for $\text{SNR} \sim 5$ to 10 with $\sigma_Q = 0.32$ zJ, or ~ 22 hours for $\sigma_Q = 0.6$ zJ at $\text{SNR} = 10$. Demonstrated SNS nanobolometers operate at $\tau_{\text{th}} \sim 30 \mu\text{s}$. [12] In the lock-in regime with $f_{\text{mod}} \sim 0.5$ to 5 kHz, the effective independent rate is typically a few kHz, stretching Landauer-scale integration to **days to weeks**. Faster bolometers or detector parallelization can recover hour-scale operation.

Table 2. Predicted signal and feasibility at $\text{SNR} = 10$ with $\sigma_Q = 0.6$ zJ (conservative, NEP-derived). t_{int} assumes 10^5 pairs/s (requires $\tau \lesssim 1 \mu\text{s}$); with demonstrated $\tau \sim 30 \mu\text{s}$ bolometers, multiply by $\sim 100\times$.

T (mK)	$k_B T \ln 2$ (zJ)	N_{pairs}	t_{int}^*
10	9.57×10^{-5}	7.86×10^9	22 hours*
50	4.79×10^{-4}	3.14×10^8	52 minutes*
100	9.57×10^{-4}	7.86×10^7	13 minutes*

9.6. Temperature Regimes

Higher temperatures increase signal but also increase thermal noise; 10 mK represents the optimal trade-off for current detector technology. Two routes can reduce the required averaging: (i) operating at elevated effective temperature (increasing $k_B T \ln 2$ linearly), and (ii) increasing $I(X; Y)$ per cycle via multi-bit record creation.

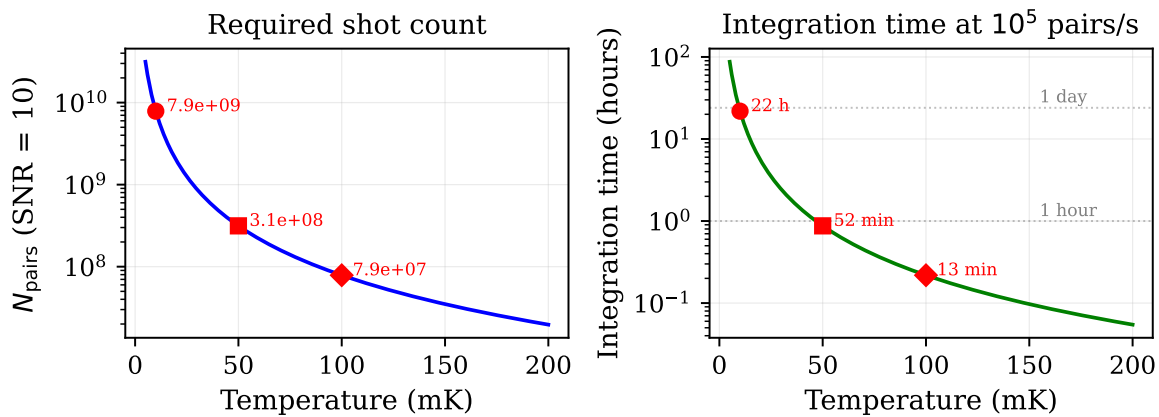


Figure 4. Temperature scaling of the Landauer-scale residual test ($\sigma_Q = 0.6$ zJ, $\text{SNR} = 10$, 10^5 pairs/s). Left: required ON/OFF pairs N . Right: integration time. Markers show the three operating points from Table 2. At 10 mK, integration requires ~ 22 hours with $\tau_{\text{th}} \lesssim 1 \mu\text{s}$ bolometers; at 50 mK, ~ 52 minutes. With demonstrated $\tau_{\text{th}} \sim 30 \mu\text{s}$ detectors, multiply times by $\sim 100\times$.

10. Lock-in Modulation Protocol

10.1. Modulation Scheme

Alternate ON/OFF branches at frequency f_{mod} chosen to lie above the $1/f$ noise corner (~ 10 Hz) and below the thermal cutoff $1/(2\pi\tau_{\text{th}})$. For the design target $\tau_{\text{th}} \sim 1 \mu\text{s}$, the cutoff is ~ 160 kHz and f_{mod} up to 5 kHz is comfortable. For demonstrated detectors with $\tau_{\text{th}} \sim 30 \mu\text{s}$, the cutoff is ~ 5.3 kHz, so $f_{\text{mod}} \lesssim 1$ kHz is appropriate.

The calorimeter voltage output is:

$$V(t) = V_{\text{DC}} + \mathcal{R} \Delta Q \cdot \cos(2\pi f_{\text{mod}} t) + n(t), \quad (30)$$

where \mathcal{R} is the calorimeter responsivity (V/J) and V_{DC} is the common-mode baseline. Lock-in demodulation extracts $\mathcal{R} \Delta Q$ while rejecting:

- Common-mode V_{DC} (premeasurement + routing losses)
- Low-frequency drift
- Out-of-band noise

10.2. Power Modulation

For ON/OFF pair repetition rate f_{rep} (pairs/s) and square-wave alternation, the fundamental power modulation amplitude is

$$P_1 = \frac{2}{\pi} \Delta Q f_{\text{rep}}, \quad (31)$$

linking the lock-in output directly to ΔQ .

10.3. Drift and Stability Control

To suppress low-frequency drift, we employ:

- Lock-in modulation as above, extracting ΔQ from the demodulated response.
- Mixing chamber temperature stabilized via PID feedback to $< 5 \mu\text{K}$ fluctuations at the modulation timescale.
- Periodic hardware null checks every 2 hours using the toggle-only and pointer-disabled protocols (Section 13).
- Allan deviation analysis of calibration pulse sequences to identify drift timescales.

Data segments are accepted only if: (1) calibration pulse drift remains $< 0.1 \cdot \Delta Q_{\text{target}}$, (2) null checks yield ΔQ consistent with zero within 2σ , and (3) no thermal excursions exceed $\pm 1 \mu\text{K}$ at the absorber.

11. Calibration Strategy

11.1. Photon-Number Calibration

Inject calibrated microwave pulses with known photon number \bar{n} :

$$Q_{\text{cal}} = \bar{n} \cdot \hbar \omega_r. \quad (32)$$

For $\omega_r/2\pi = 7 \text{ GHz}$ and $\bar{n} = 1$: $Q_{\text{cal}} = 4.64 \times 10^{-3} \text{ zJ}$. This provides an absolute energy scale.

11.2. Heater-Pulse Calibration

Use an on-chip resistive heater to inject known Joule energy $Q_{\text{heater}} = \int I^2(t) R dt$. Fit the resulting waveform to extract τ_{th} and validate the linear impulse response $h(t)$. This sets the maximum f_{rep} and determines whether deconvolution is required.

11.3. Linearity Verification

Sweep \bar{n} from 0.1 to 100 photons and verify $Q_{\text{meas}} = a \cdot \bar{n} + b$. The nonlinearity residual must be characterized at the ΔQ_{target} scale ($\sim 9.57 \times 10^{-5} \text{ zJ}$), not merely at the 0.1 zJ scale.

11.4. Branch Matching and Common-Mode Rejection

Inject identical calibration pulses upstream of the routing element and verify that the inferred energy is independent of ON/OFF state. Define a balance parameter:

$$\epsilon_{\text{bal}} = \frac{\langle \hat{Q} \rangle_{\text{ON}} - \langle \hat{Q} \rangle_{\text{OFF}}}{\langle \hat{Q} \rangle_{\text{ON}}}. \quad (33)$$

Residual common-mode bias is then $\delta Q_{\text{cm}} \approx \epsilon_{\text{bal}} Q_{\text{cm}}$, which must be propagated into the systematic error budget. We require $\epsilon_{\text{bal}} < 10^{-4}$ to keep systematic bias below the signal.

This corresponds to a common-mode rejection ratio (CMRR) of ≥ 80 dB, distinct from the ≥ 40 dB *isolation* between ON/OFF paths (which prevents thermal crosstalk). We reach $\epsilon_{\text{bal}} < 10^{-4}$ in two steps:

Step 1 (VNA pre-calibration). Match amplitude to < 0.01 dB ($\sim 10^{-3}$ fractional), phase to $< 0.5^\circ$, and path-length to < 0.1 mm (< 0.3 ps timing skew). This brings $\epsilon_{\text{bal}} \lesssim 10^{-3}$.

Step 2 (calorimetric closure). Inject identical calibration pulses through both branches and measure ΔQ via the lock-in chain. Trim with precision attenuators and phase shifters, iterate until the calorimetric differential satisfies $|\delta Q_{\text{cm}}| < 0.1 \cdot \Delta Q_{\text{target}}$, corresponding to $\epsilon_{\text{bal}} < 10^{-4}$. This second step compensates residual mismatch that VNA calibration alone cannot resolve.

Drift and stability budget. Sustaining $|\epsilon_{\text{bal}}| < 10^{-4}$ over multi-hour Landauer-scale integration is the most demanding engineering requirement. We require that ϵ_{bal} drift remains below 10^{-4} over 2-hour windows; periodic null checks (every 2 hours, Section 13) monitor and flag drift. If ϵ_{bal} degrades by a factor of $10\times$, the systematic bias δQ_{cm} grows proportionally and may exceed ΔQ_{target} ; such runs are rejected. Branch-balance stability is the hardest experimental gate in the proposed architecture.

11.5. Thermal Crosstalk Measurement

With the OFF branch blocked (forcing the pointer to always dissipate), measure any apparent heat in the OFF channel:

$$\eta_{\text{crosstalk}} \equiv \frac{Q_{\text{OFF,blocked}}}{Q_{\text{ON}}}. \quad (34)$$

We require $\eta_{\text{crosstalk}} < 0.01$ ($< 1\%$). For the absolute bias budget, the crosstalk-induced systematic is $\delta Q_{\text{xt}} = \eta_{\text{crosstalk}} \cdot Q_{\text{ON}}$. With $Q_{\text{ON}} \sim \bar{n}\hbar\omega \approx 4.64 \times 10^{-3}$ zJ (single 7 GHz photon) and $\eta < 0.01$, this gives $\delta Q_{\text{xt}} < 4.64 \times 10^{-5}$ zJ, which is below $\Delta Q_{\text{target}} \approx 9.57 \times 10^{-5}$ zJ. For multi-photon pointers ($\bar{n} \gtrsim 10$), a tighter crosstalk bound $\eta < 10^{-3}$ is needed, verified by the same blocked-branch test.

12. Reversibility Witness

Heat alone does not certify objectivity. We implement two complementary metrics.

12.1. Fidelity Metric

Reconstruct the qubit state after an OFF operation and compute

$$\mathcal{F} = |\langle \psi_{\text{initial}} | \psi_{\text{recovered}} \rangle|^2. \quad (35)$$

We distinguish two fidelity criteria. The *hardware design target* is $\mathcal{F} \gtrsim 0.99$ (Section 8), which keeps per-shot leakage below the Landauer scale. The *data-quality gate* is a looser $\mathcal{F} > 0.9$, applied to catch catastrophic reversal failures (e.g., quasiparticle events, flux jumps). Shots passing the 0.9 gate but falling below 0.99 contribute a small, characterized systematic (calibrated via Control 3). Shots below 0.9 are discarded. To control for acceptance bias, we (i) verify that the acceptance rate does not depend on the preparation label X or outcome Y , (ii) apply the same fractional random drop to ON shots to confirm $\langle Q \rangle_{\text{ON}}$ is unaffected, and (iii) report $\langle \Delta Q \rangle$ both with and without the fidelity gate to check consistency. Any systematic shift is propagated as a gating-bias uncertainty.

12.2. Operational Metric

As a real-time check, monitor the residual heat in the OFF calorimeter and qubit coherence via echo contrast. A successful uncomputation should yield calorimeter signal consistent with zero (per-shot residual $< \Delta Q_{\text{target}} \approx 9.57 \times 10^{-5}$ zJ after averaging).

12.3. Reversibility Time τ_c

Define τ_c as the smallest delay for which \mathcal{F} drops below 0.9. The calorimetric onset of $\Delta Q > 3\sigma$ at the same timescale is a *separate, independent* observable used to cross-validate the fidelity criterion,

not part of the definition. If fidelity loss and heat onset occur at different delays, that discrepancy is itself a diagnostic. τ_c empirically marks the temporal boundary of record formation.

Order-of-magnitude estimate. The reversal protocol requires qubit coherence (for the echo sequence) and pointer coherence (in the storage cavity). The limiting timescale is typically the qubit T_2 : for current transmon qubits, $T_2 \gtrsim 50 \mu\text{s}$. [21] The storage cavity lifetime is $\tau_{\text{cav}} = Q/\omega \approx 1.1 \text{ ms}$ (Section 8), which does not limit the protocol. Uncontrolled environmental coupling (substrate phonons, stray radiation, quasiparticle tunneling) degrades the pointer on timescales set by the internal quality factor, typically $\gtrsim 10 \mu\text{s}$ for well-isolated superconducting resonators at 10 mK. We therefore estimate $\tau_c \sim 8$ to $50 \mu\text{s}$, comfortably above the $\sim 100 \text{ ns}$ reversal window required by the catch-and-release protocol. The Lindblad simulation (Section 5.5) yields $\tau_c \approx 8 \mu\text{s}$ (via log-linear interpolation between bracketing fidelity points) with $T_1 = T_2 = 50 \mu\text{s}$ and $Q_{\text{storage}} = 5 \times 10^7$. The ratio $\tau_c/\tau_{\text{reversal}} \approx 78$ provides substantial margin for the OFF-branch uncomputation to complete before environmental decoherence compromises the reversal. Control 3 (Section 13) measures τ_c experimentally by sweeping the reversal delay.

13. Controls, Null Tests, and Falsification Criteria

We propose four primary discriminators. In the deep-quantum regime ($h\nu \gg k_B T \ln 2$), the absolute magnitude of ΔQ is dominated by pointer-energy thermalization and does not, by itself, discriminate record-formation thermodynamics from generic photon absorption. The scientifically distinctive content of the Phase 1 demonstration is therefore **Control 3**: the temporal coincidence of irreversible heat onset and the independently certified loss of reversibility at timescale τ_c . This timing signature is not predicted by energy conservation alone and constitutes the primary experimental signature of the boundary event. Controls 1, 2, and 4 establish the baseline, scaling, and systematic stability required to interpret Control 3 cleanly.

13.1. Control 1: Ground-State Baseline

Prepare the qubit in the ground state $|0\rangle$ across all runs. Here $H(X) = 0$ and $I(X; Y) = 0$.

Prediction: with fixed readout strength, ΔQ is *nonzero* and set by ON-branch pointer absorption, $\Delta Q \approx \bar{n}\hbar\omega_{\text{eff}} - \delta Q_{\text{leak}}$ (up to calibrated branch imbalance). Ground-state preparation sets $I(X; Y) = 0$ but does not remove pointer-energy flow to the ON absorber. This control therefore establishes the nonzero photon-absorption baseline used in residual analysis.

13.2. Control 2: Measurement-Strength Scaling

Prepare a mixed ensemble (e.g., 50% $|0\rangle$, 50% $|1\rangle$) so that $H(X) = 1$. Vary the measurement strength (pointer mean photon number \bar{n}) from $\bar{n} \approx 0$ to 1. This sweeps the realized mutual information $I(X; Y)$ from 0 to 1 bit.

Prediction: ΔQ scales linearly with \bar{n} , and thus monotonically with $I(X; Y)$. Because $I(X; Y)(\bar{n})$ is generally nonlinear, the relation between ΔQ and $I(X; Y)$ need not be globally linear; the required test is the pointwise inequality $\Delta Q \geq k_B T \ln 2 \cdot I(X; Y)$. In the deep quantum regime, $\Delta Q \sim h\nu$ typically lies far above the Landauer floor, confirming that the energetic cost scales with record-formation strength (Figure 1).

13.3. Control 3: Reversal-Delay Timing Sweep (Primary Boundary-Event Signature)

Vary the delay τ_d between premeasurement and the uncompute operation in the OFF branch. This control is the **centerpiece** of the experimental program: it directly probes the temporal boundary between reversible correlation and irreversible record formation.

Define the **boundary time**

$$t_B \equiv \inf\{t : \mathcal{F}(t) < \mathcal{F}_*\}, \quad (36)$$

with $\mathcal{F}_* = 0.9$ (the data-quality gate). This definition uses the irreversibility witness alone. The **prediction** of the boundary-event thesis is that heat onset in the ON/OFF differential coincides with

t_B : for $\tau_d < t_B$, the OFF branch successfully reverses (high fidelity, low heat), maintaining the full differential signal ($\Delta Q \approx Q_{\text{rec}}$); as τ_d exceeds t_B , reversal fails, the OFF branch begins to dissipate, and ΔQ degrades toward zero. Observing this temporal correlation between the independently measured fidelity drop and the calorimetric onset is a distinctive physical claim about when a reversible correlation becomes an irreversible classical fact, even in the deep-quantum regime where $\Delta Q \gg k_B T \ln 2$.

13.4. Control 4: Prior-Variation at Fixed Strength (Systematic Diagnostic)

Fix $\bar{n} \approx 1$ (strong measurement) but vary the preparation prior $p(1)$ from 0 to 1.

Prediction: At fixed measurement strength (fixed pointer energy), the ON/OFF differential is approximately prior-independent, $\Delta Q \approx \bar{n}\hbar\omega_{\text{eff}} - \delta Q_{\text{leak}}$, while $I(X;Y)$ varies with the prior through the realized confusion matrix. The inequality $\Delta Q \geq k_B T \ln 2 \cdot I(X;Y)$ must hold for all priors. Control 4 is primarily a **systematic diagnostic**: because ΔQ is predicted to be prior-independent at fixed \bar{n} for symmetric dispersive readout, any observed dependence on the preparation prior provides a sensitive diagnostic of state-dependent absorption, routing imbalance, or drift in the effective confusion matrix. It constrains the systematics model used in the residual analysis rather than independently testing the Landauer bound.

13.5. Auxiliary Null Protocols

Three additional nulls verify hardware systematics and bound asymmetric heat:

Toggle-only null. Alternate the routing element between ON and OFF positions while *disabling the dispersive readout pulse* (no pointer created). Any residual ΔQ bounds the heat injected by the switching operation itself.

Pointer-disabled null. Run the full pulse sequence but detune the qubit so that no system-pointer entanglement occurs (pointer in vacuum). Any residual ΔQ bounds routing-path asymmetry at the actual operating frequency.

Uncompute-only null. With the qubit detuned and the pointer in vacuum, apply the full OFF-branch uncomputation microwave pulse sequence. Because the ON branch does not receive these specific coherent reversal pulses, this null explicitly quantifies any asymmetric heat leaked from the uncomputation drive lines into the calorimetric channel.

All nulls must yield ΔQ consistent with zero within 2σ before production data is collected, ensuring drive-induced dissipation is fully characterized and subtracted.

13.6. Falsification Criteria

The record-formation heat bound is **falsified** if, under verified conditions, the pointwise inequality $\langle \Delta Q \rangle_j \geq k_B T \ln 2 \cdot I(X;Y)_j$ fails at any tested operating point j (measurement strength, prior, or temperature), after full uncertainty propagation:

1. For at least one verified point j , the residual $r_j \equiv \langle \Delta Q \rangle_j - k_B T \ln 2 \cdot I(X;Y)_j$ is negative with $> 3\sigma$ significance, where σ includes both statistical uncertainty ($\propto 1/\sqrt{N}$) and systematic contributions from Table 3.
2. A global weighted test over all points gives a negative mean residual at $> 3\sigma$.
3. Record formation is independently verified while ΔQ fails to exceed the calibrated leakage-corrected baseline expected for the chosen pointer-energy setting.

When testing multiple operating points, we treat one operating point (or a small pre-registered set of primary points) as the primary falsification test and report additional sweeps as secondary diagnostics; if multiple-comparison control is desired, a Holm-Bonferroni correction may be applied to the family of pointwise tests.

Conversely, the bound is **not falsified** (i.e., is consistent with the data) if all tested points satisfy $r_j \geq 0$ within 3σ , across measurement-strength sweeps, prior variation, and temperature scans. In the deep quantum regime ($\hbar\nu \gg k_B T$), large positive residuals are expected and indicate additional

irreversible dissipation in the ON branch not present in the OFF branch (e.g., routing asymmetry, additional thermalization channels). As a secondary diagnostic (not the primary falsification statistic), a fit of $\langle \Delta Q \rangle - \langle \Delta Q \rangle_{C1} = m I(X; Y) + b'$ can be reported after Control 1 establishes the baseline $\langle \Delta Q \rangle_{C1}$.

Table 3. Requirement-driven systematics for ΔQ .

Mechanism	Bias term	Mitigation	Verified by
Branch imbalance	$\epsilon_{\text{bal}} Q_{\text{cm}}$	S-parameter matching + calorimetric closure	Upstream pulse test
OFF leakage	δQ_{leak}	Coherent catch-release	Fidelity gating
Switch-only offset	Q_{switch}	Toggle-only null	Pointer-disabled null
Uncompute asymmetry	δQ_{uncomp}	Spatial decoupling of drive lines	Uncompute-only null
Nonlinearity	δQ_{nl}	Heater/pulse calibration	Linearity sweep
Drift	δQ_{drift}	Lock-in + PID	Baseline tracking

13.7. Systematics Summary

Tables 3 and 4 summarize the requirement-driven systematics budget and go/no-go technical targets, respectively.

The term δQ_{uncomp} captures any differential heat coupled into the calorimetric channel by the OFF-branch uncomputation drive sequence, which is not applied in the ON branch. In practice these coherent reversal pulses are short and strongly attenuated, so their expected dissipation at the absorber is small compared to pointer absorption in the deep-quantum regime. Regardless of expectation, the uncompute-only null measures this contribution directly and its uncertainty is propagated (and, if needed, subtracted) in the final $\langle \Delta Q \rangle$ budget.

Table 4. Go/no-go technical targets for the Phase 1 demonstration and near-floor residual tests.

Target	Requirement	Verified by
Branch balance	$ \epsilon_{\text{bal}} < 10^{-4}$	Upstream injection + calorimetric closure
OFF reversibility	$\mathcal{F}(100 \text{ ns}) > 0.99$	Tomography/echo contrast in OFF branch
Coupler isolation	$\geq 40 \text{ dB (design)}$	Blocked-branch leakage test
OFF leakage	$\delta Q_{\text{leak}} \ll k_B T \ln 2$	OFF-branch calorimetric baseline + leakage budget
Detector bandwidth	$f_{\text{mod}} \ll 1/(2\pi\tau_{\text{th}})$	Heater-pulse impulse response fit

14. Discussion

14.1. Relation to Prior Work

Bérut *et al.* [14] verified Landauer's principle for classical bit erasure. Our proposal extends this to *quantum* measurement, where the relevant irreversibility is the decoherence and stabilization of a classical record. The key distinction is operational: in classical systems, a pre-existing bit is erased; in record-forming quantum measurements, classical information is *created* by irreversibly distinguishing quantum alternatives into a stable record register. For pure-state binary outcomes these scales coincide numerically at $k_B T \ln 2$ per bit; for mixed or partially decohered states they need not, and the bookkeeping differs from classical reset.

Latune and Elouard [8] show that measurement can approach thermodynamic reversibility when no stable record is created; in our taxonomy this corresponds to Stage 1 alone. Mohammady and Buscemi [9] prove that efficient projective measurements are incompatible with the second and third laws, consistent with the irreversible dissipation we locate at Stage 2. Neither work proposes a concrete experimental test or provides a concise operational-conditions checklist for Landauer applicability to measurement. Our contribution fills these gaps: a matched ON/OFF architecture that operationally isolates record stabilization from reversible premeasurement, explicit conditions (C1 to C6), and calorimetric falsification criteria. To our knowledge, this is among the first proposals to combine a boundary-event framing of quantum record formation with a falsifiable differential calorimetric test and an explicit operational-conditions framework.

14.2. Applicability Summary

For a compact stage-by-stage applicability matrix, see Table 5.

Table 5. Landauer bound applicability (assumes C1 to C6).

Stage	Rev.?	$I(X;Y)?$	Landauer?
1 (Premeas.)	Yes	No	No
2 (Record)	No	Yes	Yes
3 (Reset)	No	Erased	Yes

14.3. What the Deep-Quantum Test Does and Does Not Demonstrate

The absolute magnitude of ΔQ in the deep-quantum regime ($h\nu \gg k_B T \ln 2$) does not by itself discriminate record-formation thermodynamics from generic photon absorption: a skeptic could argue that routing a photon to an absorber trivially satisfies the heat bound because the pointer energy is large. The *discriminating* content is the **temporal correlation** between heat onset and reversibility loss (Control 3), which is not predicted by energy conservation alone. A positive Phase 1 result therefore establishes that (i) the ON/OFF architecture achieves common-mode rejection at the design level, (ii) the onset of irreversible heat is temporally locked to the independently certified boundary event ($\mathcal{F} < \mathcal{F}_*$), and (iii) the null-test program closes all identified systematic channels.

14.4. Scope and Limitations

We do not claim to resolve the quantum measurement problem in the interpretational sense (e.g., wave-function collapse versus many-worlds). Our contribution is an experimentally testable *operational criterion*: objective record formation under C1 to C6 produces a measurable dissipative signature. This complements interpretational frameworks with a physical observable but does not adjudicate between them.

14.5. Roadmap: High-SNR Demonstration Versus Near-Floor Tests

We distinguish two experimentally relevant regimes. In a deep-quantum operating point ($h\nu \gg k_B T \ln 2$), the differential calorimetric signal is expected to be dominated by pointer-energy thermalization in the ON branch, so the inequality is satisfied with a large positive residual. This regime enables a high-SNR demonstration of (i) common-mode rejection in the ON/OFF architecture, (ii) the temporal link between reversibility loss and heat onset via the reversal-delay sweep, and (iii) closure of the leakage and null-test program. A near-floor residual test, in which $r = \Delta Q - k_B T \ln 2 \cdot I(X;Y)$ is small in magnitude, requires substantially increased averaging and/or operating points in which the net irreversible dissipation associated with record stabilization approaches $k_B T \ln 2$. Concrete routes include: (i) lower-frequency pointer modes ($\nu \sim 0.5$ to 1 GHz), reducing $h\nu$ toward $k_B T \ln 2$; (ii) elevated operating temperatures ($T \sim 50$ to 100 mK), increasing the Landauer floor while retaining superconducting operation; a temperature sweep would independently confirm that the residual floor scales as T , distinguishing thermodynamic content from hardware crosstalk; (iii) energy-encoded (“bright/dark”) pointer implementations in which the deposited energy is itself outcome-conditioned, making the calorimetric record variable directly interpretable as a shot-resolved outcome register; (iv) detector parallelization and improved bolometer bandwidth ($\tau_{\text{th}} < 1 \mu\text{s}$) to reduce integration times from days to hours at the Landauer scale. These extensions do not require changes to the ON/OFF differential architecture or the C1 to C6 framework.

15. Conclusion

We have provided a unified operational theory-experiment package that frames quantum measurement as a **boundary event**: a physically identifiable, irreversible transition from quantum correlation to objective classical record:

1. A three-stage taxonomy separating reversible premeasurement (Stage 1), irreversible record formation (Stage 2, the boundary event), and memory reset (Stage 3), with six explicit opera-

tional conditions (C1 to C6) specifying when the Landauer bound applies, operationalizing and extending prior stage decompositions [8].

2. A conditional record-formation heat bound $\langle Q_{\text{rec}} \rangle \geq k_B T \ln 2 \cdot I(X; Y)$, which follows from the generalized second law and is anchored by an explicit system/pointer/bath model. The model locates dissipation at record formation (Stage 2, environmental coupling), not at premeasurement (Stage 1, unitary correlation), with a protocol-dependent heat/work split made explicit by C1 to C6.
3. The first, to our knowledge, matched ON/OFF differential microcalorimetry experiment designed to isolate record-formation heat, with the temporal coincidence of reversibility loss and heat onset (Control 3) as the primary boundary-event signature.

Table 5 summarizes which stage carries irreversibility and when $I(X; Y)$ is created/erased under conditions C1 to C6.

The protocol uses demonstrated thermal-detector qubit readout and is compatible with existing circuit-QED infrastructure. Four primary controls discriminate the predicted signal from backgrounds. Sensitivity analysis using nanobolometer performance shows $\text{SNR} \gtrsim 10$ is achievable at 10 mK; integration time depends on detector bandwidth and modulation rate. Landauer-scale residual tests at 10 mK remain experimentally demanding (hours to weeks, depending on detector thermal time and effective modulation rate), while deep-quantum operation with $\Delta Q \sim h\nu \gg k_B T \ln 2$ is already comfortably accessible. All computable numerical values in this manuscript are generated or cross-validated by the accompanying scripts `Simulations/simulation.py`, `Simulations/qutip_simulation.py`, and `Simulations/run_full_audit.py`.

A positive result would establish an experimentally anchored thermodynamic criterion for objective quantum measurement.

Supplementary Materials: The following supporting information can be downloaded at the website of this paper posted on [Preprints.org](https://preprints.org).

Data Availability Statement: The numerical data and figures analyzed in this study are generated by the accompanying simulation and audit scripts in the `Simulations/` directory. Derived outputs (including audit summaries and QuTiP reports) are reproducible by running `simulation.py`, `qutip_simulation.py`, and `run_full_audit.py` with the repository defaults. The public repository is available at <https://github.com/MosesRahnama/Calorimetric-Measurement-Bound>.

References

1. R. Landauer, "Irreversibility and Heat Generation in the Computing Process," *IBM J. Res. Dev.* **5**, 183 (1961).
2. C. H. Bennett, "Notes on Landauer's principle, reversible computation, and Maxwell's Demon," *Stud. Hist. Phil. Mod. Phys.* **34**, 501 (2003).
3. J. D. Norton, "Waiting for Landauer," *Stud. Hist. Phil. Mod. Phys.* **42**, 184 (2011).
4. J. Ladyman and K. Robertson, "Landauer defended: Reply to Norton," *Stud. Hist. Phil. Mod. Phys.* **48**, 210 (2014).
5. T. Sagawa and M. Ueda, "Minimal Energy Cost for Thermodynamic Information Processing: Measurement and Information Erasure," *Phys. Rev. Lett.* **102**, 250602 (2009).
6. W. H. Zurek, "Quantum Darwinism," *Nature Phys.* **5**, 181 (2009).
7. A. Touil, B. Yan, D. Girolami, S. Deffner, and W. H. Zurek, "Quantum Darwinism, Amplification, and the Origin of Objective Classical Reality," *Phys. Rev. Lett.* **128**, 010401 (2022).
8. C. L. Latune and C. Elouard, "A thermodynamically consistent approach to the energy costs of quantum measurements," *Quantum* **9**, 1614 (2025).
9. M. H. Mohammady and F. Buscemi, "The thermodynamic trilemma of efficient measurements," *Quantum* **9**, 1250 (2025).
10. N. Katz *et al.*, "Reversal of the weak measurement of a quantum state in a superconducting phase qubit," *Phys. Rev. Lett.* **101**, 200401 (2008).
11. A. M. Gunyhó *et al.*, "Single-shot readout of a superconducting qubit using a thermal detector," *Nat. Electron.* **7**, 288 (2024).

12. R. Kokkoniemi *et al.*, "Nanobolometer with ultralow noise equivalent power," *Commun. Phys.* **2**, 124 (2019).
13. J. P. Pekola *et al.*, "Ultrasensitive Calorimetric Detection of Single Photons from Qubit Decay," *Phys. Rev. X* **12**, 011026 (2022).
14. A. Bérut *et al.*, "Experimental verification of Landauer's principle linking information and thermodynamics," *Nature* **483**, 187 (2012).
15. Y. Yin *et al.*, "Catch and Release of Microwave Photon States," *Phys. Rev. Lett.* **110**, 107001 (2013).
16. F. Wulschner *et al.*, "Tunable coupling of transmission-line microwave resonators mediated by an rf SQUID," *EPJ Quantum Technol.* **3**, 10 (2016).
17. J. R. Johansson, P. D. Nation, and F. Nori, "QuTiP 2: A Python framework for the dynamics of open quantum systems," *Comp. Phys. Commun.* **184**, 1234 (2013).
18. K. Jacobs, "The second law of thermodynamics and quantum feedback control: Maxwell's demon with weak measurements," *Phys. Rev. A* **80**, 012322 (2009); K. Jacobs, "Second law of thermodynamics and quantum feedback control: Maxwell's demon with weak measurements," *Phys. Rev. E* **86**, 040106(R) (2012).
19. D. Reeb and M. M. Wolf, "An improved Landauer principle with finite-size corrections," *New J. Phys.* **16**, 103011 (2014).
20. M. H. Mohammady, "Conditional work statistics of quantum measurements," *Quantum* **5**, 517 (2021).
21. A. P. M. Place *et al.*, "New material platform for superconducting transmon qubits with coherence times exceeding 0.3 milliseconds," *Nat. Commun.* **12**, 1779 (2021).
22. S. Deffner and C. Jarzynski, "Information Processing and the Second Law of Thermodynamics: An Inclusive, Hamiltonian Approach," *Phys. Rev. X* **3**, 041003 (2013).

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.