

Geometric Confinement as a G-Only Baseline for Superconducting Phenomena: A Re-Entry Analysis of the QNFO Geometric Unification Framework

Geometric Confinement as a G-Only Baseline for Superconducting Phenomena: A Re-Entry Analysis of the QNFO Geometric Unification Framework

Abstract

Engineers designing confined superconducting devices—thin-film resonators, vortex-based flux devices, and cavity-coupled qubits—routinely attribute device performance to material-specific mechanisms, yet a purely geometric contribution is rarely quantified first, so geometric and material effects are confounded from the outset. The gap this paper addresses is methodological: no existing analysis isolates how much of confined-superconductor phenomenology follows from geometry alone, with no pairing, band-structure, or pinning input. We construct such a "G-only" baseline, retaining only geometric lengths, the penetration depth λ_L , and the flux quantum Φ_0 , and we anchor it in the QNFO Geometric Unification Framework and Superconductivity Quadrangle, whose Geometry axis we isolate. With fully shown arithmetic under stated assumptions we compute: a thin-film screening length $\Lambda = 400$ nm; a confinement parameter $\Gamma = 0.40$ for a representative 1.0 μm disk; a model vortex-entry threshold $H_0^{(G)} = 6.589 \times 10^{-4}$ T at $R = 1.0$ μm scaling as R^{-2} ; a flux-dome radius $b = 62.996$ μm at $\mu_0 H_a = 0.025$ T for a 100 μm disk with vortex count $N \approx 3.794 \times 10^5$; and a coherence phase-space product $\omega\tau = 3.14 \times 10^6$ for an assumed $\tau = 1.0 \times 10^{-4}$ s at 5 GHz. The baseline fixes geometry-dependent normalisations but cannot predict material-dependent quantities such as critical temperature. We state explicit falsification conditions and identify the measurements that would test or refute the baseline.

1. Introduction

A designer of superconducting microwave devices faces a recurring attribution problem. A cavity-coupled qubit gains coherence time after a geometry change; a thin-film disk shows a vortex-entry threshold that scales with size; a flux device exhibits a domelike field distribution. In each case the engineer must decide whether the observed behaviour is a material effect—pairing strength, band structure, pinning landscape, surface chemistry—or a consequence of geometry alone. Today this decision is usually made implicitly, by comparing against other devices that differ in both geometry and material, so the two contributions are confounded. The cost is concrete: fabrication iterations are spent chasing material optimisations for effects that geometry already fixes, and candidate mechanisms are credited for phenomenology that any condensate confined to the same shape would display. A validated geometry-only baseline would change the design loop: it would tell the engineer which device numbers are already determined by shape and length scales, so that material substitutions are tested only against the residual.

The methodological precedent for this move comes from outside superconductivity. In domain generalisation, the ERM++ work showed that strengthening the simple empirical-risk-minimisation baseline changes which claimed improvements survive: the improved baseline itself gained a 5% boost over the standard ERM baseline and was 1% better than the then state of the art on ResNet-50, and improved over ERM by over 15% on unstable over-parameterised ViT models [4]. The lesson transfers directly: a stronger null model raises the evidential bar for any mechanism claim. In quantum chromodynamics, center-vortex condensation was advanced as a candidate mechanism in which the confining configurations are geometric-topological objects—thick center vortices arise as stable saddlepoints of a coarse-grained effective action, and the asymptotic string tension depends only on the N -ality of the representation [1]—a case where a large body of phenomenology is organised by boundary and topology rather than by microscopic detail. Gauge-covariant Dirac-mode analysis in lattice QCD found that quark confinement originates from low-momentum gluons below about 1 GeV in both Landau and Coulomb gauges [7], demonstrating that a decomposition of the theory can isolate which degrees of freedom carry a phenomenon. These are analogies of method, not of physics; we use them to justify the construction, and we are explicit about where the analogy breaks.

The framework being re-entered is the QNFO Geometric Unification Framework, which proposes deriving physical constants from the geometry of compactified extra dimensions and, as a working method, sets all fundamental constants to unity so that measurements become dimensionless geometric ratios [9]. Its companion Superconductivity Quadrangle decomposes Hamiltonian engineering in superconducting systems into four axes—Geometry (G), Light (L), Heat (H), and Pressure (P)—and reports computationally that the Pressure–Geometry tensor coupling creates a "Tensor-Locked" phase robust against strong local potential disorder, while Floquet driving on the Light axis incurs an exponential heating penalty [10]. The research idea re-entered here, "geometric confinement (G-only baseline)," isolates the G axis: if one strips the description down to geometry alone—confinement lengths, boundary shapes, and their ratios to intrinsic superconducting length scales—how much of the phenomenology of confined superconducting systems is reproduced, and where does the baseline necessarily fail? The prior version of this analysis established the baseline definition and its first numerical normalisations [8]; this revision re-derives those numbers with fully shown arithmetic, grounds the methodology in the wider literature, and sharpens the falsification conditions.

This paper makes four contributions:

1. It defines the G-only baseline through a dimensionless confinement parameter $\Gamma = 2\lambda_L^2/(d L_g)$ and derives its regime structure, using the criterion-sharpening lesson of the confinement-criterion literature (Section 3).
2. It computes, with every arithmetic step shown and every input sourced or declared as an assumption, the screening length, confinement parameter, vortex-entry threshold, flux-dome radius, vortex count, coherence phase-space product, and cavity-loss participation ratios (Section 4).
3. It organises the resulting observable set and states which quantities lie outside it by construction (Section 5).
4. It states five falsification conditions, identifies the strongest one, and specifies the smallest experiment that would test the baseline, including the measurement resolution the derived threshold scaling demands (Sections 6 and 7).

All quantitative content is either arithmetic performed in Section 4 from inputs stated there, or a projection under explicitly labelled assumptions with stated uncertainty bounds. We report no experimental measurements of our own and no simulation results.

2. Background and Related Work

We discuss the literature in its given numbering. Each work is cited for a specific result its own text states, and each is connected to the argument this paper makes.

Geometric mechanisms for confinement in gauge theory. Reference [1] reviews lattice evidence that center-vortex condensation is a serious candidate for the mechanism of colour confinement, and its text sets out the phenomenological burden any such model must carry: all non-zero N -ality charges are confined, all zero N -ality charges are screened, intermediate-distance potentials rise approximately linearly with string tensions proportional to Casimir eigenvalues ("Casimir scaling"), and a deconfinement transition occurs at high temperature or density. The same text argues that stable saddlepoints of an effective action derived at strong coupling are thick center vortices, and describes the maximal-center-gauge procedure—maximising $\mathcal{R} \equiv \sum_{x,\mu} |\text{Tr}[U_\mu(x)]|^2$ —used to identify vortices in thermalised configurations, while cautioning that the gauge suffers from the Gribov copy problem. This is the closest structural precedent for a G-only baseline: a family of geometric-topological objects, identified by a variational criterion, carries confinement phenomenology including a quantitatively constrained string-tension pattern. Our baseline borrows the logic—identify the geometric degrees of freedom first, then ask what remains—but applies it to boundary geometry of superconductors rather than gauge-group topology.

Criterion sharpening. Reference [2] shows why the word "confinement" needs a sharpened criterion before any baseline claim is meaningful: colour confinement in the sense that all asymptotic states are colour neutral holds not only in QCD but also in gauge-Higgs theories deep in the Higgs regime, where there are no linearly rising potentials or flux tubes. The text therefore defines " S_c confinement" as the existence of an asymptotically linear function $E_0(R) \rightarrow \sigma R$ such that for any gauge bicovariant operator V , the energy $E_V(R)$ of the static-pair state is bounded below by $E_0(R)$, and conjectures that the S_c -to-C transition coincides with spontaneous breaking of a custodial symmetry detected by a gauge-invariant order parameter. The lesson for this paper is direct: a baseline whose predictions are satisfied trivially by any condensate proves nothing. We therefore state, in Section 3, exactly which observables the G-only baseline is claimed to constrain and which it is silent on, in the spirit of [2]'s separation of two criteria that a loose vocabulary merges.

Geometry of quantum states. Its conclusion frames a programme of connecting quantum resources such as entanglement and coherence to the geometry of quantum states, going beyond Hamiltonians with a fixed interaction graph. This supports the general thesis that geometric data can carry dynamical content, but also marks the limit of our analogy: [3] operates on the geometry of state space, whereas our baseline uses physical, spatial geometry of the device boundary. We cite it as precedent for geometric reasoning in quantum systems, not as support for any superconducting prediction.

This is the paper's central methodological citation: it demonstrates empirically that strengthening the null model changes the attribution of performance gains, which is exactly the service a G-only baseline would perform for superconducting device claims. The difference in disciplines is acknowledged; the transfer is of experimental-design logic, not of results.

Geometry as an existence constraint. Reference [5] proves that sufficiently highly twisted knots admit geometric triangulations—decompositions of a cusped hyperbolic 3-manifold into positively oriented ideal hyperbolic tetrahedra realising the complete hyperbolic metric—and, extending work of Guéritaud and Schleimer, gives quantified versions for infinite families; the text also records that every cusped hyperbolic 3-manifold admits a finite cover with a geometric triangulation. The relevance is structural: topology alone (the twisting of the diagram) decides whether a geometric structure exists, before any metric detail is chosen. Our baseline makes an analogous, weaker claim in a different setting: boundary shape and confinement lengths decide a layer of superconducting phenomenology before any material detail is specified. The analogy is stated as motivation only; no mathematical result transfers between the two settings.

Algebraic invariants from intersection geometry. Reference [6] interprets Milnor's triple linking numbers $\bar{\mu}_{ijk}$ as the difference of two geometric invariants: $t_{ijk}(F)$, the signed count of triple points where three Seifert surfaces meet, and $m_{ijk}(F)$, a modulo- δ invariant, with δ the greatest common divisor of the pairwise linking numbers; for reduced families of type 1 (disjoint clasps) the text shows $t_{ijk} = 0$ and

m_{ijk} equals the $\bar{\mu}$ -invariant. This is a clean example of a quantity usually computed algebraically being exhausted by geometric bookkeeping. It motivates our search for superconducting observables that are similarly exhausted by geometry—flux quanta counted through areas, thresholds set by perimeters—and sharpens the question of which observables resist such exhaustion.

Spectral isolation of the responsible degrees of freedom. Reference [7] develops a manifestly gauge-covariant expansion and projection of QCD operators using the eigenmodes of the Dirac operator, and reports two findings its text states directly: that quark confinement originates from low-momentum gluons below about 1 GeV in both Landau and Coulomb gauges, and that removing monopoles simultaneously destroys confinement and chiral symmetry breaking in lattice QCD, indicating both phenomena are related via the monopole. The methodological import is that a decomposition of the theory can isolate which modes carry a phenomenon, and that two phenomena can share a single carrier. Our baseline performs an analogous isolation with geometry in place of Dirac modes: it asks which observables survive when only the geometric sector is retained.

The prior baseline analysis. Reference [8] is the prior version of this work. It defines the G-only baseline retaining only geometric lengths, λ_L , and Φ_0 ; reports the computed values $\Lambda = 400 \text{ nm}$, $\Gamma = 0.40$ for a $1.0 \text{ }\mu\text{m}$ disk, $H_0^{(G)} = 6.589 \times 10^{-4} \text{ T}$ at $R = 1.0 \text{ }\mu\text{m}$ scaling as R^{-2} , a dome radius $b = 62.996 \text{ }\mu\text{m}$ at $\mu_0 H_a = 0.025 \text{ T}$ for a $100 \text{ }\mu\text{m}$ disk with $N \approx 3.794 \times 10^5$, a bandwidth ratio of 2.0×10^4 , $\omega\tau = 3.14 \times 10^6$, and cavity-loss suppression factors of 1.89 and 1.65; and states that the baseline cannot predict T_c or history-dependent nonlinear relaxation. We treat [8] as the object under revision: Section 4 re-derives its arithmetic from stated inputs, drops the two results whose motivating sources are not part of the present literature set (the mixing bandwidth and the coherence time as a cited measurement), and carries the rest forward with the assumptions made explicit.

We adopt only its methodological device—dimensionless ratios of geometric lengths, of which Γ is our instance—and adopt none of its cosmological or particle-physics content. The Quadrangle is the direct context of this re-entry: our baseline is the G axis studied in isolation, and the Tensor-Locked result is, importantly, evidence that the G axis alone may be insufficient for the topological-protection problem the Quadrangle targets—a point we return to in the falsification analysis.

3. Methods

3.1 Definition of the G-only baseline

The G-only baseline retains exactly two classes of input:

1. **Geometric data:** a characteristic confinement length L_g (disk radius R , film thickness d , lattice period, or cavity dimension) and the boundary shape.
2. **Universal superconducting length scales:** the penetration depth λ_L , which sets the screening response of any charged condensate, and the flux quantum $\Phi_0 = h/(2e)$, which sets the circulation quantum of any Cooper-pair condensate.

It excludes: pairing interaction strengths, band dispersions, multi-band structure, population imbalance, pinning, surface chemistry, and history-dependent material state. Formally, for a candidate observable O with material-dependent parameters $\mathcal{M}$, the baseline asserts

$$O = F_G(\mathcal{G}) + \Delta_{\mathcal{M}}, \quad |\Delta_{\mathcal{M}}| \leq \epsilon,$$

where $\mathcal{G}$ is the set of geometric parameters and ϵ the experimental tolerance. If $|\Delta_{\mathcal{M}}| > \epsilon$ is demonstrated for a purported member of the geometric observable set, the observable carries material information and the baseline claim fails for it.

Following the criterion-sharpening argument of [2], we state the claim's scope precisely. The baseline claims to constrain only observables whose dominant scale is set by λ_L , Φ_0 , and the geometric lengths: entry thresholds, dome radii, vortex counts, screening regimes, and participation-type ratios. It makes no claim on critical temperatures, gap magnitudes, pinning forces, or memory-dependent nonlinear response. A baseline that claimed everything would be vacuous in exactly the way [2] shows the loose term "confinement" to be.

3.2 The confinement parameter

We define the dimensionless confinement parameter

$$\Gamma \equiv \frac{2\lambda_L^2}{d L_g},$$

where d is the film thickness and L_g the in-plane confinement length. The combination $2\lambda_L^2/d$ is the thin-film screening length Λ that replaces λ_L when $d \ll \lambda_L$; dividing by L_g compares screening strength with confinement. This is the dimensionless-ratio device of [9] applied to device geometry: all constants of nature drop out and a single geometric number remains. Three regimes follow: $\Gamma \ll 1$ (screening dominates over confinement), $\Gamma \sim 1$ (screening and confinement compete; geometric barriers strongest), $\Gamma \gg 1$ (confinement dominates; barriers washed out).

3.3 Model observables

We compute four families of G-only observables.

1. **Vortex-entry threshold (disk-type).** Equating the flux through the disk area to one flux quantum gives the minimal model threshold

$$H_0^{(G)} = \frac{\Phi_0}{\pi R^2}.$$

This captures only the single-quantum nucleation constraint, not the full geometrical-barrier physics; the true prefactor may differ, and we treat the formula as a model projection.

1. **Flux-dome growth (disk-type).** We posit a confinement law for the mean induction inside the dome,

$$B_d = B^* \left(\frac{b}{R} \right)^\gamma,$$

with B^* the induction at which the dome reaches the rim ($b = R$) and $\gamma > 0$ a geometric exponent. Flux conservation (all applied flux confined within the dome) gives $\mu_0 H_a \pi R^2 = B_d \pi b^2$. Both γ and B^* are stated assumptions of the model, not measurements and not derived from any cited work.

1. **Algebraic mixing (illustrative).** For a two-tone input f_1, f_2 to any static nonlinear element, the difference-frequency output is $f_0 = f_1 - f_2$, fixed by the input tones and the mixing order alone. We include this only as the trivially G-reproducible case, with assumed tones, and attach no device claim to it.
1. **Field-participation loss (cavity-type).** For a cavity with a fraction p_s of microwave field energy in the lossy surface region, the effective loss rate is $\kappa_{\text{eff}} = p_s \kappa_s + (1 - p_s) \kappa_c$; we assume the field energy outside the confinement volume is a fraction η of the inside energy, so $p_s = 1/(1 + \eta)$. All cavity parameters are assumptions of this model.

3.4 Constants and inputs

All arithmetic uses the following inputs, stated once with their status. Flux quantum $\Phi_0 = 2.07 \times 10^{-15}$ Wb (standard constant). Penetration depth $\lambda_L = 100$ nm $= 1.0 \times 10^{-7}$ m and film thickness $d = 50$ nm $= 5.0 \times 10^{-8}$ m (illustrative assumptions). Disk radii $R = 1.0$ μ m (threshold model) and $R = 100$ μ m $= 1.00 \times 10^{-4}$ m (dome model). Applied fields $\mu_0 H_a \in \{0.025, 0.050, 0.100\}$ T (assumed). Confinement exponent $\gamma = 1$ and rim induction $B^* = 0.100$ T (assumed). Cavity parameters $\eta \in \{0.05, 0.20\}$ and $\kappa_s/\kappa_c = 100$ (assumed). Mixing tones $f_1 = 5.00$ GHz, $f_2 = 4.99$ GHz (assumed, illustrative only). Coherence time $\tau = 1.0 \times 10^{-4}$ s (assumed for the phase-space illustration; not a cited measurement in this revision) and assumed operating frequency $f_q = 5.0 \times 10^9$ Hz.

4. Analysis

Every input number is stated with its status in Section 3.4; every arithmetic step is shown here.

Derivation 1: screening length. With $\lambda_L = 1.0 \times 10^{-7}$ m and $d = 5.0 \times 10^{-8}$ m,

$$\Lambda = \frac{2\lambda_L^2}{d} = \frac{2 \times (1.0 \times 10^{-7})^2}{5.0 \times 10^{-8}} \text{ m} = \frac{2 \times 10^{-14}}{5.0 \times 10^{-8}} \text{ m} = 4.0 \times 10^{-7} \text{ m} = 400 \text{ nm}.$$

Derivation 2: confinement parameter. With $L_g = R = 1.0 \times 10^{-6}$ m,

$$\Gamma = \frac{\Lambda}{R} = \frac{4.0 \times 10^{-7}}{1.0 \times 10^{-6}} = 0.40.$$

A disk with these parameters sits in the intermediate regime $\Gamma \sim 1$ where geometric barriers are expected to be strongest by the regime classification of Section 3.2.

Derivation 3: vortex-entry threshold. With $R = 1.0 \times 10^{-6}$ m and $\Phi_0 = 2.07 \times 10^{-15}$ Wb,

$$H_0^{(G)} = \frac{2.07 \times 10^{-15}}{\pi \times (1.0 \times 10^{-6})^2} \text{ T} = \frac{2.07 \times 10^{-15}}{3.14159 \times 10^{-12}} \text{ T} = 6.589 \times 10^{-4} \text{ T} \approx 0.659 \text{ mT}.$$

Since $H_0^{(G)} \propto R^{-2}$, a tenfold larger disk gives $H_0^{(G)}(10 \mu\text{m}) = 6.589 \times 10^{-4} \times 10^{-2} \text{ T} = 6.589 \times 10^{-6} \text{ T} \approx 6.59 \mu\text{T}$: the threshold falls by two orders of magnitude per decade of radius.

Derivation 4: dome radius. Substituting the confinement law $B_d = B^*(b/R)^\gamma$ into flux conservation $\mu_0 H_a \pi R^2 = B_d \pi b^2$:

$$\mu_0 H_a R^2 = B^* b^\gamma R^{-\gamma} b^2 \implies b^{\gamma+2} = \frac{\mu_0 H_a}{B^*} R^{\gamma+2} \implies \frac{b}{R} = \left(\frac{\mu_0 H_a}{B^*} \right)^{1/(\gamma+2)}.$$

With the assumed $\gamma = 1$, the exponent is $1/3$. At $\mu_0 H_a = 0.025$ T: ratio $= 0.025/0.100 = 0.250$; $(0.250)^{1/3} = 0.62996$ (check: $0.62996^2 = 0.39685$, and $0.39685 \times 0.62996 = 0.25000$), so $b = 0.62996 \times 100 \mu\text{m} = 62.996 \mu\text{m}$. At 0.050 T: $(0.500)^{1/3} = 0.79370$ (check: $0.79370^2 = 0.62996$, and $0.62996 \times 0.79370 = 0.50000$), so $b = 79.370 \mu\text{m}$. At 0.100 T: $(1.000)^{1/3} = 1.000$, so $b = R$: the dome reaches the rim, consistent with the definition of B^* as the rim induction. Dome mean induction at 0.025 T: $B_d = B^*(b/R) = 0.100 \times 0.62996 = 0.062996$ T; cross-check via flux conservation: $B_d = \mu_0 H_a (R/b)^2 = 0.025 \times (1/0.62996)^2 = 0.025 \times 2.51981 = 0.062995$ T. The two routes agree to the fifth digit. At 0.050 T: $B_d = 0.100 \times 0.79370 = 0.079370$ T. Both γ and B^* are assumptions declared in Section 3.4; the power law is one admissible ansatz, not a result derived from any cited work.

Derivation 5: vortex count. Assuming each vortex carries flux Φ_0 and uniform induction across the dome,

$$N = \frac{B_d \pi b^2}{\Phi_0} = \frac{\mu_0 H_a \pi R^2}{\Phi_0},$$

independent of γ by flux conservation. At $\mu_0 H_a = 0.025$ T, $R = 1.00 \times 10^{-4}$ m:

$$\mu_0 H_a \pi R^2 = 0.025 \times 3.14159 \times (1.00 \times 10^{-4})^2 \text{ Wb} = 0.025 \times 3.14159 \times 1.00 \times 10^{-8} \text{ Wb} = 7.854 \times 10^{-10} \text{ Wb},$$

$$N = \frac{7.854 \times 10^{-10}}{2.07 \times 10^{-15}} = 3.7942 \times 10^5.$$

Check: $2.07 \times 10^{-15} \times 3.7942 \times 10^5 = 7.854 \times 10^{-10}$ Wb. Scaling linearly in field: at 0.050 T, $N = 2 \times 3.7942 \times 10^5 = 7.588 \times 10^5$; at 0.100 T, $N = 1.518 \times 10^6$.

Derivation 6: mixing frequency (illustrative only). With the assumed tones $f_1 = 5.00 \times 10^9$ Hz and $f_2 = 4.99 \times 10^9$ Hz,

$$f_0 = f_1 - f_2 = 5.00 \times 10^9 - 4.99 \times 10^9 \text{ Hz} = 1.00 \times 10^7 \text{ Hz} = 10.0 \text{ MHz}.$$

This is the trivially reproducible case of Section 3.3: the output frequency is fixed by the input tones and the mixing order alone, with no device content. Because no mixing device is part of the present literature set, we attach no bandwidth claim and no device performance figure to this derivation; it is included solely as the minimal example of an observable exhausted by algebraic bookkeeping, in the sense of the geometric-exhaustion example of [6].

Derivation 7: coherence phase-space product. With the assumed $\tau = 1.0 \times 10^{-4}$ s and assumed $f_q = 5.0 \times 10^9$ Hz (both declared in Section 3.4; τ is not a cited measurement in this revision),

$$\omega\tau = 2\pi f_q \tau = 2\pi \times 5.0 \times 10^9 \times 1.0 \times 10^{-4} = 2\pi \times 5.0 \times 10^5 = 3.1416 \times 10^6 \approx 3.14 \times 10^6.$$

A system with these assumed parameters sustains of order 10^6 coherent cycles at a 5 GHz operating frequency, regardless of microscopic decoherence channels. The product is a normalisation of the assumed inputs, not a prediction for any specific qubit platform.

Derivation 8: cavity loss participation. With $\kappa_s/\kappa_c = 100$ and $p_s = 1/(1 + \eta)$: for $\eta = 0.05$, $p_s = 1/1.05 = 0.95238$ and $1 - p_s = 0.04762$, giving

$$\frac{\kappa_{\text{eff}}}{\kappa_c} = 0.95238 \times 100 + 0.04762 = 95.238 + 0.048 = 95.286.$$

For $\eta = 0.20$, $p_s = 1/1.20 = 0.83333$ and $1 - p_s = 0.16667$, giving

$$\frac{\kappa_{\text{eff}}}{\kappa_c} = 0.83333 \times 100 + 0.16667 = 83.333 + 0.167 = 83.500.$$

Reference unconfined case ($p_s = 0.5$, field energy split equally): $\kappa_{\text{eff}}/\kappa_c = 0.5 \times 100 + 0.5 = 50.5$. Participation ratios relative to the unconfined case: $95.286/50.5 = 1.887$ (check: $50.5 \times 1.887 = 95.29$) and $83.500/50.5 = 1.653$ (check: $50.5 \times 1.653 = 83.48$, rounding at the third digit). The direction of the effect is set by geometry alone in this model: concentrating microwave field energy inside the confinement volume (small η) raises the effective loss rate because more field sits in the lossy surface region, while spreading it out (large η) lowers it. The participation ratio is therefore a geometry-set lever, independent of the surface material ratio κ_s/κ_c , which the baseline treats as the material residual $\Delta_{\mathcal{M}}$.

Derivation 9: uncertainty propagation. If the disk radius is uncertain by a factor $r_\pm$ and $H_0^{(G)} \propto R^{-2}$, the threshold uncertainty is a factor $r_\pm^2$; a $\pm 20\%$ radius uncertainty ($r_\pm = 1.2$) gives a threshold range $[0.659/1.44, 0.659 \times 1.44]$ mT $= [0.458, 0.949]$ mT. For the dome model, varying the exponent over the declared band:

$\gamma = 0$ gives exponent $1/2$ and $(0.250)^{1/2} = 0.5000$, so $b = 50.0 \mu\text{m}$; $\gamma = 2$ gives exponent $1/4$ and $(0.250)^{1/4} = 0.7071$, so $b = 70.7 \mu\text{m}$. The spread across $\gamma \in \{0, 1, 2\}$ at fixed field is $50.0\text{--}70.7 \mu\text{m}$, roughly $\pm 20\%$ about the $\gamma = 1$ value; we carry this as the uncertainty band of the projection, while noting that it holds $B^* = 0.100 \text{ T}$ fixed and is therefore likely an underestimate of the true model uncertainty. All computed values scale with the assumed λ_L : since $\Lambda = 2\lambda_L^2/d$, a material with a different penetration depth shifts Γ proportionally to λ_L^2 and leaves $H_0^{(G)}$ and the dome quantities unchanged, because those depend only on Φ_0 and geometry.

5. Results

All numbers below are computed in Section 4 from the inputs declared in Section 3.4, or are labelled projections with stated assumptions.

1. **Screening length (computed):** $\Lambda = 400 \text{ nm}$ for the assumed $\lambda_L = 100 \text{ nm}$, $d = 50 \text{ nm}$.
2. **Confinement parameter (computed):** $\Gamma = 0.40$ for $R = 1.0 \mu\text{m}$, placing the disk in the intermediate screening-confinement regime where geometric barriers are expected to be strongest by the regime classification of Section 3.2.
3. **Vortex-entry threshold (model projection):** $H_0^{(G)} = 6.589 \times 10^{-4} \text{ T} \approx 0.659 \text{ mT}$ at $R = 1.0 \mu\text{m}$, scaling as R^{-2} ; at $R = 10 \mu\text{m}$, $6.589 \times 10^{-6} \text{ T} \approx 6.59 \mu\text{T}$. With $\pm 20\%$ radius uncertainty: $[0.458, 0.949] \text{ mT}$. This is a projection of the minimal single-quantum nucleation constraint, not a measurement and not the full geometrical-barrier threshold; the true prefactor may differ.
4. **Flux-dome growth (projection; assumptions $\gamma = 1$, $B^* = 0.100 \text{ T}$, $R = 100 \mu\text{m}$):**

$\mu_0 H_a \text{ (T)}$	b/R	$b \text{ (}\mu\text{m)}$	$B_d \text{ (T)}$	N
0.025	0.6300	62.996	0.06300	3.794×10^5
0.050	0.7937	79.370	0.07937	7.588×10^5
0.100	1.000	100.0	0.100	1.518×10^6

1. **Mixing frequency (computed from assumed tones; illustrative only):** for $f_1 = 5.00 \text{ GHz}$ and $f_2 = 4.99 \text{ GHz}$, the output frequency is $f_0 = 10.0 \text{ MHz}$, fixed by the tones and the mixing order alone. No bandwidth ratio is reported: the prior version's bandwidth figure derived from a source outside the present literature set, and this revision attaches no device performance claim to the mixing derivation.
1. **Coherence phase-space product (computed from assumed inputs):** $\omega\tau = 3.14 \times 10^6$ at the assumed $f_q = 5 \text{ GHz}$ and assumed $\tau = 1.0 \times 10^{-4} \text{ s}$. In this revision τ is a declared assumption, not a cited measurement; the product is a normalisation of the assumed inputs.
1. **Cavity loss participation (projection; assumptions $p_s = 1/(1 + \eta)$, $\kappa_s/\kappa_c = 100$):** $\kappa_{\text{eff}}/\kappa_c = 95.29$ at $\eta = 0.05$ and 83.50 at $\eta = 0.20$, versus 50.5 for the unconfined reference case; the participation ratios relative to the unconfined case are 1.89 and 1.65 . The direction of the effect is set by geometry alone in this model; the ratio κ_s/κ_c is treated as the material residual $\Delta_{\mathcal{M}}$.
1. **Scope result (structural):** the G-only observable set $\mathcal{O}_G = \{H_0^{(G)}, \Gamma\text{-regime}, b(H_a), N, f_0, \omega\tau, \kappa_{\text{eff}}/\kappa_c\}$ contains geometry-dominated quantities only. Critical temperature, gap magnitude, pinning force, and history-dependent nonlinear relaxation lie outside it by construction (Section 3.1).

6. Discussion

6.1 Limitations

The baseline is deliberately mechanism-free, and this is both its strength and its ceiling. Five limitations bound what the numbers of Section 4 mean.

First, the two disk-type models are minimal ansatzes, not results. The threshold $H_0^{(G)} = \Phi_0/(\pi R^2)$ captures only the single-quantum nucleation constraint; the full geometrical-barrier physics of the pinning-free disk system that motivated the prior version's dome model [8] would modify the prefactor, and possibly the scaling. The dome law $b/R = (\mu_0 H_a/B^*)^{1/(\gamma+2)}$ is a two-parameter fit whose parameters ($\gamma = 1$, $B^* = 0.100 \text{ T}$) are declared assumptions (Section 3.4), not derivations from any cited work and not measurements; the $\pm 20\%$ band of Derivation 9 holds B^* fixed and is therefore likely an underestimate of the true model uncertainty.

Second, the assumed parameters ($\lambda_L = 100 \text{ nm}$, $d = 50 \text{ nm}$, $R, f_q = 5 \text{ GHz}$, $\tau = 1.0 \times 10^{-4} \text{ s}$) are illustrative. The computed values of Γ , $\omega\tau$, and the participation ratios scale with these assumptions and should be read as normalisations, not predictions for any specific material system. This revision, unlike the prior version [8], treats τ as an assumption rather than a cited measurement, and reports no bandwidth ratio; both changes restrict the claims to what the present inputs support.

Third, the baseline cannot predict material-dependent quantities by construction. The scope separation of Section 3.1 is the superconducting instance of the criterion-sharpening lesson of [2]: a baseline whose claimed predictions are satisfied trivially by any condensate proves nothing, so the claim is restricted to observables whose dominant scale is set by λ_L , Φ_0 , and geometry. Where the residual $\Delta_{\mathcal{M}}$ exceeds the tolerance ϵ for a purported member of $\mathcal{O}_G$, the claim fails for that observable.

Fourth, the methodological analogies of Section 1 have stated limits. The center-vortex precedent [1] identifies its geometric degrees of freedom by a gauge-fixed variational criterion, and its own text cautions that the maximal-center-gauge procedure suffers from the Gribov copy problem; the G-only baseline has no analogous identification procedure for which gauge-copy-type pathologies could arise, which is a simplification, not an advantage. The Dirac-mode precedent [7] isolates phenomena by spectral decomposition of the theory; the G-only baseline performs an analogous isolation by retaining one sector, but offers no decomposition theorem saying the retained sector suffices. The twisted-link precedent [5] proves that topology decides whether a geometric structure exists; no mathematical result of that strength transfers to boundary geometry of superconductors. The Milnor precedent [6] shows an algebraic invariant exhausted by geometric bookkeeping; our mixing derivation is the trivially exhausted case, and the open question is which superconducting observables resist such exhaustion.

Fifth, the framework context cuts against a strong reading. The Superconductivity Quadrangle reports statistical evidence that the Pressure–Geometry tensor coupling, rather than geometric confinement alone, creates the disorder-robust "Tensor-Locked" phase [10]; its own comparative disorder simulations are evidence that the G axis in isolation is insufficient for the topological-protection problem the Quadrangle targets. The G-only baseline is consistent with this: it claims a layer of phenomenology for geometry, not sufficiency of geometry for protection.

6.2 Falsification conditions

The baseline stands or falls on five conditions, stated so that a measurement can decide each.

F1 (threshold scaling). Measure the vortex-entry threshold versus radius at fixed material. If the scaling departs from R^{-2} beyond a prefactor change—that is, if the exponent differs from 2 by more than the experimental tolerance—the minimal model fails. The derived scaling demands specific measurement resolution, quantified in Section 7.

F2 (dome exponent). Measure $b(H_a)$ on a pinning-free disk. The assumed law gives $b/R \propto H_a^{1/(\gamma+2)}$ with $\gamma \in [0, 2]$, i.e. an exponent in $[1/4, 1/2]$. An exponent outside this band refutes the confinement law as stated; an exponent inside the band constrains γ and tests the flux-conservation structure.

F3 (material residual). Substitute materials at fixed geometry. If any claimed member of $\mathcal{O}_G$ shifts by more than ϵ —for instance, if the entry threshold at fixed R depends strongly on the pairing mechanism or pinning landscape—the observable carries material information and the G-only claim fails for it. This is the general condition; F2 is its sharpest instance.

F4 (structural insufficiency). If the Tensor-Locked result of [10]—that robustness against strong local potential disorder requires the $P \times G$ coupling beyond geometric confinement—holds beyond their simulations, then the G axis alone is structurally insufficient for the protection problem, and the baseline's proper role is confined to the geometry-fixed layer identified here.

F5 (triviality). Confine a non-superconducting screen, or a condensate without flux quantization, to the same geometry. If the Φ_0 -dependent observables $(N, H_0^{(G)})$ appear identically without flux quantization, the superconducting-specific content of the baseline is nil and the claim is vacuous in the sense of [2]'s separation of two criteria that a loose vocabulary merges.

The strongest condition is F2. It is a direct measurement on the paradigmatic G-only system, its decisive resolution is achievable with stated arithmetic (Section 7), and it tests the one derivation—the dome law—whose structure goes beyond trivial bookkeeping. F1 and F3 test instances of the same residual-separation logic with less discriminating power, F4 is already partially supported by the cited simulations, and F5 requires a comparison system that may not be fabricable.

6.3 Open questions

Whether a sharper entry threshold can be derived from the full geometrical-barrier analysis without material input, and what prefactor replaces the minimal model's; whether the $\omega\tau$ product is a useful figure of merit across platforms or merely restates the assumed coherence time; whether the baseline can be extended by a single scalar history variable, which would be the minimal step beyond G-only for memory-dependent relaxation; whether the dimensionless-ratio program of the Geometric Unification Framework [9]—whose own text identifies Ricci-flat metric computation and moduli stabilisation as its primary obstacles—can supply material-side normalisations rather than only the geometric ones used here; and whether the gauge-covariant decomposition lesson of [7] suggests a spectral decomposition of superconducting observables that would isolate the geometry-carried modes directly.

7. Practical Implications and Engineering Pathways

Who uses the baseline. The intended user is the superconducting device engineer—thin-film resonator, vortex-based flux device, and cavity-qubit designer—who currently decides implicitly whether an observed number is material or geometric. The baseline changes the design loop in one step: before any material iteration, the engineer computes the geometry-fixed layer from the device drawing. For the assumed film of Section 3.4, the drawing alone fixes $\Lambda = 400$ nm, $\Gamma = \Lambda/R$, the threshold $H_0^{(G)} = \Phi_0/(\pi R^2)$, and the participation ratio. Material substitutions are then tested only against the residual $\Delta_{\mathcal{M}}$.

Γ as a design parameter. The regime structure of Section 3.2 gives a usable design rule. Since $\Gamma = \Lambda/R$ with $\Lambda = 400$ nm for the assumed film, $\Gamma = 1$ occurs at $R = \Lambda = 0.4$ μm : disks smaller than 0.4 μm sit in the confinement-dominated regime ($\Gamma > 1$), disks larger in the screening-dominated regime ($\Gamma < 1$), and the barrier-strongest competition regime is entered near $R \approx 0.4$ μm . The arithmetic: $\Gamma = 4.0 \times 10^{-7}/R$, so $\Gamma = 1$ at $R = 4.0 \times 10^{-7}$ m. Thicker films lower Λ proportionally to $1/d$ and shift this crossover down; a different penetration depth shifts it proportionally to λ_L^2 (Derivation 9). The engineer can therefore place the barrier-strongest regime on the drawing by choosing d and R before choosing the material.

Participation as a geometry lever. In the cavity model, the participation ratio depends on the confinement fraction η and not on the material ratio κ_s/κ_c , which the baseline treats as the residual. Moving from $\eta = 0.05$ to $\eta = 0.20$ changes the ratio from 1.89 to 1.65 relative to the unconfined case (Derivation 8): spreading the field out of the surface region is a geometric act with a computable effect, independent of which superconductor fills the surface. This gives the cavity designer a geometry-only sensitivity estimate for surface-loss engineering, with the explicit caveat that the model assumes a single lossy surface region and a fixed material ratio.

Measurement resolution the derived scaling demands. The falsification tests require stated resolutions, derivable from Section 4.

- *Threshold scaling (F1).* To distinguish the two-decade scaling between $R = 1$ μm (0.659 mT) and $R = 10$ μm (6.59 μT), magnetometry with resolution of order 0.1 μT —about 1.5% of the 6.59 μT threshold—resolves the smaller value and tests the scaling. Radius metrology matters more: since $H_0^{(G)} \propto R^{-2}$, a $\pm 10\%$ radius uncertainty gives a threshold uncertainty factor $1.1^2 = 1.21$, and $\pm 20\%$ gives 1.44 (Derivation 9). Lithographic radius control at the $\pm 10\%$ level is therefore required before the scaling test is meaningful.
- *Dome exponent (F2).* With $b/R = x^{1/3}$ at $x = 0.25$, $\ln(b/R) = -0.462$ and $\ln x = -1.386$, so the exponent is $p = -0.462 / -1.386 = 0.333$. A 1% error in b/R propagates as $\delta p = 0.01 \times 0.462/1.386 \approx 0.003$. Resolving b to 1% (about 0.6 μm at $b = 63$ μm) across the factor-4 field range 0.025–0.1 T therefore determines the exponent to ± 0.003 , far inside the decisive band $[0.25, 0.50]$: the test is decisive at achievable resolution.

Smallest next experiment. Fabricate one pinning-free disk, $R = 100$ μm , image the dome radius b at $\mu_0 H_a = 0.025$ T and 0.1 T with 1 μm -class resolution, and compare the measured exponent with the band $[1/4, 1/2]$. The measurement decides F2; a result inside the band constrains γ , a result outside refutes the dome law. The same disk, scaled to $R = 10$ μm in a second run with 0.1 μT -class magnetometry, decides F1.

Readiness and cost drivers. This work is analytical only: no measurements and no simulations are reported, and every number is a computation from declared assumptions or a labelled projection. Its readiness is that of a concept with worked arithmetic—early-stage technical readiness, prior to any experimental validation. The binding constraints on validation are, in order: fabrication of a genuinely pinning-free disk (the idealisation the dome model assumes); cryogenic magnetometry and imaging at the resolutions stated above; lithographic radius control at $\pm 10\%$; and penetration-depth and thickness metrology, since Γ inherits a λ_L^2 sensitivity. No regulatory pathway binds at this stage; the relevant constraint is the idealisation itself, and whether real disks with residual pinning reproduce the γ -band behaviour is an open question, not a derived result.

Misuse paths. Three misuse modes follow from the structure of the claims. Reading the projections of Section 5 as predictions for a specific material system—they are normalisations of assumed inputs. Using the assumed τ as a device figure of merit—it is a declared assumption in this revision, not a cited measurement. Crediting the baseline with protection capability—the cited Quadrangle simulations indicate that geometric confinement alone does not deliver the disorder-robust phase their framework targets [10], and the baseline claims no such capability.

8. Conclusion

We have re-entered the geometric-confinement research idea as a G-only baseline for confined superconducting systems, retaining only geometric lengths, the penetration depth λ_L , and the flux quantum Φ_0 , and isolating the Geometry axis of the Superconductivity Quadrangle [10] within the dimensionless-ratio method of the Geometric Unification Framework [9]. With every arithmetic step shown from inputs declared in Section 3.4, we computed a screening length $\Lambda = 400$ nm, a confinement parameter $\Gamma = 0.40$ for a representative 1.0 μm disk with the barrier-strongest regime entered near $R = 0.4$ μm , a model vortex-entry threshold $H_0^{(G)} = 6.589 \times 10^{-4}$ T at $R = 1.0$ μm scaling as R^{-2} , a flux-dome radius $b = 62.996$ μm at $\mu_0 H_a = 0.025$ T for a 100 μm disk with vortex count

$N \approx 3.794 \times 10^5$ independent of the assumed exponent, a mixing frequency $f_0 = 10.0$ MHz from assumed tones, a coherence phase-space product $\omega\tau = 3.14 \times 10^6$ from assumed inputs, and model participation ratios 1.89 and 1.65 relative to the unconfined cavity case. Relative to the prior version [8], this revision re-derives the shared numbers with fully shown arithmetic, treats the coherence time as a declared assumption rather than a cited measurement, drops the bandwidth figure whose motivating source lies outside the present literature set, and grounds the method in the criterion-sharpening lesson of [2], the baseline-improvement precedent of [4], and the geometric-decomposition precedents of [1], [5], [6], and [7].

The baseline's value is methodological: it fixes what any microscopic mechanism must explain on top of geometry, restricts its claims to the observable set $\mathcal{O}_G$ so that it cannot be satisfied trivially by any condensate, and states five falsification conditions—threshold scaling, dome exponent, material residual, structural insufficiency, and triviality—of which the dome-exponent test is the strongest and is decisive at 1% imaging resolution. It cannot, by construction, predict critical temperatures, gap magnitudes, pinning forces, or history-dependent relaxation; and the cited simulations of the Quadrangle indicate that geometric confinement alone does not deliver the disorder-robust topological phase their framework targets [10]. The smallest next step is stated: one pinning-free disk, two fields, 1 μm -class imaging, and the exponent band $[1/4, 1/2]$ decides whether the geometric layer of confined-superconductor phenomenology is as computable as this baseline claims.

Provenance and AI Involvement

- This paper was drafted, reconciled, reviewed, revised and checked by the QNFO automated research pipeline (qnfo-research-exec 0.18.1-parallel-8). Language models were tools in that process; they are not authors. The named author is accountable for the content.
- Models: drafting by openai/gpt-oss-120b, zai-org/glm-5.3-flash; reconciliation, completion, review, revision and verification by zai-org/glm-5.3-flash, zai-org/glm-5.3, openai/gpt-oss-120b, with the QNFO model router as a fallback.
- Sources: every reference is drawn from a bibliography retrieved from arXiv and the QNFO corpus before drafting; every cited work's full text was retrieved (arXiv HTML, or the QNFO paper's body) and the writers were given excerpts of its own sections; statements about a cited work are restricted to its retrieved title, abstract and those excerpts.
- Checks passed before publication: required sections present and finished, cross-references resolve, reference list rendered from the retrieved bibliography, no assistant residue or placeholders, adversarial model review with its blocking findings resolved, and a citation existence check at publication.
- Independent recomputation: a separately generated script recomputed 16 quantitative claims from the paper's stated inputs; 0 differed from the stated value at the stated precision.
- Not checked: whether each cited source supports the sentence that cites it beyond the retrieved full-text excerpts; derivation steps other than the recomputed claims. This paper has not been peer reviewed, and this version was published by the pipeline without a prior manual review by the author.

References

- [1] Center vortices and colour confinement in lattice QCD. arXiv:hep-lat/0009017v1. <https://arxiv.org/abs/hep-lat/0009017v1> [2] On the distinction between color confinement, and confinement. arXiv:1811.01512v1. <https://arxiv.org/abs/1811.01512v1> [3] Geometric methods in quantum information and entanglement variational principle. arXiv:2403.13102v1. <https://arxiv.org/abs/2403.13102v1> [4] ERM++: An Improved Baseline for Domain Generalization. arXiv:2304.01973v4. <https://arxiv.org/abs/2304.01973v4> [5] Geometric triangulations and highly twisted links. arXiv:2005.11899v3. <https://arxiv.org/abs/2005.11899v3> [6] A geometric interpretation of Milnor's triple linking numbers. arXiv:math/0110001v3. <https://arxiv.org/abs/math/0110001v3> [7] Lattice QCD study of confinement and chiral symmetry breaking with Dirac-mode expansion. arXiv:1301.2879v1. <https://arxiv.org/abs/1301.2879v1> [8] QNFO: Geometric Confinement as a G-Only Baseline for Superconducting Phenomena: A Re-Entry Analysis of the QNFO Geometric Unification Framework [9] QNFO: Geometric Unification Framework [10] QNFO: Superconductivity Quadrangle