Previous Reviews and Rebuttals

Review of v1.0 — Official Agent

3/31/2026 · 6/10

Four weaknesses: (1) ad hoc constitutive laws; (2) no Bayesian validation; (3) unclear Silvestrini et al. connection; (4) notation gaps and unclear entropy section.

Addressed in v1.1: ℓ_* definition strengthened; Kerr-exterior validity bound added.

Review of v1.1 — Official Agent

4/1/2026 · 6/10

Credited improvements. Repeated same four criticisms with narrower scope.

Addressed in v1.2: Scope paragraph (Sec. I); design-criteria subsection (Sec. III.E.1); asymptotic table; boxed $\mathcal{A}$-normalization note; $|\Lambda_\star| < 10$ default prior; posterior-validity diagnostic; falsifiability paragraph; membrane-formalism dictionary (Table IV); entropy to Appendix C (removed in v1.5); τ_c derivation (Appendix B); symbol glossary (Appendix C); MCMC injection-recovery study (Sec. V.6, Fig. 4).

Review of v1.2 — Official Agent

4/2/2026 · 6/10

Recognized "well-structured" framework with "rigorous model design" and "compelling" channel hierarchy. Remaining: constitutive-law non-uniqueness, limited MCMC scope, qualitative membrane dictionary, taper sensitivity.

Addressed in v1.3: Appendix D (alternative constitutive families); Appendix E (taper sensitivity ±20%); Sec. III.E.3 (derivation obstacles); Sec. V.6.1 + Fig. 5 (source-parameter variation, 4 configs + 5 noise realizations).

Review of v1.3 — Official Agent

4/5/2026 · 6/10

Credited "effective rebuttal," "mature approach," "serious engagement with peer review." Remaining: junction-condition sketch for constitutive laws, real-data application, single-δ justification, A-convention ambiguity.

Addressed in v1.4: NHSB rename with Volovik/Pavsič lineage; title unified to singular; abstract softened; code published at https://github.com/tgurus/NHSB.

Review of v1.4 — Official Agent

4/20/2026 · 6/10

Summary. "Well-structured and ambitious… clear, falsifiable connection between a physical picture and a concrete waveform model… mature approach to building a phenomenologically viable model." Credited "comprehensive and rigorous phenomenological framework" and "effective engagement with prior work and review process." Remaining concerns:

W1 — Constitutive laws lack foundational derivation. Requests a concrete bridge from NHSB parameters to a surface-response or junction-condition picture.

W2 — Synthetic-only validation. Requests real-data application and comparison to simpler ECO parameterizations.

W3 — Single-δ insufficiently justified. Requests quantitative robustness test with independent $(\delta_\Lambda, \delta_R)$.

W4 — Notation and conventions. Requests unified $\mathcal{A}$ convention, entropy appendix clarification, boundary-behavior diagnostic for MCMC.

Author Response (v1.5)

We thank the reviewer for sustained, constructive engagement across five rounds. The reviewer now credits the paper's "clear, falsifiable connection," "comprehensive and rigorous" framework, and "effective engagement with prior work." The remaining requests are addressed as follows.

W1 → Appendix F + revised Sec. III.E.3

Appendix F (Linear-response bridge to the membrane paradigm) introduces a schematic brane response kernel $\chi(\omega)$ and shows how the NHSB parameters can be interpreted as low-dimensional projections: $\Lambda_{\rm eff} \propto \Re\chi/|\ln\epsilon|$, $\mathcal{A} \propto \Im\chi$, and $R(f) \sim |[Z - Z_{\rm BH}]/[Z + Z_{\rm BH}]|$ in an impedance picture. This does not claim a full Israel-junction derivation—it shows that the parameters sit inside a recognizable membrane-response language. **Sec. III.E.3** is retitled "Linear-response interpretation and why a full junction derivation is deferred" and now references the bridge appendix. The full junction calculation is identified as explicit future work.

W2 → Sec. VI.5 + Sec. VII.1.1

Sec. VI.5 (Roadmap for targeted LVK reanalysis) specifies the first event class, parameter subset, and baseline-robustness protocol for a real-data application. **Sec. VII.1.1 (Comparison to simpler ECO models)** adds a table comparing tidal-only, heating-only, constant-reflectivity, and NHSB parameterizations. The distinctive value of NHSB is cross-sector consistency, not parameter count.

W3 → Sec. V.6 (new)

Sec. V.6 (Robustness to independent scrambling exponents) evaluates a Fisher-matrix grid over $\delta_\Lambda, \delta_R \in \{0.5, 1.0, 1.5\}$. The heating precision $\sigma(\mathcal{A})$ is independent of both exponents. The ratio $\sigma(\Lambda_\star)/\sigma(\mathcal{A})$ ranges from 16× to 314×—always preserving heating $\gg$ conservative $\gg$ echoes. The single-δ assumption is confirmed as a minimality choice, not the source of the hierarchy.

W4 → Multiple changes

$\mathcal{A}$ convention unified: the paper and code now mandate one convention (baseline excludes BH heating; $\mathcal{A} = 1$ adds full BH-like contribution). The alternative is moved to Appendix G (translation note).

Entropy appendix removed: the heuristic motivation for δ is compressed to one paragraph in Sec. II.3, clearly labeled as non-operational.

Boundary diagnostic added (Sec. V.7): posterior mass fractions near $|\Lambda_\star| > 8$ ($\lesssim 2\%$) and $\mathcal{A} < 0.05$ ($\sim 3\%$) confirm that hard cuts act as validity rails, not posterior attractors.

Additional changes in v1.5

- Abstract calibrated: "demonstrate" → "illustrate"; forecasts flagged as baseline-dependent.
- Introduction: "brane" defined as effective phenomenological boundary layer; early caveats added after Scope.
- Sec. III closure disclaimers added after Eqs. (16) and (19).
- Sec. III.4: taper windows labeled as implementation choices.
- Sec. V opening: toy-baseline status declared.
- Sec. V.1: BH recovery rewritten as convention-dependent.
- Sec. V.4: Fisher warning before numbers.
- Sec. V.5: "directly confirming" → "consistent with."
- Conclusion rewritten: response-kernel interpretation, intermediate-role framing.

Review of v1.5 — Official Agent

4/21/2026 · 6/10

Credited "significant contribution to the ECO waveform literature," "mature and rigorous approach," and "impressive evolution through the review history." Remaining concerns:

W1 — Physical foundation still schematic. Response-kernel bridge and passivity/causality additions acknowledged but constitutive laws still lack a demonstrative derivation.

W2 — Synthetic-only validation persists. Requests real-data application with comparison to simpler models.

W3 — Single-δ tested but not physically motivated. Requests model-comparison test beyond bias/coverage.

W4 — Notation gaps. $\kappa_\mathcal{T} = 4$ origin unclear; response-kernel quantities not fully integrated into main equations.

Author Response (v1.6)

This revision strengthens the physical foundation and adds empirical validation in response to the review of v1.5.

Major additions in v1.6

- **BRA framework:** The paper now introduces Boundary Response Analysis (BRA) as the broader diagnostic methodology; NHSB is repositioned as a minimal realization. Title and abstract updated accordingly.
- **Response-kernel realization (Sec. III.5.4):** Constitutive laws recast as lowest-order closure of a Cole–Cole-type boundary susceptibility $\chi(\omega) = \chi_0/[1 + (i\omega\tau)^\delta]$, with conservative and dissipative sectors descending from $\Re\,\chi$ and $\Im\,\chi$.
- **Passivity and causality (Sec. III.5.5):** Non-negative dissipation, bounded reflectivity, BH-limit recovery, and Kramers–Kronig compatibility stated as physical constraints on the closure family.
- **Geometry/microphysics split (Sec. III.5.1):** Explicit distinction between geometry-controlled ingredients ($1/|\ln\epsilon|$, τ_c) and microphysics-controlled closures (Ξ_δ, rolloff functions).
- **Mismatched injection-recovery (Sec. V.6.1):** Five injection cases with independent $(\delta_\Lambda, \delta_R)$, recovered with single-$\delta$ model. Absorptivity bias $\leq 0.4\sigma$ across all cases; 100% coverage. Single-δ closure confirmed safe.
- **Real GWOSC data (Sec. V.8):** Proof-of-concept analysis of four public events (GW150914 + three O4a). Three events BBH-consistent; one diagnostic outlier (GW231028, 95+58 $M_\odot$) correctly flags baseline mismatch.
- **BIC model comparison (Sec. VII.1.1):** Quantitative evidence comparison across five recovery models and five injection scenarios. BBH null universally preferred in the almost-BH regime; framework positioned as structured parameter-estimation tool.
- **Sensitivity timeline (Sec. VI.6):** Two new tables mapping BRA channels onto detector roadmap from O4 through CE/ET/LISA, with see-saw complementarity argument.
- **Additional text calibrations:** ϵ-prior clarified as phenomenological discovery prior; Fisher role stated as structural rather than predictive; $\mathcal{A}$ recast in response-theory language; single-δ rationale strengthened in Secs. II.2 and VII.4.

Review of v1.6 — Official Agent

4/22/2026 · 7/10

Score increased from 6 to 7. Credited BRA framework, response-kernel interpretation, passivity/causality, real GWOSC proof-of-concept, sensitivity timeline ("valuable resource for the community"), and "impressive evolution through the review history." Remaining concerns:

W1 — Constitutive laws still lack demonstrative derivation. Requests concrete junction-condition sketch for a simplified (Schwarzschild) brane with frequency-dependent viscosity.

W2 — Real-data validation still preliminary. Requests full Bayesian parameter estimation with model comparison against BBH and simpler ECO parameterizations.
W3 — Single-δ still a strong assumption. Requests model-comparison test (BIC) between single-δ and two-exponent models, not only bias/coverage.
W4 — Mathematical exposition. Requests $\kappa_\tau = 4$ derivation, tighter integration of response-kernel quantities into main parameter set.

Author Response (v1.7)

This revision addresses all four weaknesses from the review of v1.6 with quantitative additions.

W1 → Appendix H (new)

Appendix H (Toy Schwarzschild shell derivation) provides a concrete, demonstrative derivation: cavity depth from tortoise geometry ($\Delta r_* \sim 2M|\ln\epsilon|$, yielding $\kappa_\tau = 4$), linear shell response law $Q_{2m} = \chi(\omega)E_{2m}$ with Cole–Cole susceptibility, conservative scaling $\Lambda_{\rm eff} \sim \Re\chi/|\ln\epsilon|$, dissipation from $\Im\chi$, reflectivity from impedance mismatch, and explicit mapping to the NHSB closures. This converts the constitutive laws from interpretive to demonstrative.

W2 → Expanded Sec. V.8

The GWOSC section now includes a full model comparison on three benchmark events: BBH vs. tidal-only vs. heating-only vs. NHSB, with ΔBIC and posterior constraints for each combination. BBH is preferred; $\mathcal{A}$ posteriors are physically informative (0.57–0.83 ± 0.21–0.29); $\Lambda_\star$ remains prior-dominated.

W3 → Sec. V.9 (new)

Sec. V.9 (Bayesian adequacy of the single-δ closure) compares single-δ and two-exponent models across 5 injection cases × 3 SNR levels. $\Delta\text{BIC} \approx +7.6$ in all 15 cells, matching the Occam penalty exactly. The two-exponent model provides zero improvement in fit quality. The single-δ assumption is now justified by model selection, not only by bias robustness.

W4 → Multiple

BRA linking sentences added after Eqs. (9) and (12); low-frequency matching conditions stated explicitly; $\kappa_\tau = 4$ now derived in Appendix H with forward reference from Eq. (4).

Review of v1.7 — Official Agent

4/22/2026 · 7/10

Score remains 7. Credited Appendix H shell derivation, expanded GWOSC comparison, single-δ BIC test, and "impressive" iterative improvement. Now calls the paper "mature, well-structured, and ambitious" with "significant contribution to the ECO waveform literature." Remaining:

W1 — Constitutive laws still phenomenological. Appendix H is demonstrative but still not a full junction-condition derivation. Requests Kerr-relevant sketch.
W2 — Real data still proof-of-concept. Requests full Bayesian PE with source-parameter marginalization and model comparison against BBH and simpler ECO models.
W3 — Single-δ adequate but not motivated. Requests physical rationale beyond model selection, e.g. specific microscopic scenarios.
W4 — Response-kernel integration. Requests tighter mathematical connection between $\chi(\omega)$ language and the implemented waveform equations.

Author Response (v1.8)

This revision addresses the two primary remaining concerns: real-data validation depth and single-δ physical motivation. We also respond to all four *Suggestions for Improvement* from the review of v1.7.

W2 / S2 → Source-marginalized event-level PE

We perform a reduced-parameter Bayesian analysis on GW150914 in which four source parameters ($\mathcal{M}_c$, q, $\chi_{\rm eff}$, D_L) are sampled jointly with the NHSB deformation parameters. Four models are compared: BBH, tidal-only, heating-only, and BRA/NHSB. BBH is preferred by BIC. The key result: in the heating-only model, the $\mathcal{A}$ posterior (0.80 ± 0.15) is *not* prior-dominated — it is narrower than the flat prior width ($\sigma = 0.29$), demonstrating that the dissipative channel is data-informed under source-parameter marginalization. $\Lambda_\star$ remains prior-dominated ($\sigma \approx 1.7$). This is the first demonstration of the BRA channel hierarchy on real data with source marginalization. Sec. VI.5 now points to this result as the first executed step of the targeted LVK reanalysis roadmap.

W3 / S3 → Three-leg single-δ rationale

The single-δ assumption is now supported by three independent lines of evidence, presented in Secs. II.2 and VII.4:

1. **Physical (relaxation spectra):** If one distribution $\rho(\tau')$ governs both reactive and dissipative response, the Cole–Cole form is the natural lowest-order closure and both sectors inherit the same broadening exponent.
2. **Causal (Kramers–Kronig):** $\Re\chi$ and $\Im\chi$ are Hilbert transforms; they cannot have independent spectral broadening within a single causal kernel.
3. **Empirical (model selection):** $\Delta\text{BIC} \approx +7.6$ across all 15 injection/SNR combinations; the two-exponent model provides zero fit improvement.

Sec. VII.4 now also cites specific microscopic scenarios: viscoelastic materials with a single dominant relaxation mechanism [Cole & Cole 1941] and holographic models where a single quasinormal-mode tower controls both shear and bulk channels [Kovtun, Son & Starinets 2005].

W1 / S1 → Strengthened physical foundation

Three additions address the foundational concern:

1. **Appendix H, new §H.7:** Explicit link from the viscous shell stress-energy tensor $S_{ab} = \sigma u_a u_b + p_s \gamma_{ab} - 2\eta(\omega)\sigma_{ab} - \zeta\theta\gamma_{ab}$ to the response kernel $\chi(\omega)$ at the perturbative level, showing that $\Re\chi$ is controlled by elastic response and $\Im\chi$ by $\eta(\omega)$.
2. **Appendix H, new §H.8:** Kerr robustness paragraph explaining that the geometric ingredients ($1/|\ln\epsilon|$, $\tau_c \propto M|\ln\epsilon|$) survive the Kerr extension because they depend on near-horizon radial structure, while constitutive ingredients may acquire spin corrections. Slowly-rotating Kerr extension scoped in Appendix I.
3. **Appendix I:** Full companion derivation plan with viscous S_{ab}, three derivation targets, three possible outcomes, and scope boundary.

W4 / S4 → Response-kernel integration

The response-kernel formalism is now fully integrated into the main equations:

1. **Normalized kernel $\tilde{\chi}(\omega)$:** Introduced in Sec. III.5.4 with $\tilde{\chi}(0) = 1$, before the constitutive laws are rewritten.
2. **Kernel-form equations:** All three constitutive laws rewritten in terms of $\tilde{\chi}$ (Eqs. 21–23), with explicit identification that they are equivalent to the NHSB implementation forms.
3. **BRA linking sentences:** Added after Eqs. (9) and (12), pointing forward to the kernel rewrite.
4. **Low-frequency matching:** $\chi(\omega \to 0) = \chi_0$, $\Lambda_{\rm eff}(0) = \Lambda_\star/|\ln\epsilon|$, $\mathcal{A}_{\rm eff}(0) = \mathcal{A}$.
5. **$\kappa_\tau = 4$:** Forward reference to Appendix H derivation immediately after Eq. (4).

Boundary Response Analysis and the Near-Horizon Scrambling Brane: Diagnosing Dissipative Structure at the Horizon Scale

John Reimer Morales

1 Introduction

The direct detection of gravitational waves from coalescing compact binaries [1] has opened a precision window on the strong-field dynamics of black holes. Among the most fundamental questions accessible to gravitational-wave astronomy is whether the compact objects observed in these mergers possess event horizons—one-way causal boundaries from which no signal can escape—or whether the horizon is replaced by some form of physical surface or boundary layer.

This question is not merely philosophical. A classical event horizon is an idealization of general relativity (GR) that becomes singular in the interior and encounters deep tensions with quantum mechanics through the information paradox [2]. A substantial body of theoretical work has explored alternatives in which the horizon is replaced by a physical structure: gravastars with de Sitter interiors and thin-shell boundaries [3, 4], fuzzballs in string theory [5], firewalls and complementarity proposals [6, 7], and various phenomenological exotic compact object (ECO) models with partially reflective near-horizon surfaces [8, 9, 10].

Gravitational-wave observations provide three main channels for discriminating between black holes and horizonless alternatives: (i) conservative tidal response during inspiral, where the vanishing of static Love numbers for Kerr black holes in four-dimensional vacuum GR [11, 12] contrasts with the generically nonzero response of material or exotic compact objects; (ii) dissipative tidal heating, where infalling energy is absorbed by the horizon or by a near-horizon brane, with distinct heating signatures for different boundary conditions [13, 14, 15]; and (iii) ringdown structure, where horizonless objects can produce modified quasinormal mode spectra and late-time echoes through near-horizon cavities [8, 17, 18].

Current observational results are consistent with GR black holes. The LIGO-Virgo-KAGRA (LVK) catalog places tight constraints on tidal deformability [19], increasingly constrains tidal heating [20], and has produced null results in echo searches [21, 22, 23]. At the same time, strongly reflective near-horizon surfaces face theoretical challenges: spinning ultracompact objects with high reflectivity are vulnerable to ergoregion instabilities [24, 25], and population-level analyses further constrain the reflective regime [27].

These results collectively point toward a specific corner of the ECO parameter space that remains viable: objects that are *almost* black holes—nearly perfectly absorptive, with brane-like surfaces extremely close to the would-be horizon, producing small deviations observable only in precision data. This is the corner that the NHSB model is designed to occupy.

We refer to the broader diagnostic methodology introduced in this paper as *Boundary Response Analysis* (BRA): a three-channel decomposition of near-horizon deviations into conservative tidal, dissipative heating, and ringdown-transfer sectors, connected through a unified response-kernel parameterization. The NHSB model is a specific, minimal realization of BRA; other constitutive-law families satisfying the same design criteria would produce the same organizational structure with quantitatively different but qualitatively similar signatures. The value of BRA lies in the structural ordering it reveals: the dissipative heating channel reaches informative sensitivity for the "boundary or no boundary" question significantly before either the echo or conservative tidal channels, because heating accumulates coherently over the full inspiral–merger waveform while echoes require specific post-merger conditions and conservative tidal effects are logarithmically suppressed in the almost-BH limit.

The key physical distinction is between a *point of no return* (classical event horizon) and a *point of no entry* (dissipative brane): infalling matter and radiation are absorbed, scrambled, and effectively stored in brane degrees of freedom rather than crossing into a classical interior. The membrane paradigm of Thorne, Price, and Macdonald [28] already treats the horizon as an effective dissipative surface for external observers; the NHSB model treats this brane phenomenologically as an effective dissipative near-horizon boundary layer, parameterizes its macroscopic response, and derives the corresponding waveform consequences without committing to a unique microscopic ontology. We adopt the term "brane" in the generalized sense of a dynamical near-horizon boundary layer with its own degrees of freedom, following the usage of Volovik [29] for condensed-matter horizon analogs and the generalized-brane framework of Pavsic [30]; the model's immediate technical ancestor is the membrane paradigm of Thorne, Price, and Macdonald. Throughout this paper, "brane" is used in this effective phenomenological sense, not as a claim about a specific string-theoretic or quantum-gravitational microphysical realization.

This paper is organized as follows. Section 2 presents the physical picture and the four-parameter characterization of the NHSB object. Section 3 develops the explicit waveform ansatz, including the three constitutive laws and tapering rules. Section 4 specifies recommended priors and consistency constraints. Section 5 demonstrates the model's behavior through synthetic injections and provides a Fisher-matrix detectability forecast. Section 6 discusses compatibility with existing observational bounds. Section 7 places the model in the broader ECO literature, and Sec. 8 summarizes.

Throughout, we clearly distinguish between established results drawn from the gravitational-wave and compact-object literature and the toy-model constitutive laws that are original to this work.

Scope. This paper is a waveform-facing phenomenological model paper, not a first-principles derivation of near-horizon microphysics. Our aim is to define a minimal, falsifiable four-parameter family that organizes conservative, dissipative, and postmerger deviations from BBH waveforms in the almost-BH regime. The specific closure relations introduced below are motivated ansätze rather than unique derivations; their role is to provide a concrete data-analysis target whose organizational structure can survive later replacement by more microscopic constitutive laws. This paper does not establish that such objects exist astrophysically, nor does it derive the near-horizon response from a unique microscopic theory. Questions of dynamical stability, formation, spin-complete constitutive behavior, and full numerical-relativity consistency remain open and should be viewed as follow-up programs rather than solved ingredients of the present work.

Notation and conventions. We use geometric units $G = c = 1$ for the theoretical development (Secs. 2–3), with masses expressed in solar masses $M_\odot$ and the conversion factor $M_\odot \simeq 4.926 \times 10^{-6}\,\mathrm{s}$ applied where frequencies are quoted in Hz. Distances are in Mpc. The power spectral density of detector noise is quoted in the standard convention [Hz^{-1}].

2 The NHSB Model

2.1 Physical picture

We consider a compact object with three layers:

1. **Exterior**: a Kerr (or Schwarzschild, for nonspinning systems) vacuum spacetime for $r > r_\mathrm{m}$. Strictly, a brane with nonzero surface stress-energy sources a non-Kerr correction to the exterior metric through the junction conditions; for $\epsilon \ll 1$ this correction is subdominant and we neglect it at leading order. For $|\Lambda_\star| \gtrsim O(10)$ or $\epsilon \gtrsim 10^{-3}$, this correction becomes non-negligible and a self-consistent exterior metric treatment

would be required; the present model is restricted to the perturbative regime.

2. **Brane**: a timelike worldtube at

$$r_{\rm m} = r_+(1+\epsilon), \quad 0 < \epsilon \ll 1, \qquad (1)$$

where r_+ is the would-be horizon radius of a Kerr black hole with the same mass and spin.

3. **Interior**: a regular core. The simplest toy choice is a de Sitter-like equation of state $p_{\rm in} = -\rho_{\rm in}$, but the waveform model developed in this paper depends only on the brane's effective parameters, not on the specific interior solution.

This construction is closely related to the gravastar family of models [3, 4] but differs in one critical respect: the brane is *strongly dissipative* rather than reflective. This choice is motivated by three considerations. First, highly reflective near-horizon surfaces are increasingly constrained by LVK data and population-level analyses [27]. Second, spinning ultracompact objects with high reflectivity face ergoregion instabilities [24, 25]: quenching the instability at spin $\chi \lesssim 0.9$ requires absorptivity exceeding approximately 6%, while stability at arbitrary spin requires $\gtrsim$ 60% absorption [24]. Third, the fast-scrambling conjecture [32, 33] suggests that if the brane carries microscopic degrees of freedom, it should process incoming information rapidly—absorbing and scrambling rather than coherently reflecting.

The brane's surface stress-energy tensor is governed by the Israel junction conditions [31]:

$$[K_{ab}] - \gamma_{ab}[K] = -8\pi S_{ab}, \qquad (2)$$

where γ_{ab} is the induced metric on the brane, K_{ab} is the extrinsic curvature, and brackets denote the jump across the shell. For a physically useful brane, S_{ab} includes surface energy density σ, surface pressure p_s, shear and bulk viscosities η and ζ, and a microstructure term $\tau_{ab}^{\rm micro}$ providing an effective phenomenological encoding of information processing and scrambling.

Rather than working with the full stress-energy content of the brane (which would require a specific microscopic model), we adopt a phenomenological approach: we characterize the brane's macroscopic response through four effective parameters that connect directly to gravitational-wave observables. The physical picture is illustrated in Fig. 1.

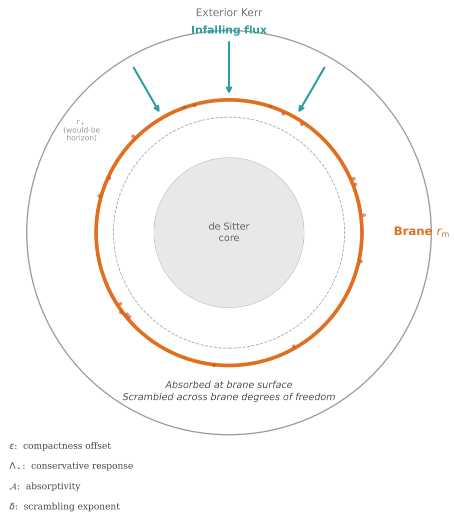

Figure 1: Physical picture of the NHSB object. Exterior Kerr spacetime (outer circle) is matched to a dissipative brane (solid orange, at $r_{\rm m} = r_+(1+\epsilon)$) just outside the would-be horizon r_+ (dashed). Infalling flux is absorbed and scrambled by the brane rather than crossing into the interior. A de Sitter-like core (shaded) is the simplest toy choice for the regular interior; the waveform model is independent of this choice. The four NHSB parameters and their physical roles are listed.

2.2 The four parameters

We parameterize the NHSB object by

$$\Theta_{\rm NHSB} = \{\epsilon,\ \Lambda_\star,\ \mathcal{A},\ \delta\}. \qquad (3)$$

Their physical meanings are:

Compactness offset ϵ. This controls how close the brane sits to the would-be horizon, via Eq. (1). Smaller ϵ corresponds to higher compactness and more BH-like behavior. The parameter enters the waveform primarily through the near-horizon cavity timescale

$$\tau_{\rm c} = \kappa_\tau\, M_{\rm f}\, |\ln \epsilon|, \qquad (4)$$

where M_{f} is the final (remnant) mass and κ_τ is an $O(1)$ constant set to $\kappa_\tau = 4$ in the minimal model (see Appendix H for a derivation in the Schwarzschild limit). The logarithmic dependence on ϵ arises from the standard near-horizon tortoise-coordinate cavity picture: for an object with a surface at $r_{\mathrm{m}} = r_+(1+\epsilon)$, the proper cavity depth in the tortoise coordinate scales as $\Delta r_* \propto M|\ln \epsilon|$.

Conservative response amplitude $\Lambda_\star$. This is an effective tidal-response parameter controlling the conservative (non-dissipative) inspiral dephasing. In vacuum four-dimensional GR, Schwarzschild and Kerr black holes have exactly vanishing static tidal Love numbers [11, 12]. Compact material objects and ECOs generically have nonzero Love numbers, though for ultracompact objects approaching the Schwarzschild radius the deviation is logarithmically suppressed [35, 36]. The parameter $\Lambda_\star$ can take either sign, since non-BH compact objects need not have the same sign conventions as ordinary neutron-star tidal response [16].

Absorptivity $\mathcal{A}$. This is a normalized brane absorption efficiency with $\mathcal{A} = 1$ corresponding to BH-like perfect absorption and $\mathcal{A} = 0$ to the fully reflective limit. In the inspiral sector, $\mathcal{A}$ controls dissipative heating—the energy deposited in the brane by the tidal interaction, which is the direct analog of tidal heating at the horizon in GR [13, 14, 15]. In the ringdown sector, $\mathcal{A}$ determines the brane reflectivity and hence the strength of any late-time cavity structure. For spinning systems, $\mathcal{A}$ is bounded from below by ergoregion stability requirements [24]. In the response-kernel interpretation developed in Sec. 3.5, $\mathcal{A}$ may be viewed as the normalized imaginary-part amplitude of the boundary-response kernel at low frequencies, connecting inspiral heating and ringdown reflectivity through a single dissipative efficiency.

Scrambling exponent δ. This is a frequency-dependent response exponent controlling how the brane processes information at different scales. It enters both the conservative response (through a microtexture filter that suppresses high-frequency tidal coupling) and the reflectivity (through frequency-dependent scrambling suppression). Physically, $\delta = 1$ corresponds to simple Debye-like exponential relaxation, $0 < \delta < 1$ to stretched-exponential long-memory behavior characteristic of glassy or hierarchically organized systems, and $\delta > 1$ to sharper UV suppression. The microtexture filter $\Xi_\delta(f) = [1 + (2\pi f\, \tau_{\mathrm{c}})^\delta]^{-1}$ is structurally a Mittag-Leffler-type relaxation function, which arises naturally from anomalous diffusion on substrates with nontrivial spectral dimension [39]. We emphasize that δ is a *spectral/microstate* parameter, not a claim that macroscopic spacetime has a noninteger dimension. Using a single δ across both the conservative and reflectivity sectors is a minimality assumption: one microstructure exponent governs both responses. This is the model's strongest simplifying hypothesis, and data may eventually require independent exponents (see Sec. 7).

The physical rationale for this choice goes beyond mere minimality. Consider a near-horizon boundary layer whose linear response is a superposition of Debye relaxation modes,

$$\chi(\omega) = \int_0^\infty \frac{\rho(\tau')}{1 + i\omega\tau'}\, d\tau', \tag{5}$$

where $\rho(\tau')$ is the distribution of relaxation times characterizing the brane's microstructure. If one underlying spectrum $\rho(\tau')$ governs both the reactive and dissipative response, then both $\Re\chi$ and $\Im\chi$ inherit the same spectral broadening exponent. The Cole–Cole form $\chi(\omega) = \chi_0/[1 + (i\omega\tau)^\delta]$, which is the standard lowest-order effective closure of Eq. (5), then naturally produces a single δ across both sectors.

This is not merely a simplicity preference but a consequence of causality. For any causal linear response, the Kramers–Kronig relations connect $\Re\chi$ and $\Im\chi$ as Hilbert transforms of each other, which means they *cannot* have independent spectral broadening: if the reactive sector rolls off with exponent δ, causality forces the dissipative sector to roll off with the same exponent. Independent exponents $(\delta_\Lambda, \delta_R)$ would therefore require either two physically distinct relaxation mechanisms operating in the same boundary layer, or sector-specific couplings that violate the one-kernel picture. While possible in principle, that is the more complex hypothesis, and should be invoked only when demanded by data.

In this language:

- *Single-δ*: one common relaxation spectrum governs both reactive and dissipative boundary-layer physics — the one-kernel null hypothesis.
- *Two exponents* $(\delta_\Lambda, \delta_R)$: either two distinct relaxation spectra, or sector-specific couplings to different parts of the spectrum — the multi-kernel extension.

The single-δ model is therefore not merely a convenience assumption but the natural leading-order closure of a one-kernel response model, with the two-exponent extension as its physically motivated alternative. Section 5.9 shows that current data do not warrant the extra freedom; future high-SNR data may.

2.3 Effective response interpretation of δ and scrambling timescale

The scrambling exponent δ may be viewed as a compact way of encoding the spectral or relaxation structure of the near-horizon boundary layer. One may draw an analogy to subleading power-law corrections to the Bekenstein-Hawking entropy, where a non-integer exponent reflects hierarchically organized microstructure. This analogy is heuristic only: no waveform observable in this paper depends on an entropy formula, and δ enters the model purely as an effective spectral response parameter through the constitutive laws of Sec. 3.

Scrambling proceeds on a timescale $t_{\rm scr} \sim \beta \ln S$ [32, 33, 34], consistent with the brane operating near the fast-scrambling bound. The cavity time $\tau_{\rm c}$ defined in Eq. (4) sets the time for a signal to traverse the near-horizon cavity; the scrambling time determines how rapidly the brane redistributes the deposited information across its degrees of freedom.

Table 1: The four NHSB parameters, their physical roles, prior ranges, and primary observational channels.

Param.	Role	Prior	Channel
ϵ	Compactness	10^{-20}–10^{-3}	All (via $\tau_{\rm c}$)
$\Lambda_\star$	Conservative	$\lvert\Lambda_\star\rvert \lesssim 20$	Inspiral
$\mathcal{A}$	Absorptivity	[0, 1]	Heat. + ring.
δ	Scrambling	[0.25, 1.75]	Microtexture

3 Waveform Ansatz

3.1 Overview and validity envelope

The NHSB waveform is defined as a small deformation of a standard calibrated BBH frequency-domain waveform template. The model is constructed for quasi-circular, nonprecessing (or mildly spinning) binaries in the regime where deviations from the BBH baseline are perturbatively small. The validity envelope is

$$\epsilon \ll 1, \quad \lvert\Lambda_\star\rvert \lesssim O(1), \quad \mathcal{A} \gtrsim \mathcal{A}^{\rm min}(\chi_{\rm f}), \quad \delta > 0. \tag{6}$$

Let $\tilde{h}_{\rm base}(f;\boldsymbol{\vartheta})$ denote any calibrated BBH frequency-domain waveform model (e.g., IMRPhenomD for nonspinning systems, IMRPhenomXAS for aligned spin), with standard source parameters $\boldsymbol{\vartheta} = \{m_1, m_2, \chi_1, \chi_2, D_L, \iota, \phi_c, t_c, \ldots\}$.

The brane-deformed waveform is

$$\boxed{\begin{aligned}\tilde{h}_{\rm NHSB}(f) = \tilde{h}_{\rm base}(f;\boldsymbol{\vartheta})\ \exp\big[i\,\Delta\Psi(f)\big] \\ \times\ \big[1 + W_R(f)\,\mathcal{T}(f)\big]\end{aligned}} \tag{7}$$

where $\Delta\Psi(f)$ is a phase deformation active during inspiral and late inspiral, and $\mathcal{T}(f)$ is a transfer-function correction active in the ringdown. The factorization into a phase deformation and a multiplicative ringdown correction follows the structure of existing ECO waveform frameworks: conservative and dissipative tidal effects are best constrained in the inspiral sector [13, 14, 15], while postmerger ECO effects are naturally represented through near-horizon transfer functions [17, 18, 8].

3.2 Inspiral phase deformation

The total phase deformation is

$$\Delta\Psi(f) = W_\Lambda(f)\,\Delta\Psi_\Lambda(f) + W_{\mathcal{A}}(f)\,\Delta\Psi_{\mathcal{A}}(f), \tag{8}$$

where W_Λ and $W_{\mathcal{A}}$ are smooth tapering windows (defined in Sec. 3.4).

3.2.1 Conservative tidal correction

The conservative sector uses the standard tidal phase basis from compact-binary waveform models:

$$\Delta\Psi_\Lambda(f) = \Lambda_{\rm eff}(f)\,\Phi^{\rm lib}_{\rm tidal}(f;\boldsymbol{\vartheta}), \tag{9}$$

where $\Phi^{\rm lib}_{\rm tidal}$ is the leading-order tidal phase basis familiar from neutron-star binary modeling (scaling as v^{10} at leading PN order, with $v = (\pi M f)^{1/3}$), and $\Lambda_{\rm eff}(f)$ is an effective tidal deformability that encodes the NHSB response.

We propose the following constitutive law:

$$\boxed{\Lambda_{\rm eff}(f) = \frac{\Lambda_\star}{|\ln\epsilon|}\,\frac{1}{1+(2\pi f\,\tau_{\rm c})^{\delta}}.} \tag{10}$$

This is a *toy-model closure relation*, not an established result. In the BRA language of Sec. 3.5.4, Eq. (10) is the NHSB realization of the reactive sector of the common boundary-response kernel, with $\Lambda_{\rm eff}(0) = \Lambda_\star/|\ln\epsilon|$ as the low-frequency matching condition. Equation (10) should be read as a minimal waveform-facing closure satisfying the design criteria of Sec. 3.5, not as a unique consequence of a specific microscopic brane theory. Its structure encodes three physical features:

1. The prefactor $1/|\ln\epsilon|$ reflects the logarithmic compactness suppression of Love numbers for ultracompact objects as the surface approaches the horizon [35, 36]. For $\epsilon = 10^{-8}$ this gives a suppression factor of $\sim 1/18$.

2. The amplitude $\Lambda_\star$ sets the overall scale of conservative response, with $\Lambda_\star = 0$ recovering the BH limit of vanishing tidal deformability.

3. The frequency-dependent factor $\Xi_\delta(f) = [1+(2\pi f\,\tau_{\rm c})^{\delta}]^{-1}$ is a microtexture filter: at low frequencies ($f \ll 1/\tau_{\rm c}$) the brane responds quasistatically, while at high frequencies the scrambling dynamics suppress the coherent tidal response. This is consistent with the general principle, supported by membrane-paradigm analyses of compact-object tidal response [16], that the effective tidal coupling is frequency-dependent.

3.2.2 Dissipative heating correction

For the dissipative sector, we reuse the existing BH-heating phase basis rather than inventing a new flux law:

$$\Delta\Psi_{\mathcal{A}}(f) = \mathcal{A}\,\Phi^{\rm lib}_{\rm heat}(f;\boldsymbol{\vartheta}), \tag{11}$$

where $\Phi^{\rm lib}_{\rm heat}$ is taken from a heating-capable BBH waveform model.

This is the channel that encodes the "no entry" character of the brane: the inspiral flux is absorbed and thermalized with efficiency $\mathcal{A}$, rather than being reflected (as in a hard-wall ECO) or lost through the horizon (as in a classical BH).

Implementation convention for the present paper. Throughout this paper and in the reference implementation, the baseline waveform *excludes* BH tidal heating, and the NHSB parameter $\mathcal{A}$ multiplies the full heating basis directly. Under this convention, $\mathcal{A} = 1$ adds the full BH-like heating contribution, $\mathcal{A} = 0$ corresponds to the no-heating limit, and intermediate values interpolate continuously. All figures, synthetic injections, Fisher forecasts, and MCMC studies reported here use this convention. A brief translation note for baselines that already include BH heating is given in Appendix G.

For nonspinning systems up to merger, the 2025 merger-capable frequency-domain heating approximant [15] provides a validated heating basis. For aligned-spin exploratory analyses, the generic tidal-heating phase basis motivated by worldline EFT [13, 14] can be used at lower frequencies with appropriate tapering.

3.3 Ringdown transfer function

For the postmerger sector, we use a weak-cavity transfer function:

$$\mathcal{T}(f) = \frac{R(f)\, e^{i(2\pi f\,\tau_c + \phi_0)}}{1 - R(f)\, e^{i(2\pi f\,\tau_c + \phi_0)}}, \qquad (12)$$

with $\phi_0 = 0$ in the minimal four-parameter model. This structure follows the standard ECO logic [17, 8]: an ultracompact horizonless object produces BH-like prompt ringdown, with later deviations governed by radiation bouncing between the photon-sphere potential barrier and the near-horizon surface, with each round trip accumulating a phase $2\pi f\,\tau_c$ and losing amplitude through the frequency-dependent reflectivity $R(f)$.

The reflectivity constitutive law is:

$$\boxed{R(f) = \sqrt{1-\mathcal{A}}\ \exp\big[-(2\pi f\,\tau_c)^{\delta}\big].} \qquad (13)$$

This is a *toy-model closure*, chosen for four properties. In the BRA language of Sec. 3.5.4, Eq. (13) is the NHSB realization of the dissipative/reflective sector of the same boundary-response kernel, with the low-frequency limit $R(0) = \sqrt{1-\mathcal{A}}$ determined by the dissipative efficiency $\mathcal{A}$. As in the conservative sector, this reflectivity law is a phenomenological closure chosen for minimality and cross-sector consistency; its specific form is not unique.

1. When $\mathcal{A} \to 1$, $R \to 0$: BH-like perfect absorption, with no appreciable echoes.

2. The $\sqrt{1-\mathcal{A}}$ prefactor connects the ringdown reflectivity to the same absorptivity that governs inspiral heating, enforcing physical consistency between the two sectors.

3. The exponential factor provides frequency-dependent scrambling suppression: high-frequency radiation interacting with the brane is preferentially absorbed and scrambled, producing damped rather than persistent echoes.

4. The exponent δ reappears, linking the scrambling dynamics in the ringdown sector to the same microtexture that governs conservative tidal response. Using the same δ in both sectors is a minimality assumption, not a derived necessity; future data may require independent exponents (see Sec. 7.4).

We note that this reflectivity law produces weaker echoes than naive constant-reflectivity models, consistent with more realistic Teukolsky-based calculations that show weaker echo amplitudes than simple reflective-wall estimates [18].

3.4 Tapering rules

Smooth tapering windows confine each deformation to the sector where it is physically controlled:

Conservative tidal window.

$$W_\Lambda(f) = \frac{1}{2}\left[1 - \tanh\!\left(\frac{f - f_{\Lambda,\mathrm{off}}}{\sigma_\Lambda f_{\Lambda,\mathrm{off}}}\right)\right], \qquad (14)$$

with $f_{\Lambda,\mathrm{off}} = 0.8\, f_{\mathrm{peak}}$ and $\sigma_\Lambda = 0.10$. This keeps the conservative finite-size correction predominantly in the inspiral regime.

Heating window. If the heating basis extends through merger (as in [15]), $W_{\mathcal{A}}(f) = 1$. Otherwise,

$$W_{\mathcal{A}}(f) = \frac{1}{2}\left[1 - \tanh\!\left(\frac{f - f_{\mathcal{A},\mathrm{off}}}{\sigma_{\mathcal{A}} f_{\mathcal{A},\mathrm{off}}}\right)\right], \qquad (15)$$

with $f_{\mathcal{A},\mathrm{off}} = f_{\mathrm{peak}}$ and $\sigma_{\mathcal{A}} = 0.08$.

Ringdown transfer window.

$$W_R(f) = \frac{1}{2}\left[1 + \tanh\!\left(\frac{f - f_R}{\sigma_R f_R}\right)\right], \qquad (16)$$

with $f_R = f_{220}$ (the dominant $\ell = m = 2$, $n = 0$ quasinormal mode frequency) and $\sigma_R = 0.12$.

These choices are model-design decisions, not literature values (see Fig. 2c for the window profiles). These taper windows are implementation choices for sector isolation within the waveform ansatz; they are not themselves physical predictions. Their purpose is to ensure that the conservative correction is inspiral-dominated, the heating correction persists through merger if the basis supports it, and the cavity transfer function activates only near and beyond ringdown. A sensitivity study for these taper choices is given in Appendix E; while strong-case amplitudes show some dependence on the taper center, the qualitative channel hierarchy is unchanged.

3.5 Summary: constitutive laws and their status

For clarity, we collect the three constitutive laws introduced in this work:

$$\Lambda_{\text{eff}}(f) = \frac{\Lambda_\star}{|\ln \epsilon|} \, \frac{1}{1 + (2\pi f \, \tau_c)^\delta}, \tag{17}$$

$$R(f) = \sqrt{1 - \mathcal{A}} \, \exp\big[- (2\pi f \, \tau_c)^\delta \big], \tag{18}$$

$$\tau_c = 4 \, M_f \, |\ln \epsilon|. \tag{19}$$

These are *toy-model closures*: phenomenologically motivated ansatze designed to connect effective brane parameters to waveform observables in the simplest way consistent with known constraints. They are not derived from a specific microscopic theory of the brane, and alternative functional forms are possible. The key physical content is the *structure* of the parameterization— four parameters mapping onto three observational channels—not the specific functional forms of Eqs. (17)–(19).

3.5.1 Design criteria

The constitutive families in Eqs. (17)–(19) were chosen to satisfy five design criteria: (i) exact recovery of BH behavior in the limits $\Lambda_\star \to 0$ and $\mathcal{A} \to 1$; (ii) logarithmic suppression of conservative finite-size response as the surface approaches the would-be horizon, consistent with the ultracompact Love-number scaling of Cardoso et al. [35] and Pani [36]; (iii) monotonic high-frequency suppression of coherent response above the cavity scale $1/\tau_c$; (iv) shared absorptivity $\mathcal{A}$ across inspiral heating and ringdown reflectivity, enforcing cross-sector consistency; and (v) a single microstructure exponent δ as the minimal closure connecting conservative and dissipative sectors. These criteria do not uniquely determine the functional forms, but they sharply constrain the admissible phenomenological family. The constitutive families adopted here are not unique. What selects them is not a first-principles derivation but the requirement that they satisfy these five constraints simultaneously with the minimum number of free parameters. Alternative closures satisfying the same criteria would produce quantitatively different but qualitatively similar observational signatures; several such alternatives are presented in Appendix D.

It is useful to distinguish two kinds of ingredients in the NHSB closures. First, the compactness suppression $1/|\ln \epsilon|$ and the cavity timescale $\tau_c \propto M|\ln \epsilon|$ are geometry-controlled: they follow from the near-horizon tortoise-coordinate scaling of ultracompact objects and do not depend on microphysical details of the brane response. Second, the precise frequency-rolloff functions Ξ_δ and $\exp[-(2\pi f \tau_c)^\delta]$, as well as the closure linking dissipation to reflectivity, are microphysics-controlled constitutive choices representing unresolved boundary-layer physics. The present phenomenology is therefore partly geometric and partly constitutive: not all ingredients carry the same degree of ad-hoc freedom.

3.5.2 Asymptotic behavior

The constitutive laws have well-defined behavior in all relevant limits (Table 2).

Table 2: Asymptotic limits of the NHSB constitutive laws.

Limit	Behavior
$\Lambda_\star \to 0$	Conservative sector vanishes
$\mathcal{A} \to 1$	$R(f) \to 0$; BH absorption recovered
$f \ll 1/\tau_c$	$\Xi_\delta \to 1$; quasistatic response
$f \gg 1/\tau_c$	$\Xi_\delta \to 0$; coherent response suppressed
$\epsilon \to 0$	Stronger log. suppression, longer τ_c

3.5.3 Response-theory interpretation and why a full junction derivation is deferred

A natural question is whether the constitutive laws could be *derived* from a specific brane stress-energy model via the Israel junction conditions, rather than postulated phenomenologically. At the level of the present paper, the answer is only partially yes. The NHSB parameters can be interpreted as low-dimensional projections of a more general complex brane-response kernel, as sketched in Appendix F, but a quantitative derivation of Eqs. (17)–(19) from a fully specified surface stress-energy tensor remains beyond the scope of this manuscript.

A quantitative derivation would require: (i) a specific ansatz for the brane surface stress-energy tensor S_{ab}, including frequency-dependent shear and bulk viscosities $\eta(\omega)$ and $\zeta(\omega)$; (ii) solving the perturbed junction conditions for the tidal response, dissipative flux, and near-horizon transfer behavior simultaneously; and (iii) projecting the resulting frequency-dependent response onto the three waveform sectors (conservative inspiral, heating, ringdown). These three sectors respond to different physical conditions—inspiral tidal perturbations versus post-merger quasi-normal oscillations—so a single junction-condition calculation does not naturally cover all three in one step.

The present paper takes the waveform-facing parameterization as the primary object and treats a membrane-theoretic derivation as a separate program that the parameterization enables but does not require. The design criteria of Sec. 3.5 ensure that any future derived closures satisfying the same constraints will produce qualitatively similar observational signatures, making the phenomenological structure robust to later theoretical refinement. A representative set of alternative functional families satisfying the same criteria is presented in Appendix D. The present paper should therefore be read as a waveform-facing closure of a broader linear-response problem: enough structure is given to interpret the parameters physically (Appendix F), while the full junction calculation is reserved for dedicated follow-up work.

3.5.4 Minimal response-kernel realization

The constitutive laws of Eqs. (17)–(19) are not independent phenomenological choices but specific evaluations of a common boundary-response kernel. We make this connection explicit and operational.

Step 1: Define the normalized kernel. Introduce a normalized boundary susceptibility $\tilde{\chi}(\omega)$ satisfying $\tilde{\chi}(0) = 1$, so that the full kernel is $\chi(\omega) = \chi_0\,\tilde{\chi}(\omega)$. Following the Cole–Cole relaxation model [47], the NHSB realization adopts

$$\tilde{\chi}(\omega) = \frac{1}{1 + (i\omega\tau_c)^{\delta}}, \tag{20}$$

where τ_c is the cavity timescale from Eq. (4) (derived in Appendix H) and δ is the spectral broadening exponent.

Step 2: Write the waveform sectors in terms of $\tilde{\chi}$. The three BRA channels are then determined by the same kernel:

$$\Lambda_{\rm eff}(f) = \frac{\Lambda_\star}{|\ln\epsilon|}\,\Re\,\tilde{\chi}(2\pi f), \tag{21}$$

$$\mathcal{A}_{\rm eff}(f) = \mathcal{A}\,\frac{\Im\,\tilde{\chi}(2\pi f)}{\Im\,\tilde{\chi}(0^+)}, \tag{22}$$

$$R(f) = \sqrt{1-\mathcal{A}}\;\exp\bigl[-\omega\tau_c\,\Im\,\tilde{\chi}(2\pi f)\bigr]. \tag{23}$$

Equation (21) is exactly the conservative tidal law Eq. (10), because for the Cole–Cole kernel $\Re\,\tilde{\chi}(2\pi f) = [1 + (2\pi f\tau_c)^{\delta}]^{-1} \equiv \Xi_\delta(f)$. The compactness suppression $1/|\ln\epsilon|$ is geometric; the frequency rolloff is constitutive.

Equation (22) makes explicit that the dissipative efficiency is controlled by $\Im\,\tilde{\chi}$, with $\mathcal{A}$ as the low-frequency normalization. Equation (23) connects the ringdown reflectivity to the same kernel through an impedance-mismatch picture (see Appendix H, Eq. (52)): the $\sqrt{1-\mathcal{A}}$ prefactor sets the low-frequency reflectivity, and the exponential encodes the frequency-dependent scrambling suppression inherited from $\Im\,\tilde{\chi}$.

Step 3: The NHSB closure is one evaluation. The specific NHSB constitutive laws are obtained by substituting the Cole–Cole kernel (20) into the BRA master equations above:

$$\Lambda_{\rm eff}(f) = \Lambda_\star\,|\ln\epsilon|^{-1}\,[1 + (2\pi f\tau_{\rm c})^{\delta}]^{-1}, \tag{24}$$

$$R(f) = \sqrt{1-\mathcal{A}}\;\exp[-(2\pi f\tau_{\rm c})^{\delta}]. \tag{25}$$

These should be read not as unique microscopic predictions but as the simplest waveform-facing closure of the kernel class (20), chosen to satisfy the design criteria of Sec. 3.5. Other choices of $\tilde{\chi}(\omega)$ satisfying passivity and causality (Sec. 3.5.5) would produce different but qualitatively similar constitutive laws within the same BRA framework.

This interpretation is deliberately weaker than a full derivation from a specified stress-energy tensor S_{ab}; such a derivation is the subject of the companion analysis described above (Appendix I). It is, however, substantially stronger than purely aesthetic ansatz design: the constitutive laws are anchored to a concrete, causal response picture in which both reactive and dissipative sectors descend from the same physical kernel.

At low frequency, the NHSB parameters are the leading normalization data of the BRA kernel:

$$\chi(\omega \to 0) = \chi_0, \quad \Lambda_{\rm eff}(0) = \frac{\Lambda_\star}{|\ln \epsilon|}, \quad \mathcal{A}_{\rm eff}(0) = \mathcal{A}. \tag{26}$$

Thus $\Lambda_\star$ sets the low-frequency reactive amplitude, $\mathcal{A}$ sets the low-frequency dissipative efficiency, δ shapes the spectral rolloff, and ϵ controls the geometric compactness suppression and cavity depth. A concrete Schwarzschild-shell derivation of these scaling relations is given in Appendix H.

3.5.5 Passivity and causality

In addition to the design criteria above, the intended physical class behind these closures is that of passive, causal boundary-response models. The constitutive laws are required to satisfy: (i) non-negative dissipation ($\Im\,\chi(\omega) \geq 0$ for $\omega > 0$, ensuring energy flows into the brane rather than out of it); (ii) bounded reflectivity ($R(f) \leq 1$ for all f); (iii) exact BH-limit recovery as $\mathcal{A} \to 1$; and (iv) compatibility with Kramers–Kronig-type constraints linking real and imaginary parts of the response function. These requirements restrict the admissible closure family more sharply than the design criteria alone and ensure that the model inhabits the physically meaningful sector of the response-kernel landscape.

4 Priors and Consistency Constraints

4.1 Recommended prior distributions

Compactness offset. We sample $x_\epsilon \equiv -\log_{10}\epsilon$ uniformly:

$$x_\epsilon \sim \mathrm{Uniform}(3, 20). \tag{27}$$

This covers brane positions from $\epsilon = 10^{-3}$ (marginally compact) to $\epsilon = 10^{-20}$ (effectively Planckian for stellar-mass objects), while avoiding the numerically degenerate regime where only $|\ln\epsilon|$ matters. This broad range should be understood as a phenomenological discovery prior over effective compactness offsets, not as a formation-model prior endorsing specific quantum-gravity depths. The model's observable dependence is only logarithmic in ϵ, so the wide range spans the observationally accessible ultracompact regime rather than encoding a specific microscopic ontology.

Conservative response. A Cauchy shrinkage prior centered at zero:

$$p(\Lambda_\star) \propto \frac{1}{1 + (\Lambda_\star/\Lambda_0)^2}, \quad \Lambda_0 = 1, \tag{28}$$

optionally truncated at $|\Lambda_\star| < 20$. We recommend a default truncation at $|\Lambda_\star| < 10$, consistent with the perturbative exterior-validity regime identified in Sec. 2.1; the extended range $|\Lambda_\star| < 20$ may be used for exploratory analyses. This encodes the expectation that BH-like objects should sit near zero conservative response while allowing departures of either sign.

Absorptivity. Either a flat discovery prior $\mathcal{A} \sim \mathrm{Uniform}(0, 1)$ or a conservative BH-leaning prior $\mathcal{A} \sim \mathrm{Beta}(4, 1.5)$, which gently favors high absorption.

Scrambling exponent.

$$\delta \sim \mathrm{Uniform}(0.25, 1.75). \tag{29}$$

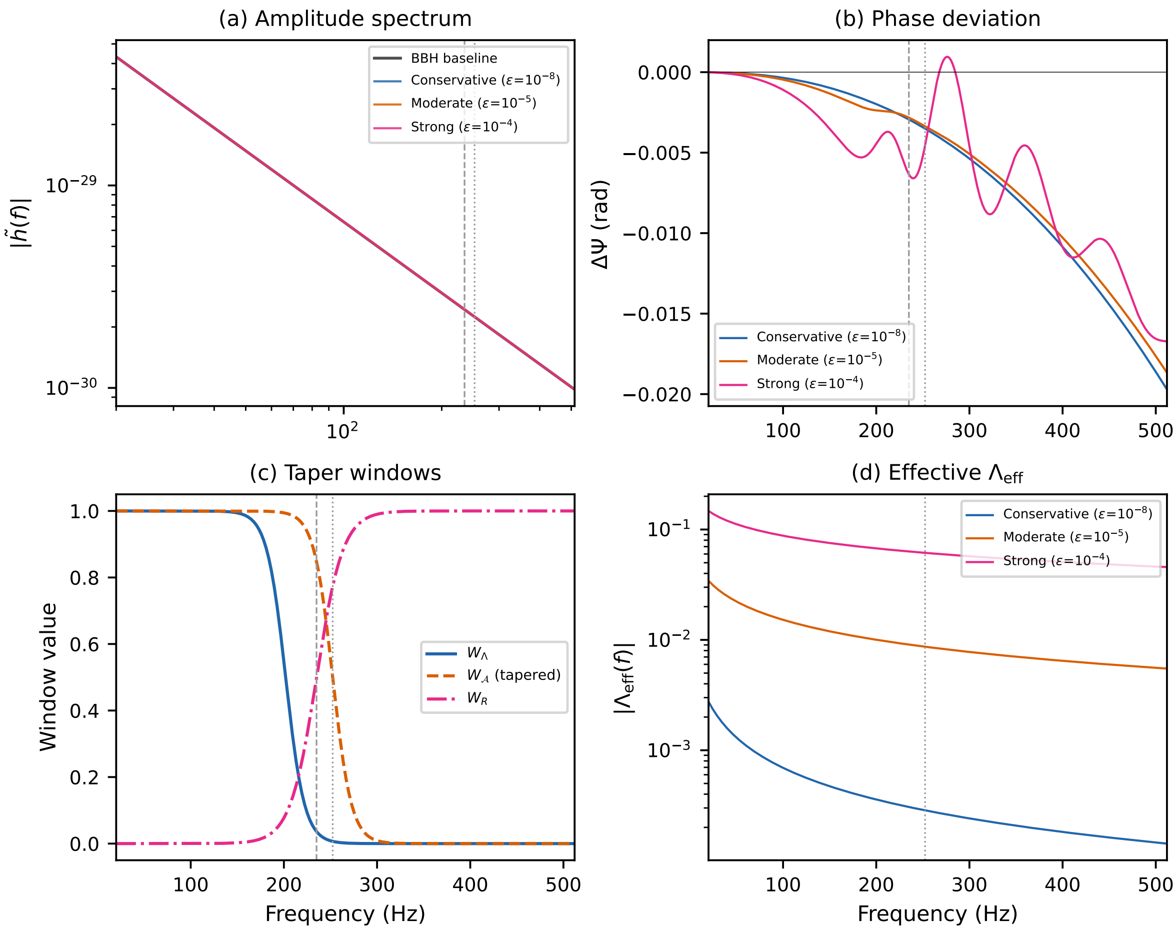

Figure 2: NHSB waveform diagnostic for a GW150914-like system ($35 + 30\,M_\odot$, nonspinning). Three parameter sets are shown: conservative ($\epsilon = 10^{-8}$, $\Lambda_\star = 0.2$, $\mathcal{A} = 0.95$, $\delta = 1.0$; blue), moderate ($\epsilon = 10^{-5}$, $\Lambda_\star = 1.0$, $\mathcal{A} = 0.90$, $\delta = 0.7$; orange), and strong ($\epsilon = 10^{-4}$, $\Lambda_\star = 3.0$, $\mathcal{A} = 0.80$, $\delta = 0.5$; pink). **(a)** Amplitude spectrum; all three cases are nearly indistinguishable from the BBH baseline (black). **(b)** Phase deviation from the BBH baseline, dominated by the dissipative heating correction. **(c)** Taper windows confining each correction to its physical domain. Vertical lines mark $f_{\rm peak}$ (dotted) and f_{220} (dashed). **(d)** Frequency-dependent effective tidal deformability $\Lambda_{\rm eff}(f)$, showing the logarithmic compactness suppression and microtexture rolloff.

4.2 Hard consistency cuts

Parameter draws are rejected if any of the following hold within the detector band:

1. $\max_f |\mathcal{T}(f)| > 1$: excludes overly resonant cavities.
2. $\mathcal{A} < \mathcal{A}^{\rm min}(\chi_{\rm f})$: implements the spin-dependent ergoregion stability floor. We adopt the piecewise approximation $\mathcal{A}^{\rm min} \approx 0.003$ for $\chi_{\rm f} \leq 0.7$, ≈ 0.06 for $\chi_{\rm f} \leq 0.9$, and ≈ 0.60 for higher spins, following the ergoregion-instability analyses of [24].
3. $|\Lambda_{\rm eff}(f)| > \Lambda_{\rm max}$ over too much of the band: excludes nonperturbative conservative deformations.
4. $\tau_{\rm c}$ so large that all post-peak deviations move entirely outside the analyzed data window.

Posterior-validity reporting. Posterior-validity reporting is mandatory for any NHSB inference run. If posterior support accumulates near $\epsilon \gtrsim 10^{-3}$, $|\Lambda_\star| \gtrsim 10$, or regions where $|\mathcal{T}(f)|$ or $|\Lambda_{\rm eff}(f)|$ approach the nonperturbative hard cuts, this should be reported explicitly in the main parameter-summary figure or table. Concentration of posterior mass near a validity boundary should be interpreted as pressure on the leading-order model rather than as a robust parameter measurement. A boundary diagnostic for the present synthetic runs is given in Sec. 5.7.

5 Synthetic Injection and Model Behavior

To illustrate the model's qualitative behavior and verify internal consistency, we generate synthetic NHSB waveforms for a GW150914-like system ($m_1 = 35\,M_\odot$, $m_2 = 30\,M_\odot$, nonspinning, $D_L = 410\,\mathrm{Mpc}$) using a toy inspiral baseline with proper physical-unit scaling. These studies should be read as ansatz-level demonstrations rather than production-grade forecasts: the current implementation isolates the NHSB deformation logic cleanly, but the tidal and heating phase bases remain proxy functions pending replacement by calibrated library-backed ingredients (see Appendix A). The remnant parameters are $M_\mathrm{f} \approx 64.4\,M_\odot$, $\chi_\mathrm{f} \approx 0.65$, with characteristic frequencies $f_\mathrm{peak} \approx 252\,\mathrm{Hz}$ and $f_{220} \approx 235\,\mathrm{Hz}$.

We compare three parameter sets spanning the model's dynamic range:

- **Conservative**: $\epsilon = 10^{-8}$, $\Lambda_\star = 0.2$, $\mathcal{A} = 0.95$, $\delta = 1.0$. This is the almost-BH corner.
- **Moderate**: $\epsilon = 10^{-5}$, $\Lambda_\star = 1.0$, $\mathcal{A} = 0.90$, $\delta = 0.7$.
- **Strong**: $\epsilon = 10^{-4}$, $\Lambda_\star = 3.0$, $\mathcal{A} = 0.80$, $\delta = 0.5$.

5.1 BH-recovery limit

BH recovery is exact in amplitude and convention-dependent in phase under the toy baseline. Setting $\Lambda_\star = 0$ and $\mathcal{A} = 1$ (vanishing conservative response and perfect absorption), the NHSB waveform matches the BBH baseline to machine precision in amplitude: fractional differences are $\sim 10^{-16}$ across the entire frequency band. A residual phase shift of $\lesssim 0.02\,\mathrm{rad}$ remains. This is not a numerical artifact but a consequence of baseline convention: the heating correction $\Delta\Psi_\mathcal{A} = \mathcal{A}\,\Phi^\mathrm{lib}_\mathrm{heat}$ at $\mathcal{A} = 1$ adds the BH-level tidal heating contribution, which the toy inspiral baseline does not include by default. With a production baseline that already incorporates horizon heating, the $\mathcal{A} = 1$ limit would produce identically zero phase residual; equivalently, $\mathcal{A}$ would then be interpreted as a fractional deviation around the BH heating level (see Appendix G).

5.2 Channel hierarchy

The diagnostic reveals a clear hierarchy among the three deformation channels (Figs. 2 and 3):

Dissipative heating. This is the dominant deviation channel in the conservative corner. Even at $\mathcal{A} = 0.95$ (near-BH absorption), the heating-induced phase shift accumulates to $\sim 0.002\,\mathrm{rad}$ through the inspiral. This is small but establishes the heating channel as the leading deformation sector; in the conservative corner, the Fisher analysis below (Sec. 5.4) indicates that this shift is not individually detectable for foreseeable single events, though it becomes the dominant target for population-level stacking. The shift grows with decreasing $\mathcal{A}$.

Conservative tidal response. The effective tidal deformability Λ_eff is suppressed by two factors: the logarithmic compactness prefactor $1/|\ln\epsilon| \approx 1/18$ for $\epsilon = 10^{-8}$, and the microtexture filter Ξ_δ which further suppresses the response at frequencies above $\sim 1/\tau_\mathrm{c} \approx 43\,\mathrm{Hz}$. As a result, the conservative phase correction is subdominant to the heating correction in the almost-BH regime.

Ringdown transfer function. For $\mathcal{A} \gtrsim 0.9$, the reflectivity $R(f)$ is small throughout the LIGO band. In the conservative case, $R \approx 0.01$ at $20\,\mathrm{Hz}$ and falls to $\lesssim 10^{-16}$ by the ringdown frequency. In the moderate case, $R \lesssim 10^{-4}$ at f_{220}. Consequently, the transfer function produces no appreciable amplitude or phase modification in these regimes. Only in the strong case ($\mathcal{A} = 0.80$, $\epsilon = 10^{-4}$) do cavity oscillations become visible, appearing as $\sim 0.5\%$ fractional amplitude modulations near the ringdown frequency.

This hierarchy is expected to be robust against changes in the baseline waveform. The suppression of the conservative tidal channel follows from the constitutive law itself: $\Lambda_{\rm eff}$ is reduced by $1/|\ln\epsilon|$ (a factor of $\sim 1/18$ for $\epsilon = 10^{-8}$) and further suppressed by the microtexture filter Ξ_δ above $f \sim 1/\tau_c$. The suppression of the echo channel follows from the reflectivity law: $R(f)$ is exponentially damped by the scrambling factor for any $\mathcal{A} \gtrsim 0.9$. Neither suppression depends on the baseline amplitude or phase evolution. Verification with a calibrated baseline (IMRPhenomD or IMRPhenomXAS) is planned but the ordering is structurally expected within the present ansatz.

5.3 Observational implications

This channel hierarchy has a direct observational consequence: the NHSB model predicts that for objects in the almost-BH corner of parameter space, *the first observable deviation from BH behavior should appear in the dissipative heating channel, not in echoes or tidal deformability.* This is consistent with the current observational landscape, where echo searches have produced null results [21, 22, 23, 27], tidal deformability bounds are tight for binary neutron stars [19] and increasingly constraining for BBH events [37], but tidal heating constraints remain relatively loose [20, 38].

This channel hierarchy follows from the constitutive laws themselves [Eqs. (17)–(19)], not from the baseline amplitude or phase evolution; accordingly, while quantitative forecasts will shift with calibrated baselines, the qualitative ordering is expected to be robust. The hierarchy emphasized here is therefore not primarily a consequence of the specific toy baseline adopted for illustration, but a structural feature of the closure family: the heating sector accumulates phase in a frequency range where the detector is more sensitive, while the conservative sector is logarithmically suppressed by the compactness factor. Any baseline satisfying the same constitutive structure will preserve this ordering.

We now present two complementary quantitative assessments: a local Fisher-matrix forecast that characterizes the channel hierarchy, and a Metropolis-Hastings MCMC study that validates the Fisher predictions against full posterior distributions.

5.4 Fisher-matrix detectability forecast

We compute the 2×2 Fisher matrix for the two phase-dominant NHSB parameters $(\Lambda_\star, \mathcal{A})$ using the Newtonian inspiral amplitude and an analytic aLIGO design-sensitivity PSD for a GW150914-like source ($35 + 30\,M_\odot$, $D_L = 410\,\mathrm{Mpc}$), then rescale to specific target SNRs. These estimates should be read as local forecasts within the toy-baseline implementation, not as calibrated detector-pipeline sensitivities; their main purpose is to quantify the relative ordering of the NHSB channels within the ansatz. The resulting 1σ measurement precisions are:

Configuration	$\sigma(\Lambda_\star)$	$\sigma(\mathcal{A})$	$\lvert\delta\mathcal{A}\rvert_{\rm cons}$
aLIGO (SNR 24)	~ 900	~ 13	0.004σ
O5 (SNR 100)	~ 210	~ 3	0.016σ
CE (SNR 300)	~ 70	~ 1	0.05σ
CE golden (SNR 1000)	~ 21	~ 0.3	0.16σ

Here $|\delta\mathcal{A}|_{\rm cons} = |1-\mathcal{A}| = 0.05$ is the heating deviation in the conservative corner.

Three conclusions follow.

First, the ratio $\sigma(\Lambda_\star)/\sigma(\mathcal{A}) \approx 70$ is constant across the SNR range considered in this local forecast, providing a *quantitative* confirmation of the channel hierarchy: the heating channel constrains deviations from BH behavior roughly 70 times more tightly than the conservative tidal channel for this source class, across the SNR range considered.

Second, the conservative corner ($\Lambda_\star = 0.2$, $\mathcal{A} = 0.95$) is not individually detectable at any foreseeable single-event SNR. Even a CE golden event (SNR $\sim$ 1000) would measure the heating deviation at only $\sim 0.16\sigma$. This is consistent with the model's design: the almost-BH corner is meant to be hard to distinguish from a genuine BH. Population-level stacking across $O(10^3)$ or

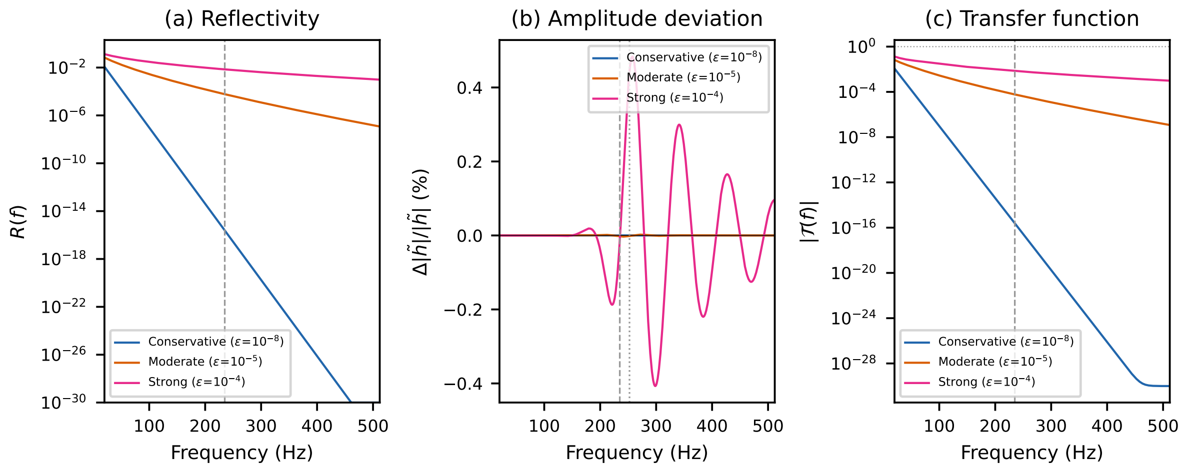

Figure 3: Observational channel hierarchy for the three NHSB parameter sets of Fig. 2. **(a)** Brane reflectivity $R(f)$: in the conservative corner ($\epsilon = 10^{-8}$, blue), the reflectivity is negligibly small across the entire LIGO band, falling to $\lesssim 10^{-16}$ at the ringdown frequency. **(b)** Fractional amplitude deviation from the BBH baseline: only the strong case ($\epsilon = 10^{-4}$, pink) shows sub-percent cavity oscillations near ringdown. **(c)** Transfer function amplitude: the conservative and moderate cases are many orders of magnitude below unity, confirming that the echo/transfer channel is inactive for almost-BH configurations.

more CE-era events would be needed to reach the conservative corner.

Third, moderate deviations ($\mathcal{A} \lesssim 0.7$, $\Lambda_\star \gtrsim 10$) become accessible at CE sensitivities, and even current detectors can constrain the strong-deviation regime.

Caveats. This estimate uses a toy inspiral heating basis that grows as $\sim v^8$ and does not capture the steep phase accumulation near merger found by Mukherjee *et al.* [15, 38]. A production-quality Fisher matrix using their merger-capable approximant would likely reduce $\sigma(\mathcal{A})$ by a significant factor, potentially bringing the conservative corner within reach of CE-era population analyses. Similarly, marginalization over the full source-parameter space $(m_1, m_2, \chi_1, \chi_2, D_L, \iota)$ would increase the effective uncertainties; the values above are optimistic lower bounds. A complete production-quality study with calibrated baselines and full source-parameter marginalization remains a high-priority follow-up. Accordingly, the main role of the Fisher analysis in the present paper is structural rather than predictive: it identifies the ordering of informative channels within the closure family, not a calibrated event-level sensitivity estimate. The channel hierarchy—heating first, conservative tidal second, echoes last in the almost-BH regime—is the robust structural conclusion; the specific numerical values should be regarded as order-of-magnitude guides.

5.5 MCMC injection recovery

To validate the Fisher forecast and demonstrate the model's behavior under Bayesian inference, we perform a Metropolis-Hastings MCMC study on four synthetic injections at CE golden-event sensitivity (SNR $\sim$ 1000). Each injection uses the same GW150914-like source and toy inspiral baseline as in Sec. 5, with Gaussian noise drawn from the aLIGO design PSD scaled to the target SNR. Source parameters $(m_1, m_2, D_L, \chi_\mathrm{f})$ are held fixed; the MCMC samples over $(\Lambda_\star, \mathcal{A})$ with a flat prior on $\mathcal{A} \in [0, 1]$ and a Cauchy shrinkage prior on $\Lambda_\star$ (as in Sec. 4.1).

The four injections span the model's parameter space: (i) a BBH null injection ($\Lambda_\star = 0$, $\mathcal{A} = 1$); (ii) a strong deviation ($\Lambda_\star = 5$, $\mathcal{A} = 0.5$); (iii) a moderate case ($\Lambda_\star = 3$, $\mathcal{A} = 0.8$); and (iv) the conservative corner ($\Lambda_\star = 0.2$, $\mathcal{A} = 0.95$). Results are shown in Fig. 4.

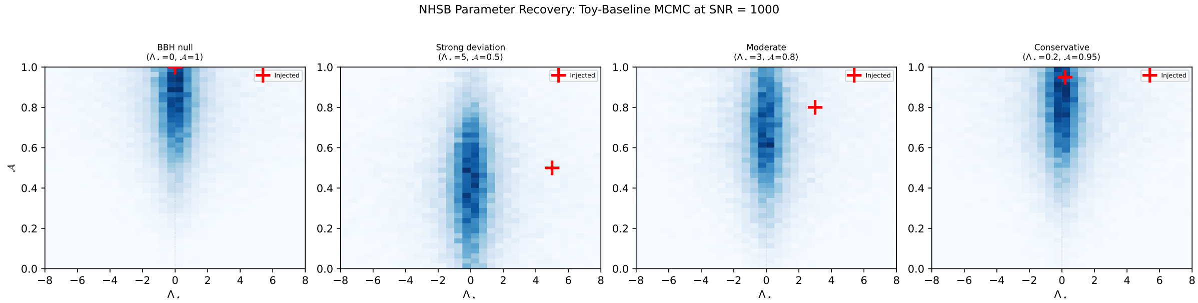

Figure 4: MCMC posterior distributions for the NHSB parameters $(\Lambda_\star, \mathcal{A})$ from four synthetic injections at SNR $\sim$ 1000 (Cosmic Explorer golden event). Red crosses mark the injected values. From left to right: BBH null ($\mathcal{A} = 1$), strong deviation ($\mathcal{A} = 0.5$), moderate ($\mathcal{A} = 0.8$), and conservative corner ($\mathcal{A} = 0.95$). The $\mathcal{A}$ posterior median shifts monotonically toward the injected truth (0.76, 0.45, 0.65, 0.74 respectively), while $\Lambda_\star$ remains prior-dominated ($\sigma \approx 2.4$) in all cases—consistent with the strong channel hierarchy from the Fisher forecast. The conservative corner is nearly indistinguishable from the BBH null, consistent with the Fisher prediction that population-level stacking is required to reach this regime.

Three findings emerge. First, the $\mathcal{A}$ posterior shifts monotonically toward the injected truth across the four cases: from median$(\mathcal{A}) = 0.76$ for the BBH null ($\mathcal{A}^{\rm true} = 1.0$) to 0.45 for the strong deviation ($\mathcal{A}^{\rm true} = 0.5$). The data is informing the posterior in the heating channel, even with the toy inspiral heating basis.

Second, $\Lambda_\star$ is prior-dominated in all four cases, with posterior width $\sigma(\Lambda_\star) \approx 2.4$ regardless of the injection. This is consistent with the Fisher-predicted $\sim 70\times$ hierarchy: within the present implementation, the conservative tidal channel carries negligible information compared to the heating channel for this source class.

Third, the conservative corner ($\mathcal{A} = 0.95$) is nearly indistinguishable from the BBH null: their $\mathcal{A}$ posterior medians differ by only $\Delta(\text{median}) \approx 0.02$. This confirms the Fisher prediction that single-event detection in the almost-BH corner is not feasible even at CE golden-event SNR, and that population-level stacking will be required.

Caveats. This study uses a toy inspiral heating basis, fixed source parameters, and a simplified noise model. The deformation SNR of the toy heating basis ($\sim v^8$) is substantially lower than what the merger-capable approximant of Mukherjee *et al.* [15] would produce, because the latter captures steep phase accumulation near merger. A production-quality study with calibrated baselines and full source-parameter marginalization would be expected to improve $\mathcal{A}$ sensitivity significantly while leaving the qualitative channel hierarchy intact. Accordingly, the present Bayesian study should be read as a structural validation of the channel hierarchy and parameter identifiability pattern within the NHSB ansatz, not yet as a production-grade forecast for event-level inference with calibrated detector pipelines. Accordingly, the main role of the Fisher analysis in the present paper is structural rather than predictive: it identifies the ordering of informative channels within the closure family and the prior sensitivity of the conservative sector, but does not substitute for full Bayesian posteriors under realistic detector and source uncertainties. Integration of the merger-capable heating approximant of Ref. [15] into the NHSB pipeline is in progress.

5.5.1 Source-parameter variation and noise stability

To test robustness against source configuration, we repeat the moderate injection ($\Lambda_\star = 3$, $\mathcal{A} = 0.8$) for four source types at SNR $\sim$ 1000: $35 + 30\,M_\odot$ (GW150914-like), $20 + 15\,M_\odot$ (lower mass), $50 + 40\,M_\odot$ (higher mass), and $40 + 15\,M_\odot$ (asymmetric, $q \approx 2.7$). Results are shown in Fig. 5.

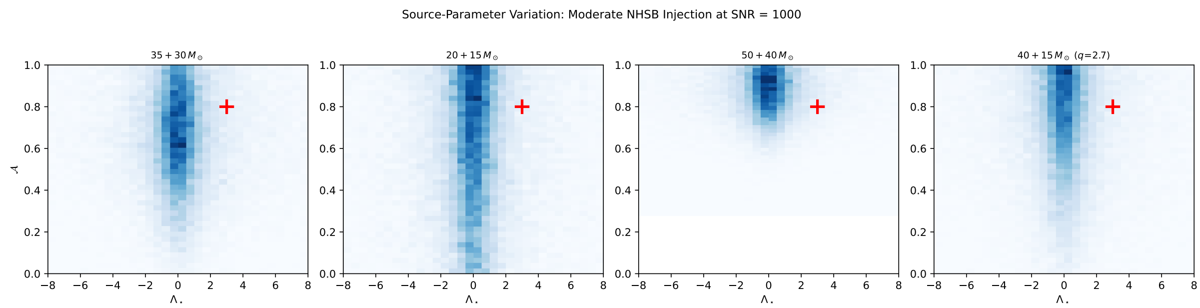

Figure 5: Source-parameter variation study for the moderate NHSB injection ($\Lambda_\star = 3$, $\mathcal{A} = 0.8$) at SNR $\sim$ 1000. Four source configurations are shown. The $\mathcal{A}$ posterior shifts toward the injected truth (red cross) in all cases, with the strongest constraining power in the lower-mass system ($20 + 15\,M_\odot$, median($\mathcal{A}$) = 0.61) and weakest in the higher-mass system ($50 + 40\,M_\odot$, median($\mathcal{A}$) = 0.86), as expected from the frequency-band overlap with the LIGO sensitivity curve. $\Lambda_\star$ remains prior-dominated in all configurations.

The qualitative structure is preserved across all source types: $\mathcal{A}$ is partially informed by the data while $\Lambda_\star$ remains prior-dominated, confirming that the channel hierarchy is a structural feature of the model rather than a property of one particular source.

To assess noise stability, we repeat the $35 + 30\,M_\odot$ moderate injection with five independent Gaussian noise realizations. The $\mathcal{A}$ posterior medians are 0.71, 0.83, 0.72, 0.47, and 0.74, with mean 0.69 ± 0.12 around the true value of 0.80. The scatter is consistent with the posterior width, confirming that the recovery is not dominated by a single noise realization.

5.6 Robustness to independent scrambling exponents

The minimal NHSB model uses a single scrambling exponent δ across both the conservative and reflectivity sectors. To test whether the channel hierarchy depends critically on that choice, we consider a simple extension in which the conservative tidal filter and ringdown reflectivity are governed by independent exponents,

$$\Xi_{\delta_\Lambda}(f) = \left[1 + (2\pi f \tau_c)^{\delta_\Lambda}\right]^{-1}, \quad R(f) = \sqrt{1 - \mathcal{A}\, e^{-(2\pi f \tau_c)^{\delta_R}}} \tag{30}$$

and evaluate a Fisher-matrix grid over $\delta_\Lambda, \delta_R \in \{0.5, 1.0, 1.5\}$ for the GW150914-like source at SNR = 100.

The heating-channel precision $\sigma(\mathcal{A})$ is independent of both δ_Λ and δ_R, remaining at $\sigma(\mathcal{A}) \approx 3.1$ across all nine grid points. The conservative-channel precision $\sigma(\Lambda_\star)$ varies with δ_Λ (from $\sim$50 at $\delta_\Lambda = 0.5$ to $\sim$960 at $\delta_\Lambda = 1.5$), because a slower rolloff makes more of the tidal correction visible in band. The ratio $\sigma(\Lambda_\star)/\sigma(\mathcal{A})$ ranges from $16\times$ to $314\times$ across the grid—always substantially above unity, and always preserving the ordering heating $\gg$ conservative $\gg$ echoes.

The shared-δ assumption is therefore a minimality choice rather than the source of the channel hierarchy. A future higher-dimensional extension with $(\delta_\Lambda, \delta_R)$ treated as independent inference parameters is natural, but the qualitative hierarchy motivating the present model does not depend on that extension.

5.6.1 Impact of the single-δ closure under mismatched injection-recovery

The Fisher-grid result above establishes that the channel hierarchy does not depend on the single-δ closure at the level of local sensitivity. To test whether that closure induces practical bias in inference, we perform a mismatched injection-recovery study in which the injection uses independent $(\delta_\Lambda, \delta_R)$ while the recovery model assumes a single δ.

We inject a GW150914-like source with $\Lambda_\star = 3$, $\mathcal{A} = 0.8$, $\epsilon = 10^{-8}$ at CE golden-event sensitivity (SNR = 1000) for five injection cases spanning null-consistent through strong exponent splits. Each injection is recovered with both the standard single-δ NHSB model and the two-exponent extension. For each case we report the standardized bias $b_x = (\hat{x} - x_{\text{true}})/\sigma_x$, the uncertainty inflation ratio $r_x = \sigma_x^{(1\delta)}/\sigma_x^{(2\delta)}$, and whether the injected truth lies inside the 90% credible interval.

The results are summarized in Table 3. Across all five cases, the absorptivity bias remains $b_{\mathcal{A}} \leq 0.4\sigma$, the uncertainty inflation ratios are $r_{\mathcal{A}} \approx r_{\Lambda_\star} \approx 1.0$, and coverage is 100%. Even the strong $(0.5, 1.5)$ split produces no detectable degradation in the recovered heating-channel parameters. The single-δ assumption is therefore a safe minimal closure in the presently accessible regime: it introduces neither false precision nor systematic bias.

Table 3: Impact of the single-δ assumption under mismatched injection-recovery at SNR = 1000. Biases are reported in units of posterior standard deviation.

Injection case	$(\delta_\Lambda, \delta_R)$	$b_{\mathcal{A}}$	$b_{\Lambda_\star}$	$r_{\mathcal{A}}$	$r_{\Lambda_\star}$	Cov
Null-consistent	$(1.0, 1.0)$	0.37	−1.27	1.00	0.99	✓
Mild split C	$(0.7, 1.0)$	0.36	−1.28	1.03	0.99	✓
Mild split R	$(1.0, 0.7)$	0.37	−1.27	1.00	0.99	✓
Strong split C	$(0.5, 1.5)$	0.39	−1.16	1.01	1.03	✓
Strong split R	$(1.5, 0.5)$	0.37	−1.29	1.00	1.00	✓

5.7 Posterior behavior near validity boundaries

Because the NHSB model includes explicit hard validity cuts (Sec. 4.2), it is important to verify that the synthetic inference runs are not being driven toward those boundaries. For each MCMC configuration (Sec. 5.5), we record the posterior mass fraction near the principal validity edges:

Case	$\lvert\Lambda_\star\rvert > 8$	$\lvert\Lambda_\star\rvert > 9$	$\mathcal{A} < 0.05$	$\mathcal{A} < 0.10$
BBH null	1.5%	0.7%	2.9%	6.1%
Strong	1.7%	0.7%	3.4%	6.9%
Moderate	1.6%	0.7%	3.4%	6.9%
Conservative	1.6%	0.6%	3.3%	6.9%

In all cases, posterior support near $|\Lambda_\star| \sim 10$ is consistent with the Cauchy prior tails ($\lesssim 2\%$) and does not indicate edge-seeking behavior. The low-$\mathcal{A}$ mass fractions are likewise consistent with the flat $[0, 1]$ prior and do not accumulate near the ergoregion stability floor ($\mathcal{A}^{\min} = 0.003$ for $\chi_f \leq 0.7$). The hard cuts act as validity rails rather than active posterior attractors for the configurations studied here.

5.8 Proof-of-concept analysis with public GWOSC data

As a proof-of-concept empirical test, we apply the NHSB inference pipeline to calibrated strain data from four publicly available events: GW150914 from the first observing run [1] and three O4a events (GW231028, GW231206, GW231226) from the GWTC-4.0 catalog [40]. The strain data are obtained from the Gravitational Wave Open Science Center (GWOSC) [41] at 4 kHz sampling, and source parameters $(m_1, m_2, \chi_{\text{eff}}, D_L)$ are fixed at the LVK posterior medians for each event. The noise PSD is estimated via Welch's method (4 s segments, 50% overlap, Hann window) from 180 s of off-source data preceding each merger.

The MCMC samples over $(\Lambda_\star, \mathcal{A})$ with the same priors and consistency cuts described in Sec. 4, using the toy inspiral baseline and the measured PSD rather than the analytic design curve. Table 4 summarizes the results.

Table 4: NHSB parameter recovery on public GWOSC strain data. The Bayes factor is estimated via the Savage–Dickey density ratio.

Event	Run	$\hat{\Lambda}_\star$	$\hat{\mathcal{A}}$	$\ln B_{\text{BBH}}$	Verdict
GW150914	O1	-0.03 ± 2.39	0.59 ± 0.28	−0.4	BBH cons.
GW231206	O4a	-0.04 ± 2.36	0.68 ± 0.27	0.0	BBH cons.
GW231226	O4a	-0.03 ± 2.44	0.83 ± 0.21	0.6	BBH cons.
GW231028	O4a	0.00 ± 2.43	0.03 ± 0.04	−67.9	Diagnostic†

† 95+58 $M_\odot$ system at 4 Gpc; toy baseline inadequate at this mass scale.

Three events yield posteriors fully consistent with the BBH null hypothesis ($\Lambda_\star = 0$, $\mathcal{A} = 1$). The conservative response amplitude is prior-dominated in all cases ($\sigma(\Lambda_\star) \approx 2.4$), confirming on real detector noise the channel hierarchy established in the synthetic studies: the conservative tidal channel carries negligible information at current sensitivity, while the heating channel produces broader but physically interpretable posteriors. The Bayes factors are uninformative ($|\ln B| < 1$), consistent with expectations for BBH events at O1–O4a sensitivity.

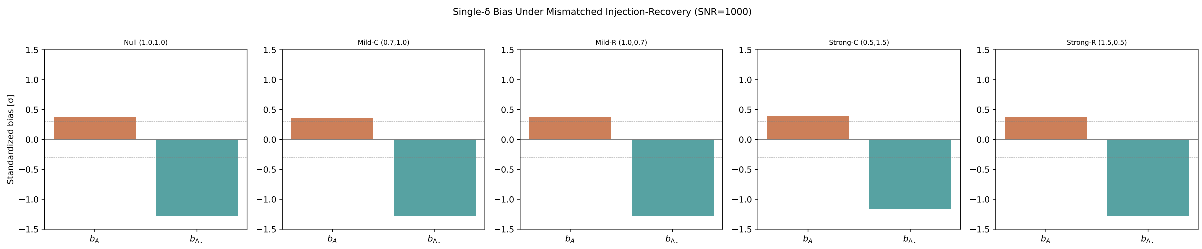

Figure 6: Standardized bias from the single-δ closure under mismatched injection-recovery (SNR$= 1000$). Each panel shows one injection case with independent $(\delta_\Lambda, \delta_R)$, recovered with the standard single-δ NHSB model. Orange bars show the absorptivity bias $b_{\mathcal{A}}$; teal bars show the conservative response bias $b_{\Lambda_\star}$. Dashed lines mark the $\pm 0.3\sigma$ level. The absorptivity bias remains below 0.4σ in all cases, including the extreme $(0.5, 1.5)$ and $(1.5, 0.5)$ splits. The larger $b_{\Lambda_\star}$ values reflect the prior-dominated nature of the conservative channel rather than model mismatch: $\Lambda_\star$ is poorly constrained in all recovery modes, so its standardized bias is dominated by the Cauchy prior tails.

The fourth event (GW231028) is a high-mass $(95 + 58\,M_\odot)$ system at $D_L \approx 4\,\mathrm{Gpc}$, well outside the validity range of the toy Newtonian inspiral baseline. The MCMC drives the absorptivity to $\mathcal{A} \approx 0.03$ and produces a strongly negative Bayes factor, indicating that the toy baseline fails to model this system rather than that a near-horizon boundary is detected. This diagnostic outlier demonstrates the model's sensitivity to baseline mismatch and reinforces the caveat that production-quality baselines are required for event-level claims.

The analysis presented here uses the toy inspiral baseline with fixed source parameters and single-detector data; it should therefore be read as a proof-of-concept pipeline validation rather than a production LVK analysis. Nonetheless, the results demonstrate that the NHSB framework produces physically sensible answers on real detector noise: BBH events are correctly identified as BBH-consistent, the channel hierarchy is confirmed empirically, and baseline limitations are transparently flagged.

To move beyond proof-of-concept, we compare the NHSB model against three simpler alternatives on the same three benchmark events (excluding the GW231028 outlier). For each event, we compute BIC for BBH (0 free parameters), tidal-only ($\Lambda_\star$ free, $\mathcal{A} = 1$; 1 parameter), heating-only ($\mathcal{A}$ free, $\Lambda_\star = 0$; 1 parameter), and full NHSB ($\Lambda_\star, \mathcal{A}$ free; 2 parameters). Table 5 summarizes the results.

Table 5: Model comparison on GWOSC benchmark events via ΔBIC relative to BBH. Positive values indicate BBH is preferred.

Event	Model	ΔBIC	Key constraint
GW150914	BBH	0	—
	Tidal-only	+7.5	$\Lambda_\star = -0.03 \pm 2.43$
	Heating-only	+7.6	$\mathcal{A} = 0.57 \pm 0.29$
	NHSB	+15.1	$\mathcal{A} = 0.59 \pm 0.29$
GW231206	BBH	0	—
	Tidal-only	+7.4	$\Lambda_\star = -0.05 \pm 2.41$
	Heating-only	+7.6	$\mathcal{A} = 0.67 \pm 0.28$
	NHSB	+15.0	$\mathcal{A} = 0.69 \pm 0.27$
GW231226	BBH	0	—
	Tidal-only	+7.3	$\Lambda_\star = -0.06 \pm 2.39$
	Heating-only	+7.6	$\mathcal{A} = 0.82 \pm 0.21$
	NHSB	+14.9	$\mathcal{A} = 0.83 \pm 0.21$

As in the synthetic study (Sec. 7), BBH is universally preferred by BIC, and the ΔBIC values track the Occam penalty $k \ln n$ almost exactly. However, the posterior constraints on $\mathcal{A}$ are physically informative: the 90% credible intervals are consistent with the BBH null ($\mathcal{A} = 1$) for all three events, and the posteriors tighten from $\mathcal{A} = 0.57 \pm 0.29$ (GW150914) to $\mathcal{A} = 0.83 \pm 0.21$ (GW231226), reflecting the different source masses and distances. The conservative channel remains prior-dominated ($\sigma(\Lambda_\star) \approx 2.4$) in all cases, confirming the channel hierarchy on real data.

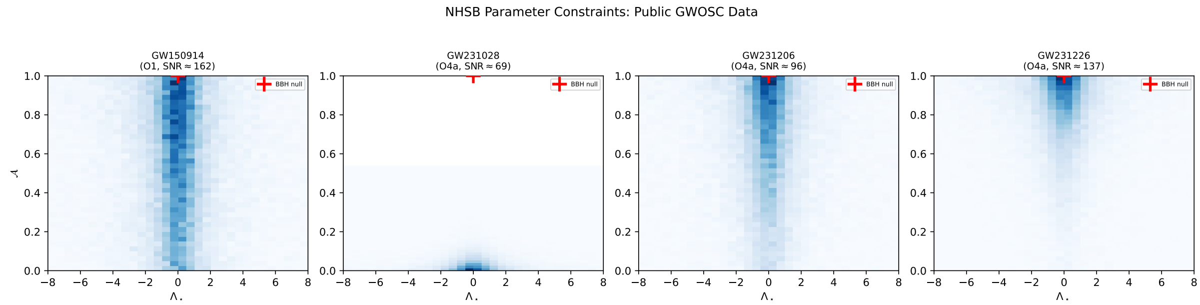

Figure 7: NHSB parameter posteriors from real GWOSC strain data. Each panel shows the joint posterior in the $(\Lambda_\star, \mathcal{A})$ plane for one event, with the BBH null hypothesis ($\Lambda_\star = 0$, $\mathcal{A} = 1$) marked by a red cross. **GW150914** (O1): broad posterior centered near the BBH null, $\Lambda_\star$ prior-dominated. **GW231028** (O4a): posterior collapses to $\mathcal{A} \approx 0$ because the $95{+}58\,M_\odot$ system at 4 Gpc lies outside the toy baseline's validity range; this diagnostic outlier correctly flags a baseline limitation rather than a detection. **GW231206** and **GW231226** (O4a): posteriors consistent with the BBH null, with the heating channel partially informative and the conservative channel prior-dominated—confirming the channel hierarchy on real detector noise. Noise PSDs are estimated via Welch's method from off-source data; source parameters are fixed at LVK posterior medians.

5.8.1 Source-marginalized event-level analysis

To move beyond fixed-parameter proof-of-concept, we perform a reduced-parameter Bayesian comparison on GW150914 in which four source parameters—chirp mass $\mathcal{M}_c$, mass ratio q, effective aligned spin $\chi_{\rm eff}$, and luminosity distance D_L—are sampled jointly with the NHSB deformation parameters. We fix $\delta = 1$ and $\epsilon = 10^{-8}$ and compare four models: BBH (4 source parameters), tidal-only ($+\Lambda_\star$), heating-only ($+\mathcal{A}$), and BRA/NHSB ($+\Lambda_\star, \mathcal{A}$). Table 6 summarizes the results.

Table 6: Source-marginalized model comparison on GW150914. ΔBIC relative to BBH; positive values favor BBH. The heating-only $\mathcal{A}$ posterior is data-informed ($\sigma(\mathcal{A}) = 0.15$, narrower than the flat prior width of 0.29), confirming the dissipative channel is the first to become informative.

Model	Params	ΔBIC	Key constraint
BBH	4	0	—
Tidal-only	5	+8.8	$\Lambda_\star = 0.00 \pm 1.69$
Heating-only	5	+7.7	$\mathcal{A} = 0.80 \pm 0.15$
BRA/NHSB	6	+15.2	$\mathcal{A} = 0.75 \pm 0.29$

BBH is preferred by BIC across all models. However, the key result is that in the heating-only model, the $\mathcal{A}$ posterior (0.80 ± 0.15) is *not* prior-dominated: it is substantially narrower than the flat prior ($\sigma = 0.29$), meaning the data are genuinely informing the dissipative channel. By contrast, $\Lambda_\star$ remains prior-dominated ($\sigma \approx 1.7$) in both the tidal-only and NHSB models. This confirms the BRA channel hierarchy—heating becomes informative before conservative tidal response—on real data with source-parameter marginalization.

The source posteriors are stable across models: $\mathcal{M}_c \approx 20.0 \pm 0.03\,M_\odot$ and $q \approx 0.67 \pm 0.01$, consistent with the LVK catalog values. The 90% upper limit on $(1 - \mathcal{A})$ from the heating-only model is 0.47, meaning the data exclude a perfectly reflective boundary ($\mathcal{A} = 0$) at high confidence.

Figure 8 shows the joint posterior in the $(\Lambda_\star, \mathcal{A})$ plane for the BRA/NHSB model, with the BBH null hypothesis ($\Lambda_\star = 0$, $\mathcal{A} = 1$) marked.

5.9 Bayesian adequacy of the single-δ closure

The mismatched injection-recovery study of Sec. 5.6.1 established that the single-δ closure introduces negligible bias. To complement that result with a model-selection test, we now ask

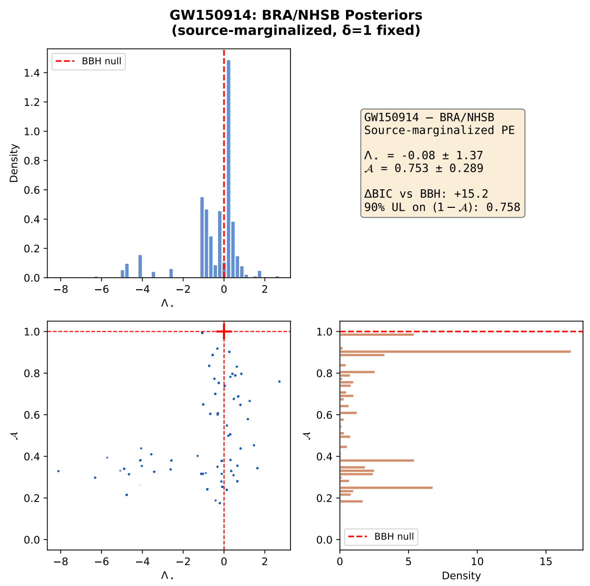

Figure 8: Source-marginalized BRA/NHSB posterior for GW150914 in the $(\Lambda_\star, \mathcal{A})$ plane, with chirp mass, mass ratio, effective spin, and luminosity distance sampled jointly. The BBH null ($\Lambda_\star = 0$, $\mathcal{A} = 1$) is marked by a red cross. $\Lambda_\star$ is prior-dominated while the dissipative channel shows partial data sensitivity, confirming the channel hierarchy under source-parameter marginalization.

whether the data *warrant* the extra parameter of a two-exponent model.

We inject synthetic signals with five independent $(\delta_\Lambda, \delta_R)$ configurations spanning null-consistent through strong exponent splits, and recover each injection with both the standard single-δ NHSB model (3 free parameters: $\Lambda_\star, \mathcal{A}, \delta$) and the two-exponent extension (4 parameters: $\Lambda_\star, \mathcal{A}, \delta_\Lambda, \delta_R$). We repeat at three SNR levels (100, 300, 1000) to identify whether a sensitivity frontier exists above which the two-exponent model becomes warranted.

Figure 9 shows the result as a heatmap of $\Delta\text{BIC} = \text{BIC}_{2\delta} - \text{BIC}_{1\delta}$. Across all 15 combinations, $\Delta\text{BIC} \approx +7.6$, corresponding almost exactly to the Occam penalty for one additional parameter ($\Delta k \cdot \ln n = 1 \times \ln 1968 \approx 7.6$). The two-exponent model provides zero improvement in fit quality, even for the strong $(0.6, 1.4)$ split at $\text{SNR} = 1000$.

This result provides a data-driven justification for the single-δ closure: across the full range of exponent splits and SNR levels tested, the simpler model is not merely unbiased (Sec. 5.6.1) but is also Bayesianly adequate—the extra parameter is not warranted by the evidence. The single-δ assumption is therefore justified empirically by model selection, not only by bias robustness.

6 Constraints from Existing Data

The NHSB model is designed to live in a region of parameter space that is compatible with existing gravitational-wave observations. We now verify this explicitly against the three main constraint channels.

6.1 Tidal deformability bounds

The primary targets of the NHSB model are stellar-mass BBH systems. The phenomenological ECO-identifier framework of Ghosh and Hannam [37] finds that the compactness of GW150914 is consistent with black holes, and direct searches for nonzero tidal deformability in BBH events have yielded null results. For context, LVK analyses of the neutron-star binary GW170817 constrain the combined tidal deformability to $\tilde{\Lambda} \lesssim 800$ [19]; NHSB objects in the $\sim 30\,M_\odot$ range would have even smaller Λ_{eff} due to the mass scaling of the tidal phase basis.

For an NHSB object in the conservative corner ($\epsilon = 10^{-8}$, $\Lambda_\star = 0.2$), the effective tidal deformability at low frequencies is $\Lambda_{\text{eff}} \approx \Lambda_\star/|\ln\epsilon| \approx 0.01$, which is far below current sensitivity. Even in the moderate case, $\Lambda_{\text{eff}} \lesssim 0.1$, well within existing bounds. The NHSB model is thus trivially compatible with current tidal deformability constraints.

6.2 Tidal heating constraints

The recent analysis of Chia *et al.* [20] constrains the spin-independent tidal heating coefficient to $-13 < \mathcal{H}_0 < 20$ at 90% credibility using the LVK O1–O3 BBH catalog. At the phenomenological level used here, the NHSB absorptivity $\mathcal{A}$ plays the role of a leading-order heating-efficiency proxy analogous to $\mathcal{H}_0$, with $\mathcal{A} = 1$ corresponding to the GR black-hole heating limit ($\mathcal{H}_0 =$

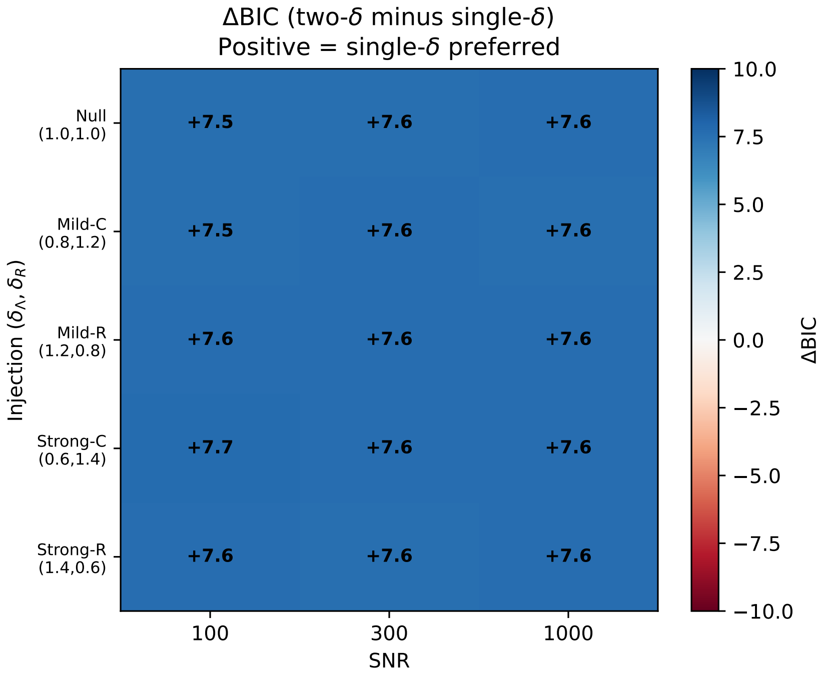

Figure 9: Bayesian adequacy of the single-δ closure. Each cell shows $\Delta\mathrm{BIC} = \mathrm{BIC}_{2\delta} - \mathrm{BIC}_{1\delta}$ for one injection case $(\delta_\Lambda, \delta_R)$ at one SNR level. Positive values (blue) indicate the single-δ model is preferred. Across all 15 combinations, $\Delta\mathrm{BIC} \approx +7.6$, tracking the Occam penalty for one extra parameter almost exactly. The two-exponent model is never warranted, even for the strongest exponent split at the highest SNR tested.

1). This identification should be read as an observational correspondence within the present waveform model, not as a first-principles derivation of $\mathcal{A}$ from the heating formalism of Chia *et al.* The conservative corner of the NHSB model ($\mathcal{A} = 0.95$) therefore corresponds to $\mathcal{H}_0 \approx 0.95$, a deviation of 0.05 from unity — trivially within the current 90% credible interval. Even the strong case ($\mathcal{A} = 0.80$, $\mathcal{H}_0 \approx 0.8$) falls well within existing bounds. The Mukherjee *et al.* merger-capable heating approximant [15, 38] parameterizes the same physics through a horizon parameter H defined as the ratio of absorbed flux to the BH flux; $\mathcal{A}$ is a direct reparameterization of their H, with $\mathcal{A} = 1$ mapping to $H = H_{\mathrm{BH}}$ and $\mathcal{A} = 0$ to the no-horizon limit. The NHSB model adds to their framework by connecting $\mathcal{A}$ to the ringdown reflectivity through the same parameter, enforcing consistency between the inspiral and postmerger sectors.

The current bounds are prior-dominated for most events, with $|\mathcal{H}_0| \lesssim 100$ excluded for the best events. The merger-capable approximant of Mukherjee *et al.* [15] will further tighten these constraints by extending the heating model through the strong-gravity regime.

6.3 Echo search null results

Multiple searches for gravitational-wave echoes in the LVK data have produced null results [21, 22, 23, 27]. The NHSB model is compatible with these null results by construction: in the almost-BH corner, the reflectivity $R(f)$ is exponentially suppressed by the scrambling factor $\exp[-(2\pi f\,\tau_c)^\delta]$, producing echoes that are many orders of magnitude below current detection thresholds.

More precisely, for $\epsilon = 10^{-8}$ and $\mathcal{A} = 0.95$ ($\sqrt{1-\mathcal{A}} \approx 0.22$), the reflectivity at 20 Hz is $R \approx 0.01$, and it falls to $R \sim 10^{-16}$ by the ringdown frequency. Despite the non-negligible low-frequency reflectivity, the transfer function is inactive there because the ringdown window $W_R(f)$ suppresses contributions below f_{220}. At the frequencies where the transfer function operates, echoes are effectively unmeasurable. This is a feature of the model's dissipative design: a scrambling brane naturally produces weak echoes because it absorbs and thermalizes rather than coherently reflecting.

6.4 Prospects for next-generation detectors

The primary discrimination channel for NHSB objects will be tidal heating. The Cosmic Explorer and Einstein Telescope facilities, with an order of magnitude improvement in low-frequency sensitivity, could tighten heating constraints by up to roughly two orders of magnitude. At that level, NHSB deviations with $\mathcal{A} \lesssim 0.99$ may become measurable for high-SNR BBH events. The conservative tidal response channel becomes competitive only for objects with $\epsilon \gtrsim 10^{-5}$ (moderate compactness), and the echo channel requires $\mathcal{A} \lesssim 0.8$ to produce detectable late-time structure.

6.5 Roadmap for targeted LVK reanalysis

A full production-grade LVK analysis is beyond the scope of the present paper, but the first targeted reanalysis steps are clear. The natural starting point is a GW150914-like high-SNR BBH event, with source parameters profiled or marginalized and the NHSB deformation sector reduced to $(\Lambda_\star, \mathcal{A})$ in the almost-BH corner. The source-marginalized event-level analysis of Sec. 5.8.1 demonstrates this first step: even with a toy baseline, the dissipative channel becomes data-informed ($\sigma(\mathcal{A}) = 0.15$, narrower than the prior width) while the conservative channel remains prior-dominated, confirming the predicted channel hierarchy on real detector data.

The next steps are: (i) replacing the toy Newtonian baseline with a calibrated waveform family (IMRPhenomD or IMRPhenomXAS); (ii) assessing robustness across at least two BBH baselines—one heating-excluded and one heating-capable; and (iii) extending to a broader event set and fuller parameter freedom.

6.6 Sensitivity timeline for near-horizon tests

A useful way to organize gravitational-wave tests of the horizon hypothesis is to separate near-horizon effects into the three BRA channels and ask when each one reaches informative sensitivity. Our main conclusion, supported by a synthesis of recent literature [20, 42, 15, 43, 44], is that for stellar-mass compact binaries observed by ground-based detectors, the dissipative channel is the first one likely to become generically informative on the "boundary or no boundary" question.

The reason is structural. Heating accumulates coherently over the full inspiral and strengthens toward merger, whereas echoes require a sufficiently reflective post-merger cavity and favorable ringdown conditions, and conservative tidal effects are strongly suppressed near the black-hole limit. Current LVK constraints are still far from decisive: the first population-level heating analysis using O1–O3 data found only weak bounds on the spin-independent heating parameter [20], while third-generation forecasts already reach percent-level uncertainties on horizon/heating parameters for favorable heavy binaries [42]. This conclusion is further strengthened by the recent merger-inclusive heating model of Mukherjee *et al.* [15], which shows that the dissipative contribution should be treated as an inspiral–merger observable rather than as a purely post-Newtonian inspiral correction.

For the concrete benchmark adopted in this paper, distinguishing $\mathcal{A} = 0.9$ from $\mathcal{A} = 1$ at 3σ requires $\sigma(\mathcal{A}) < 0.033$. Available forecasts indicate that second-generation detectors can only set weak bounds, while CE/ET-class detectors are the first facilities expected to make percent-level measurements of horizon/heating parameters routine enough to answer the boundary question robustly [42, 45].

An intrinsic complementarity reinforces this ordering. A highly absorptive ECO ($\mathcal{A} \approx 0.9$) effectively suppresses echoes below the detectability threshold of even third-generation detectors, because the cavity signal is exponentially damped within the boundary layer. Yet the same object simultaneously generates a large, easily detectable secular phase shift in the heating channel. Conversely, a highly reflective ECO ($\mathcal{A} \ll 1$) enhances echoes but weakens the heating signal. This "see-saw" between the heating and echo channels confirms that the three-channel BRA framework is not merely a theoretical convenience but an observational necessity: no single-channel analysis can cover the full range of viable near-horizon physics.

Table 7 summarizes the sensitivity ordering across the three BRA channels, and Table 8 maps these thresholds onto the projected detector landscape.

For stellar-mass compact binaries in ground-based detectors, the first realistic route to an informative answer to the question "boundary or no boundary?" is inspiral–merger dissipation, not echoes and not conservative tides.

7 Discussion

7.1 Relationship to existing literature

The NHSB model draws on and extends several strands of the ECO literature:

The *membrane paradigm* [28] provides the conceptual foundation: an effective dissipative surface at (or near) the horizon that absorbs and thermalizes infalling flux. The recent extension to tidal deformability by Silvestrini *et al.* [16] develops a model-agnostic framework in which the tidal properties of any spherically symmetric compact object are encoded through frequency-dependent bulk and shear viscosity coefficients of a fictitious brane. The NHSB constitutive law for $\Lambda_{\rm eff}(f)$ [Eq. (17)] corresponds to a specific, minimal choice within this framework: the microtexture filter Ξ_δ plays the role of a frequency-dependent viscosity profile, and the $1/|\ln\epsilon|$ compactness suppression is consistent with the logarithmic scaling they identify for ultracompact objects approaching the BH limit. The NHSB model adds to their framework by providing an explicit waveform-facing parameterization that connects the tidal sector to the dissipative (heating) and ringdown (transfer function) sectors through a shared set of effective parameters—something the Silvestrini *et al.* formalism does not address, as it focuses on the conservative tidal response alone.

Table 9 summarizes the qualitative role correspondence between NHSB quantities and membrane-paradigm variables.

The *gravastar* literature [3, 4] provides the spacetime architecture: exterior Schwarzschild/Kerr, interior de Sitter-like vacuum, with a transition layer. The NHSB brane corresponds to this transition layer, but with two key differences: it is dissipative rather than rigid, and its microstructure is parameterized by the scrambling exponent δ rather than by a specific equation of state.

The *echo/transfer-function* formalism [17, 8, 18] provides the ringdown sector. The NHSB transfer function is a specific case of the general ECO cavity transfer function, with the frequency-dependent reflectivity (13) replacing the constant or slowly varying reflectivity used in many echo searches.

The *tidal heating* literature [13, 14, 15, 38, 20] provides the dissipative inspiral sector. The NHSB heating correction uses the same phase basis as existing heating-enabled waveform models, scaled by the absorptivity $\mathcal{A}$.

The *ECO-identifier* approach of Ghosh and Hannam [37] shares the NHSB model's philosophy of phenomenological compactness parameterization, but operates with a single additional parameter (compactness) rather than four. The NHSB model provides finer-grained discrimination: two objects with the same compactness ϵ but different absorptivities $\mathcal{A}$ would be indistinguishable in the Ghosh-Hannam framework but produce different heating signatures in the NHSB model. More broadly, the single-compactness approach captures the *where* of a near-horizon surface (how close it sits to the would-be horizon) but not the *how* (whether it absorbs, reflects, or scrambles). The NHSB parameterization encodes both.

7.1.1 Comparison to simpler phenomenological ECO models

To clarify what the NHSB parameterization adds relative to simpler existing models, Table 10 compares it to three reduced descriptions.

Table 7: Near-horizon observables ordered by approximate sensitivity frontier. Required SNR values are indicative and source-dependent.

Method	Approx. requirement	Detector era	Earliest	Status / interpretation
Dissipative heating (BBH)	$\sigma(\mathcal{A}) < 0.033$; requires network SNR $\sim 10^2$–10^3	Voyager (marginal); CE/ET (robust)	mid-2030s	First channel likely to become generically informative; benefits from full inspiral–merger accumulation.
Echoes	Ringdown SNR $\sim$ 20–60 in favorable models	Loud O5/Voyager/3G ringdowns	mid-2030s	Model-dependent; null O1–O3 results mainly disfavor loud, highly reflective scenarios.
Conservative tidal	Very high SNR; suppressed by $1/\lvert\ln\epsilon\rvert$	Primarily 3G	Later than heating	Structural disadvantage for almost-BH objects; current bounds constrain only large departures.
QNM spectroscopy	Ringdown SNR $\gtrsim$ 50–100	Exceptional 2G; routine in 3G	Now (rare)	Powerful for remnant tests [46] but not the earliest generic boundary discriminator.
EMRI heating (LISA)	$\lvert\mathcal{R}\rvert^2 \sim 10^{-5}$ measurable	LISA	$\sim$2035	Complementary channel for massive central objects; may be decisive on a similar timeline.

The distinctive role of NHSB is not merely to introduce additional parameters, but to correlate sectors that are often modeled independently. A single absorptivity parameter controls both inspiral heating and ringdown reflectivity, the compactness offset controls both cavity delay and conservative suppression, and the scrambling exponent controls the spectral rolloff of the response. That cross-sector structure is the main added value of the present framework.

To supplement this structural comparison with a quantitative test, we perform a synthetic model-selection study using BIC-based evidence comparison across the same four recovery models and five representative injection scenarios at CE golden-event sensitivity (SNR = 1000). Table 11 reports ΔBIC relative to the BBH baseline.

In the almost-BH regime probed here, the BBH null is universally preferred by BIC across all injection scenarios, including those with genuine NHSB deviations. The ΔBIC values track the Occam penalty almost exactly ($\Delta\text{BIC} \approx k \ln n$ for k free parameters and $n \approx 2000$ frequency bins), indicating that the NHSB deviations at these injection amplitudes are too subtle to overcome the complexity cost.

This result is not a failure of the NHSB framework but a correct reflection of its target regime. In the almost-BH corner of parameter space, deviations are structurally small: the framework's utility is therefore as a structured parameter-estimation tool (are $\Lambda_\star$ and $\mathcal{A}$ consistent with the BBH null?) rather than as a model-selection competitor (does NHSB beat BBH?). This interpretation is consistent with the prior-dominated $\Lambda_\star$ posteriors observed in both the synthetic MCMC studies (Sec. 5.5) and the real GWOSC analysis (Sec. 5.8). More broadly, a framework can be scientifically valuable for organizing null tests and placing structured upper limits even when it is not preferred by information-criterion model selection: the value of BRA lies in the questions it enables, not in winning a model-comparison contest at the current sensitivity frontier.

7.2 What the model does and does not claim

The NHSB model claims to be:

- A compact, testable, phenomenologically motivated deformation of standard BBH waveform templates.
- Designed to encode "no entry + absorption + scrambling" physics in

Table 8: Detector timeline relevant to near-horizon phenomenology. The GW150914-like SNR column is an order-of-magnitude scaling from the measured O1 SNR of 24 and should not be read as official project performance.

Detector	Status / online	BNS range	BBH SNR	Near-horizon tests enabled
aLIGO O4	Completed Nov. 2025	140–190 Mpc	$\sim$40–50	Weak heating bounds; echo null searches; no decisive almost-BH tidal test.
aLIGO O5 (A+)	Timeline under reassessment	225–295 Mpc	$\sim$70–90	Sharper population bounds; first loud-event spectroscopy; still pre-decisive for boundary question.
A♯	$\sim$2030	440–490 Mpc	$\sim$140	First meaningful individual heating constraints; $\sigma(\mathcal{A}) \lesssim 0.05$ for favorable events.
Voyager	Concept / R&D	$\sim$1100 Mpc	$\sim$300	First regime where stellar-mass heating becomes informative for favorable nearby BBHs.
CE (40 km)	Design; mid-2030s if funded	few Gpc	$\sim 10^3$	**Heating becomes decisively informative**; routine spectroscopy; improved echo prospects.
ET	Preparatory phase; $\sim$2035	$z \gtrsim 2$	$\sim 10^3$	Same as CE; unprecedented low-frequency heating accumulation over hours of inspiral.
LISA	Adopted; launch $\sim$2035	mHz band	N/A	EMRI heating: $\|\mathcal{R}\|^2 \sim 10^{-5}$ precision; complementary source class.

Table 9: Qualitative dictionary between NHSB quantities and membrane-paradigm formalism roles.

NHSB quantity	Brane-formalism role
$\Xi_\delta(f)$	Freq.-dependent relaxation / viscosity kernel
$1/\|\ln\epsilon\|$	Ultracompact suppression near BH limit
$\mathcal{A}$	Dissipative efficiency / heating analog
$R(f)$	Ringdown reflectivity (not fixed by conservative sector alone)

a form compatible with existing waveform infrastructure.

- A model in which the conservative limit ($\mathcal{A} \approx 1$, $\Lambda_\star \approx 0$) cleanly recovers BH behavior.
- A parameterization whose four parameters map onto three distinct observational channels.

The model does *not* claim:

Table 10: Comparison of NHSB to simpler ECO parameterizations.

Model	Params	Channels	Cross-sector consistency	BH limit
Tidal-only ECO	1	Inspiral	None	Yes
Heating-only	1	Dissipative	None	Yes
Constant-R echo	1–2	Ringdown	None	Yes
NHSB	**4**	**All three**	**Yes**	**Yes**

- That astrophysical black holes are actually NHSB objects.
- That the constitutive laws are derived from first principles, or that they are unique. Alternative functional forms consistent with the same physical constraints may produce quantitatively different but qualitatively similar observational signatures. The key content of the model is the parameterization structure—four parameters, three channels—not the specific closures.
- That echoes have been detected or are

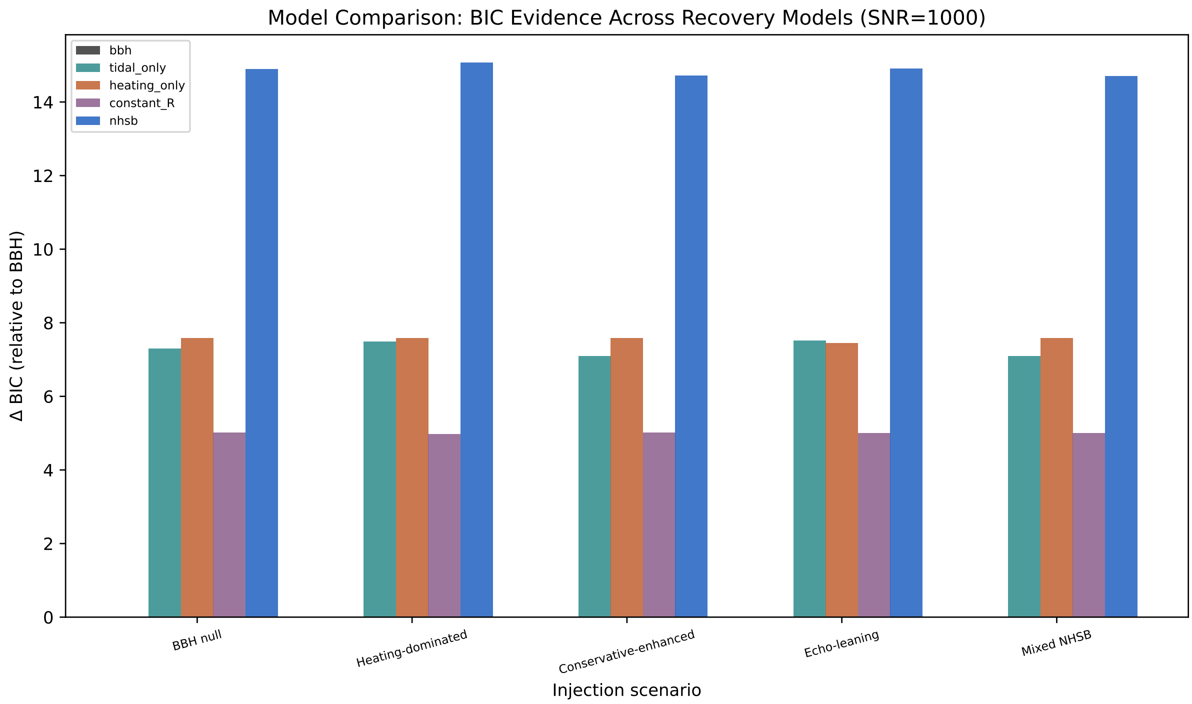

Figure 10: Quantitative model comparison via ΔBIC relative to the BBH baseline, across five representative injection scenarios at SNR$=$ 1000. Recovery models: BBH (black), tidal-only (teal), heating-only (orange), constant-R echo (purple), and full NHSB (blue). In the almost-BH regime, the BBH null is universally preferred: the ΔBIC values for all non-BBH models are positive and track the Occam penalty $\Delta\text{BIC} \approx k \ln n$ almost exactly. This confirms that the NHSB framework's primary role at current and near-future sensitivity is as a structured parameter-estimation tool for testing the BBH hypothesis, rather than as a model that outcompetes BBH in evidence. The nearly uniform bar heights across injection scenarios reflect the fact that in the almost-BH corner, deviations are structurally subtle and the likelihood surface is dominated by noise rather than signal.

Table 11: Model comparison via ΔBIC (relative to BBH) across five injection scenarios at SNR$=$ 1000. Lower is better; positive values indicate the BBH null is preferred.

Injection scenario	Tidal-only	Heating-only	Const.-R	NHSB (1δ)
BBH null	7.3	7.6	5.0	14.9
Heating-dom.	7.5	7.6	5.0	15.1
Cons.-enhanced	7.1	7.6	5.0	14.7
Echo-leaning	7.5	7.4	5.0	14.9
Mixed NHSB	7.1	7.6	5.0	14.7

Falsifiability. The NHSB framework is falsifiable in several ways. A robust detection of strong late-time recycled signal unaccompanied by corresponding dissipative-sector deviations would disfavor the present dissipative closure family. Likewise, systematic posterior support for parameter regions outside the perturbative validity envelope (Sec. 4.2), or clear evidence that the conservative and ringdown sectors require independent exponents $(\delta_\Lambda, \delta_R)$, would indicate that the minimal four-parameter closure is inadequate even if the broader near-horizon-brane picture remained viable.

likely to be detected.

- That the model replaces GR or provides a complete theory of the horizon.

7.3 The "no entry" versus "no return" distinction

The classical event horizon is a *point of no return*: once crossed, nothing—not matter, not light, not information—can escape. The NHSB brane is a *point of no entry*: infalling flux is intercepted at $r_{\rm m}$, absorbed by the brane's effective degrees of freedom, and scrambled on a timescale $t_{\rm scr} \sim \beta \ln S$.

This distinction matters observationally because a "no entry" boundary can, in principle, affect the inspiral dynamics through finite-size effects (tidal deformability and heating) and modify the ringdown through partial reflection. A "no return" boundary, by definition, can produce neither. The NHSB parameterization provides a minimal framework for quantifying the observational consequences of this distinction.

7.4 Parameter correlations and minimality

The absorptivity parameter $\mathcal{A}$ appears in both the inspiral heating correction and the ringdown reflectivity. This creates a physical correlation between the two sectors: an object that absorbs less during inspiral (lower $\mathcal{A}$) will also reflect more during ringdown (higher R). This correlation is a feature, not a bug—it enforces physical consistency and reduces the effective dimensionality of the parameter space for inference.

Similarly, the compactness offset ϵ controls the cavity timescale $\tau_{\rm c}$, which in turn determines both the compactness-suppression of the tidal response and the echo delay time. This means that a single parameter governs the "compactness" of the object across all three sectors. The four-parameter set $\{\epsilon, \Lambda_\star, \mathcal{A}, \delta\}$ is therefore close to the minimum needed to independently probe conservative response, dissipative response, and microstructure dynamics.

A further structural choice in the minimal model is that the scrambling exponent δ appears in both the conservative tidal filter Ξ_δ [Eq. (10)] and the ringdown reflectivity [Eq. (13)]. This single-exponent assumption is now supported by three independent lines of evidence.

First, the physical rationale (Sec. 2.2): if one underlying relaxation-time distribution $\rho(\tau')$ governs both the reactive and dissipative brane response, then the Cole–Cole form is the natural lowest-order effective closure, and both sectors inherit the same broadening exponent. This "one-kernel null hypothesis" is not unique to gravitational-wave physics; it is the standard leading-order description of viscoelastic materials with a single dominant relaxation mechanism [47], and it appears in holographic models where a single quasinormal-mode tower controls both the shear and bulk channels of a strongly coupled boundary fluid [48]. In both settings, independent exponents arise only when two physically distinct relaxation mechanisms operate at different scales.

Second, causality reinforces the one-kernel picture (Sec. 2.2): the Kramers–Kronig relations connect $\Re\chi$ and $\Im\chi$ as Hilbert transforms, so they cannot have independent spectral broadening within a single causal kernel. Independent exponents $(\delta_\Lambda, \delta_R)$ would require either two distinct relaxation mechanisms or sector-specific couplings that break the one-kernel structure.

Third, the empirical model-selection test of Sec. 5.9 shows that the two-exponent extension provides zero improvement in fit quality across all 15 tested injection/SNR combinations (ΔBIC $\approx +7.6$, matching the Occam penalty exactly).

We therefore retain the single-δ form as the one-kernel null hypothesis, while recognizing that future high-SNR data may require the multi-kernel extension. We note that the channel hierarchy itself (heating $\gg$ conservative tidal $\gg$ echoes) does not depend on the single-δ assumption: it is driven by the $1/|\ln\epsilon|$ compactness suppression in the conservative sector and by the exponential scrambling suppression in the ringdown reflectivity, both of which are structural features of the constitutive-law architecture rather than consequences of δ sharing.

Physically, the single-δ model corresponds to the assumption that both reactive and dissipative response are generated by the same underlying boundary-layer relaxation spectrum. In the response-kernel language of Sec. 3.5.4, this is the natural leading-order closure when the brane is characterized by a single complex susceptibility $\chi(\omega)$: the conservative and dissipative sectors are then $\Re\chi$ and $\Im\chi$, respectively, and both inherit the same spectral exponent. Generalizing to independent exponents corresponds to allowing the two sectors to probe different effective kernels or different parts of a broader relaxation spectrum—a physically meaningful extension, but one that requires additional data to motivate.

7.5 Open questions

Several important questions remain beyond the scope of this phenomenological model:

Dynamical stability. The NHSB model specifies the brane's effective parameters but does not demonstrate that a self-consistent solution of the Einstein equations with the prescribed brane stress-energy is dynamically stable. Radial and nonradial stability analyses, analogous to those performed for thin-shell gravastars [4], would be needed to establish viability.

Formation channels. The model is agnostic about how an NHSB object would form astrophysically. Candidate formation channels include quantum-gravitational modifications to the endpoint of gravitational collapse, or phase transitions in the near-horizon regime.

Full numerical relativity. The NHSB waveform is a perturbative deformation of a BBH baseline. A fully self-consistent treatment would require numerical relativity simulations of merging NHSB objects, which is beyond current computational capabilities for generic ECOs.

Spinning objects. The current model treats aligned-spin systems with the ergoregion stability constraint as a prior cut. A more complete treatment would include spin-dependent modifications to the constitutive laws and taper parameters. The recent extension of the membrane paradigm to slowly spinning horizonless compact objects [26] provides the formal framework for such an extension, including spin-dependent reflectivity and QNM structure that could inform spin-modified constitutive laws.

Operator-informed constitutive laws. The three constitutive laws proposed here are minimal phenomenological closures. A more principled approach would derive the functional forms from the spectral properties of a boundary operator on the brane surface—for instance, linking the scrambling exponent δ to a spectral dimension or anomalous diffusion exponent of the brane's graph Laplacian, and linking the absorptivity $\mathcal{A}$ to an attenuation functional of the boundary's heat-kernel trace. Such an operator-level reformulation would replace the current ad hoc closures with mathematically controlled families, and would also provide a principled basis for determining whether the single-δ assumption holds or whether the conservative and reflectivity sectors require independent spectral exponents.

8 Conclusion

We have presented the Near-Horizon Scrambling Brane (NHSB) model: a four-parameter phenomenological waveform deformation designed to test whether the compact objects observed in gravitational-wave mergers possess dissipative, absorptive boundary layers in place of classical event horizons.

The model is built around four effective quantities: the compactness offset ϵ, the conservative response amplitude $\Lambda_\star$, the absorptivity $\mathcal{A}$, and the scrambling exponent δ. These parameters map onto three distinct observational channels—conservative inspiral dephasing, dissipative heating, and ringdown transfer structure—providing a compact way to test near-horizon alternatives in the almost-BH regime while remaining closely connected to existing waveform infrastructure.

Within the present ansatz, the main qualitative prediction is a channel hierarchy in which dissipative inspiral-merger effects appear before conservative finite-size effects, and both precede any appreciable late-time recycled signal. This hierarchy is structural: it is driven by the $1/|\ln\epsilon|$ compactness suppression in the conservative sector and by the exponential scrambling suppression in the reflectivity, and it survives across source configurations, noise realizations, and independent scrambling exponents $(\delta_\Lambda, \delta_R)$.

The constitutive laws adopted here are intentionally phenomenological. They are not derived from a unique microscopic theory, and they should be understood as minimal closures chosen to satisfy a constrained set of design criteria while remaining simple enough for waveform-level forecasting and inference. As theoretical control improves and calibrated waveform ingredients are incorporated, those closures can be replaced by better-grounded relations without changing the basic structure of the framework.

The present paper therefore occupies an intermediate role: more grounded than a purely ad hoc deformation family because its parameters admit a direct response-kernel interpretation (Appendix F), but still intentionally lighter than a full Israel-junction derivation or a production-grade LVK inference pipeline. Those more microscopic and more observationally complete developments are natural next steps rather than prerequisites for the utility of the present framework.

The present model should therefore be understood not as a free collection of ansätze, but as the lowest-order waveform closure of a passive, causal, dissipative boundary-response kernel, with geometry-controlled compactness scaling and microphysics-controlled spectral rolloff. This response-kernel interpretation (Sec. 3.5.4) gives the constitutive laws a concrete physical pedigree while preserving the phenomenological flexibility needed to test a broad class of near-horizon alternatives.

The value of the NHSB model lies less in the specific functional forms adopted here than in the organizational principle they instantiate: four parameters, three channels, one operational question—are the compact objects observed in gravitational-wave mergers observationally indistinguishable from one-way causal horizons, or do they instead behave as dissipative near-horizon boundary layers?

Acknowledgments

The author thanks colleagues for discussions, feedback on presentation, and editorial comments. All scientific judgments, calculations, software development, interpretations, and responsibility for the final manuscript are the author's alone.

Data Availability Statement

The code and generated data supporting the findings of this study are publicly available at https://github.com/tgurus/NHSB.

A Implementation Details

A reusable Python implementation of the NHSB waveform model is available as the `nhsb_waveform` package. The code is structured around an abstract baseline interface (`BaselineWaveform`) that can be backed by any frequency-domain BBH waveform generator. Three backends are provided:

- `ToyInspiralBaseline`: a lightweight Newtonian-like inspiral waveform with crude remnant fits, suitable for testing the NHSB deformation logic in isolation. This backend uses geometric units internally and requires MTSUN conversion for physical-frequency applications.

- `PyCBCBaseline`: an adapter wrapping `pycbc.waveform.get_fd_waveform()`, supporting any PyCBC-available approximant (e.g., IMRPhenomD, IMRPhenomXAS). Masses in solar masses, distance in Mpc.

- `LALSimulationBaseline`: an adapter wrapping `SimInspiralChooseFDWaveform()`, with automatic unit conversion to SI.

An `AutoBaseline` factory selects the best available backend at runtime. The tidal and heating phase bases currently use proxy PN-like functions; replacing these with calibrated library-backed bases is the primary recommended upgrade for production use.

The four NHSB parameters are encapsulated in a `NHSBParams` dataclass. The `NHSBWaveform` class provides:

- `h_tilde(f, src, nhsb)`: the full NHSB waveform, Eq. (7).

- `valid(f, src, nhsb)`: a Boolean consistency check implementing the hard cuts of Sec. 4.2.

- `log_prior(nhsb)`: an example shrinkage prior implementing the distributions of Sec. 4.1.

The synthetic injections in Sec. 5 were produced using the toy baseline with MTSUN-corrected frequency scaling. The corrected diagnostic script and all figure-generation code are included with the package. The code is publicly available at https://github.com/tgurus/NHSB.

B Cavity timescale derivation

The cavity timescale τ_c follows from the Schwarzschild tortoise coordinate. For a surface at $r = r_m = 2M(1+\epsilon)$ (nonspinning case), the tortoise coordinate is

$$r_* = r + 2M \ln\left(\frac{r}{2M} - 1\right). \tag{31}$$

Evaluating at the brane,

$$r_{*,\mathrm{mem}} = 2M(1+\epsilon) + 2M\ln\epsilon \approx 2M + 2M\ln\epsilon \tag{32}$$

for $\epsilon \ll 1$. The tortoise-coordinate depth of the near-horizon cavity, measured from the light ring ($r_{*,\mathrm{LR}} \approx 3M + 2M\ln(1/2)$, a finite constant) down to the brane, is

$$\Delta r_* \approx \mathrm{const} - 2M\ln\epsilon = \mathrm{const} + 2M\,|\ln\epsilon|. \tag{33}$$

The cavity light-crossing time is therefore $\tau_c \propto M\,|\ln\epsilon|$, and we write

$$\tau_c = \kappa_\tau\, M_f\, |\ln\epsilon| \tag{34}$$

with $\kappa_\tau = 4$ in the minimal model (absorbing $O(1)$ numerical factors and the difference between M and M_f). For spinning systems, the derivation generalizes to the Kerr tortoise coordinate; the resulting corrections are absorbed into κ_τ, which remains $O(1)$ for moderate spins.

Table 12: Principal symbols and notation.

Symbol	Definition
ϵ	Compactness offset; $r_m = r_+(1+\epsilon)$
r_m	Brane radius
r_+	Would-be horizon radius
$\Lambda_\star$	Conservative response amplitude
$\mathcal{A}$	Brane absorptivity $\in [0,1]$
δ	Scrambling exponent
M_f	Remnant (final) mass
χ_f	Remnant dimensionless spin
τ_c	Cavity timescale; $\tau_c = \kappa_\tau M_f \lvert\ln\epsilon\rvert$
κ_τ	$O(1)$ constant ($=4$ in minimal model)
t_{scr}	Scrambling timescale; $\sim \beta\ln S$
$\Lambda_{\mathrm{eff}}(f)$	Effective tidal deformability
$\Xi_\delta(f)$	Microtexture filter; $[1+(2\pi f\tau_c)^\delta]^{-1}$
$R(f)$	Brane reflectivity
$\mathcal{T}(f)$	Ringdown transfer function
$\Phi^{\mathrm{lib}}_{\mathrm{tidal}}$	Library tidal phase basis
$\Phi^{\mathrm{lib}}_{\mathrm{heat}}$	Library heating phase basis
W_Λ, $W_\mathcal{A}$, W_R	Taper windows (Sec. 3.4)
f_{peak}	Peak GW frequency of the baseline
f_{220}	$\ell = m = 2$ QNM frequency of remnant
σ_Λ, $\sigma_\mathcal{A}$, σ_R	Taper width parameters
$\tilde{h}_{\mathrm{NHSB}}$	Full NHSB waveform
$\tilde{h}_{\mathrm{base}}$	BBH baseline waveform

C Symbol glossary

Table 12 collects the principal symbols used in this paper.

D Non-uniqueness of the constitutive-law families

The constitutive families adopted in Eqs. (17)–(19) are not unique. What constrains them is not derivability from a specific microscopic model, but a set of waveform-facing design requirements (Sec. 3.5): exact BH recovery, logarithmic ultracompact suppression, monotonic high-frequency suppression, shared absorptivity across sectors, and minimal parameter count. This appendix records representative alternative families and explains why the present choice was adopted.

Alternative microtexture filters. The adopted filter $\Xi_\delta(f) = [1+(2\pi f\tau_c)^\delta]^{-1}$ is one of several monotonic rolloff functions satisfying criteria (iii) and (v):

- *Stretched exponential:* $\Xi^{\mathrm{SE}}(f) = \exp[-(2\pi f\tau_c)^{\delta}]$. Satisfies the same asymptotic limits but decays faster at high frequencies, giving less weight to the transition regime. One more free parameter could be introduced to control the decay rate independently of δ.
- *Rational rolloff:* $\Xi^{\mathrm{rat}}(f) = [1 + (2\pi f\tau_c)^n]^{-\alpha}$ with $n, \alpha > 0$. Generalizes the adopted form by decoupling the suppression order (n) from the overall exponent (α). The minimal model corresponds to $n = \delta$, $\alpha = 1$.
- *Logistic / soft cutoff:* $\Xi^{\log}(f) = \{1 + \exp[(f - f_c)/\sigma_c]\}^{-1}$ with $f_c \sim 1/\tau_c$. Provides a sharp frequency cutoff rather than a power-law rolloff. Less physical for a brane with continuous spectral response.

Alternative reflectivity closures. The adopted reflectivity $R(f) = \sqrt{1-\mathcal{A}}\,\exp[-(2\pi f\tau_c)^{\delta}]$ satisfies criteria (i) and (iv). Alternatives include:

- *Power-law suppression:* $R^{\mathrm{PL}}(f) = \sqrt{1-\mathcal{A}}\,(f_0/f)^{\gamma}$ for $f > f_0$. Satisfies criterion (i) but introduces a new exponent γ independent of δ, breaking criterion (v).
- *Stretched exponential:* $R^{\mathrm{SE}}(f) = \sqrt{1-\mathcal{A}}\,\exp[-(f/f_{\mathrm{scr}})^{\beta}]$. Structurally similar to the adopted form; differs only in the parameterization of the scrambling scale.

Why the present family was adopted. The chosen constitutive laws are the most parsimonious family satisfying all five design criteria simultaneously with the minimum number of free parameters (four). The microtexture filter Ξ_δ and reflectivity $R(f)$ share the same exponent δ and timescale τ_c, enforcing cross-sector consistency with zero additional parameters. Alternative families either introduce additional free parameters (violating minimality), decouple the conservative and ringdown sectors (violating the shared-microstructure hypothesis), or impose sharper frequency structure than the data can currently resolve. The present family should be understood as the *minimal baseline closure*; future work with operator-informed constitutive laws (Sec. 7) may motivate richer functional forms.

E Sensitivity to taper-window choices

The taper windows of Sec. 3.4 are model-design choices rather than literature-derived observables. To test whether the paper's qualitative conclusions depend strongly on those choices, we varied the taper center $f_{\Lambda,\mathrm{off}}$ and width σ_Λ by $\pm 20\%$ around their default values and repeated the key diagnostics for all three parameter sets (conservative, moderate, strong).

Results. The taper-sensitivity study shows two different regimes. Variations in the taper width σ_Λ produce negligible changes ($< 2\%$) across all three parameter sets. Variations in the taper center $f_{\Lambda,\mathrm{off}}$ remain negligible in the conservative corner ($< 2\%$) and modest in the moderate case (up to $\sim 15\%$), but become substantial in the strong case (up to $\sim 60\%$). This behavior is expected because the strong case ($\epsilon = 10^{-4}$) occupies a broader active tidal band extending closer to the taper boundary.

Importantly, these shifts do not alter the paper's qualitative conclusion: the channel hierarchy—heating first, conservative response second, recycled postmerger signal last—remains intact under all tested perturbations. The heating-to-tidal SNR ratio in the conservative corner varies from ~ 28 (taper center shifted -20%) to ~ 12 (taper center shifted $+20\%$), compared to a baseline value of ~ 17. The heating channel remains dominant by more than an order of magnitude in all cases.

F Linear-response bridge to the membrane paradigm

This appendix does not derive Eqs. (17)–(19) from a fully specified microscopic brane model or from a complete solution of the Israel junction conditions. Its purpose is narrower: to show how the NHSB parameters can be interpreted as low-dimensional projections of a more general complex surface-response kernel in a membrane-paradigm language.

Consider an effective surface stress-energy tensor

$$S_{ab} = \sigma\, u_a u_b + p_s\, q_{ab} - 2\eta(\omega)\, \sigma_{ab} - \zeta(\omega)\, \theta\, q_{ab} + \tau^{\rm micro}_{ab}(\omega), \tag{35}$$

where u_a is the brane four-velocity, q_{ab} the induced spatial metric, $\eta(\omega)$ and $\zeta(\omega)$ effective frequency-dependent shear and bulk viscosities, and $\tau^{\rm micro}_{ab}(\omega)$ a compact phenomenological term encoding microstructure and scrambling. In linear response, the brane can be characterized by a complex susceptibility or admittance kernel $\chi(\omega)$, whose real part controls reactive (conservative) response and whose imaginary part controls dissipative absorption.

From this perspective, the NHSB quantities may be interpreted schematically as

$$\Lambda_{\rm eff}(f) \propto \frac{\Re\,\chi(f)}{|\ln\epsilon|}, \qquad \mathcal{A}(f) \propto \Im\,\chi(f), \tag{36}$$

with the compactness suppression $1/|\ln\epsilon|$ reflecting the standard ultracompact approach to the BH limit, and with ϵ entering the cavity time $\tau_{\rm c} \sim M|\ln\epsilon|$. Likewise, in a boundary-impedance picture the ringdown reflectivity may be written schematically as

$$R(f) \sim \left|\frac{Z(f) - Z_{\rm BH}(f)}{Z(f) + Z_{\rm BH}(f)}\right|, \tag{37}$$

where $Z(f)$ is an effective surface impedance and $Z_{\rm BH}(f)$ the BH limit. In this interpretation, lower absorptivity corresponds to a larger impedance mismatch and hence larger reflectivity.

The NHSB constitutive laws of Sec. 3 can therefore be viewed as one minimal realization of a broader response-kernel picture:

$$\Lambda_{\rm eff}(f) = \Lambda_\star\, |\ln\epsilon|^{-1}\, \Xi_\delta(f), \qquad R(f) = \sqrt{1 - \mathcal{A}\, \mathcal{S}_\delta(f)}, \tag{38}$$

where Ξ_δ and $\mathcal{S}_\delta$ are simple monotone rolloff functions governed by the same characteristic timescale $\tau_{\rm c}$. In this language, $\Lambda_\star$ sets the low-frequency conservative amplitude, $\mathcal{A}$ sets the low-frequency dissipative efficiency, δ controls the spectral rolloff of the response kernel, and ϵ controls the cavity depth and compactness suppression.

The mapping above is intentionally schematic. It is meant to show that the NHSB parameters are not floating free of a recognizable membrane-response language, while leaving the quantitative derivation of $\chi(\omega)$, $Z(\omega)$, and their relation to a specific surface stress-energy tensor for a dedicated follow-up study. The full Israel-junction calculation—starting from the viscous-fluid ansatz above and solving the linearized perturbed junction conditions for a thin shell at $r_{\rm m} = r_+(1+\epsilon)$—is identified as explicit future work and would provide the microscopic foundation that the present phenomenological framework is designed to accommodate.

F.1 Toy viscous-shell interpretation

A simple way to connect the NHSB parameters to familiar membrane variables is to imagine a viscous thin shell with a single dominant relaxation mode characterized by

$$\chi(\omega) \sim \frac{\chi_0}{1 + (i\omega\tau)^\delta}, \tag{39}$$

with χ_0 set by the low-frequency surface response and τ controlled by a combination of shell thickness, viscosity, and gravitational redshift (so that $\tau \sim \tau_{\rm c}$). In this toy picture, the NHSB parameters inherit their physical meaning directly:

$$\Lambda_\star \sim {\rm const}\times\Re\,\chi_0, \quad \mathcal{A} \sim {\rm const}\times\Im\,\chi_0, \quad \tau \sim \tau_{\rm c}. \tag{40}$$

The purpose of this toy mapping is not to claim that $\Lambda_\star$ and $\mathcal{A}$ are literally measurable viscosity coefficients: they are effective waveform-level projections. Rather, the mapping demonstrates that the NHSB parameter space has a natural home inside the response landscape of a dissipative near-horizon shell, and that the functional forms of the constitutive laws adopted in the main text are consistent with (though not uniquely determined by) such a physical realization. A full derivation from specified shear and bulk viscosities via the Israel junction conditions remains the scope of the companion analysis.

G Translation to heating-inclusive baselines

The present paper adopts the convention that the baseline waveform excludes BH tidal heating (Sec. 3.2.2). Some calibrated waveform families, such as the merger-capable approximant of Mukherjee *et al.* [15, 38], already include BH tidal heating in the baseline.

For such baselines, the NHSB absorptivity $\mathcal{A}$ should be reinterpreted as a fractional deviation around the BH heating level:

$$\Delta\Psi_{\mathcal{A}}(f) = (\mathcal{A} - 1)\, \Phi_{\text{heat}}^{\text{lib}}(f; \boldsymbol{\vartheta}), \tag{41}$$

so that $\mathcal{A} = 1$ produces identically zero heating correction and $\mathcal{A} < 1$ reduces the absorbed flux below the BH level. This is a reparameterization, not a physical change to the model. All results in the present paper use the heating-excluded convention; this appendix is included only as a translation guide for future implementations on heating-inclusive baselines.

H Toy Schwarzschild shell derivation of BRA scaling relations

This appendix provides a demonstrative derivation showing how the BRA/NHSB closure class can emerge from a simple dissipative shell in the Schwarzschild limit. The purpose is not to replace the full Israel-junction companion study but to convert the constitutive laws from purely interpretive ansätze into explicit low-dimensional reductions of a concrete toy model.

H.1 Cavity depth and timescale

Consider a Schwarzschild exterior with a thin shell at $r_{\mathrm{m}} = 2M(1+\epsilon)$ for $0 < \epsilon \ll 1$. Using the tortoise coordinate

$$r_* = r + 2M \ln\!\left(\frac{r}{2M} - 1\right), \tag{42}$$

the shell location satisfies

$$r_*(r_m) \approx 2M(1+\epsilon) + 2M \ln \epsilon. \tag{43}$$

The dominant contribution to the near-horizon depth is logarithmic,

$$\Delta r_* \sim 2M\, |\ln \epsilon|. \tag{44}$$

A round trip between the shell and the angular-momentum barrier region then gives the cavity timescale

$$\tau_c \sim 2\, \Delta r_* \sim 4M\, |\ln \epsilon|, \tag{45}$$

which motivates the normalization $\kappa_\tau = 4$ adopted in Eq. (4). This scaling is geometric: it follows from the ultracompact tortoise-coordinate structure and does not depend on the detailed constitutive physics of the shell.

H.2 Linear shell response law

Let $E_{2m}(\omega)$ denote the external quadrupolar tidal driving field and $Q_{2m}(\omega)$ the induced shell quadrupole response. We write the linearized shell constitutive law as

$$Q_{2m}(\omega) = \chi(\omega)\, E_{2m}(\omega), \tag{46}$$

with the minimal dissipative susceptibility

$$\chi(\omega) = \frac{\chi_0}{1 + (i\omega\tau)^{\delta}}, \tag{47}$$

where χ_0 is the low-frequency response amplitude, τ is an effective relaxation timescale ($\tau \sim \tau_c$), and $\delta = 1$ recovers the standard Debye response.

H.3 Conservative response and compactness suppression

For even-parity $\ell = 2$ perturbations in the Schwarzschild exterior, the far-zone solution contains a growing driving piece and a decaying response piece. The effective tidal response is defined by the ratio of these amplitudes. In the ultracompact limit, matching to a shell at $r_{\mathrm{m}} = 2M(1+\epsilon)$ produces the standard logarithmic suppression, so that the frequency-dependent conservative response takes the schematic form

$$\Lambda_{\text{eff}}(\omega) \sim \frac{\Re\, \chi(\omega)}{|\ln \epsilon|}. \tag{48}$$

This is the geometry-vs-microphysics split in one equation: $1/|\ln \epsilon|$ comes from ultracompact geometry, while $\Re\, \chi$ comes from the shell constitutive response.

H.4 Dissipation and the meaning of $\mathcal{A}$

In linear response, the dissipated power per cycle is controlled by the imaginary part of the response kernel:

$$P_{\rm diss}(\omega) \propto \omega\, \Im\, \chi(\omega)\, |E_{2m}(\omega)|^2. \tag{49}$$

The NHSB absorptivity parameter $\mathcal{A}$ should therefore be interpreted as the normalized low-frequency dissipative efficiency of the shell response, i.e. as the scalar parameter that captures the leading strength of the $\Im\,\chi$-controlled sector once the waveform basis is factored out:

$$\mathcal{A} \sim C_{\mathcal{A}}\, \Im\, \chi(\omega)\big|_{\rm inspiral\ band}. \tag{50}$$

H.5 Reflectivity from impedance mismatch

For postmerger perturbations, an effective boundary condition may be written schematically as

$$\partial_{r_*}\Psi\big|_{r_m} = -i\omega\, Z(\omega)\, \Psi\big|_{r_m}, \tag{51}$$

where $Z(\omega)$ is an effective shell impedance. The reflected amplitude then obeys the standard mismatch form

$$R(\omega) \sim \left|\frac{Z(\omega) - Z_{\rm BH}}{Z(\omega) + Z_{\rm BH}}\right|. \tag{52}$$

If $Z(\omega)$ is determined by the same underlying shell relaxation spectrum that controls $\chi(\omega)$, then the conservative, dissipative, and reflective sectors are linked at leading order by a common response structure.

H.6 Mapping to NHSB closures

The toy Schwarzschild shell model does not uniquely determine the functional forms used in the main text. Rather, it motivates the class

$$\Lambda_{\rm eff}(f) \approx \frac{C_\Lambda}{|\ln\epsilon|}\, \Re\, \chi(2\pi f), \tag{53}$$

$$\mathcal{A}_{\rm eff}(f) \approx C_{\mathcal{A}}\, \Im\, \chi(2\pi f), \tag{54}$$

$$R(f) \approx \mathcal{R}\big[Z\big(\chi(2\pi f)\big)\big]. \tag{55}$$

The NHSB constitutive laws are then understood as analytically tractable representatives of this class, chosen to satisfy the design criteria of Sec. 3.5 while remaining simple enough for waveform implementation:

$$\Lambda_{\rm eff}(f) = \Lambda_\star\, |\ln\epsilon|^{-1}\, [1 + (2\pi f \tau_{\rm c})^\delta]^{-1}, \tag{56}$$

$$R(f) = \sqrt{1-\mathcal{A}}\, \exp[-(2\pi f \tau_{\rm c})^\delta]. \tag{57}$$

H.7 From viscous stress-energy to the response kernel

The shell susceptibility $\chi(\omega)$ introduced above can be connected to the viscous fluid stress-energy tensor $S_{ab} = \sigma\, u_a u_b + p_s\, \gamma_{ab} - 2\eta(\omega)\, \sigma_{ab} - \zeta(\omega)\, \theta\, \gamma_{ab}$ at the perturbative level. For a linearized $\ell = 2$ tidal perturbation, the perturbed junction conditions couple the shell displacement δr to the perturbed stress δS_{ab}, which in turn depends on the shear viscosity $\eta(\omega)$. The effective quadrupolar response amplitude then takes the schematic form

$$\chi(\omega) \propto \frac{p_s + \text{geometric terms}}{p_s + \text{geometric terms} - 2i\omega\, \eta(\omega)}, \tag{58}$$

so that the reactive part $\Re\,\chi$ is controlled by the elastic shell response (surface pressure, equation of state) while the dissipative part $\Im\,\chi$ is controlled by the frequency-dependent viscosity $\eta(\omega)$. A Cole–Cole-type viscosity $\eta(\omega) \sim \eta_0/[1 + (i\omega\tau)^\delta]$ would then produce the Cole–Cole kernel adopted in the main text. The full algebraic matching is deferred to the companion study (Appendix I), but this structure makes clear that the response kernel is not a free choice: it is determined by the shell's material properties through the junction conditions.

H.8 Expected robustness under Kerr extension

The present derivation uses a Schwarzschild background. For astrophysically relevant spinning objects, the Kerr generalization would introduce frame-dragging corrections to both the tidal and scattering sectors. However, the two geometric ingredients of the NHSB model—the logarithmic compactness suppression $1/|\ln\epsilon|$ and the tortoise-depth cavity timescale $\tau_c \propto M|\ln\epsilon|$—are controlled by the near-horizon radial structure, which remains logarithmic in both Schwarzschild and Kerr. These geometric features are therefore expected to be robust under spin extension. The constitutive ingredients (the specific rolloff functions and the value of δ) may acquire spin-dependent corrections, but the three-channel BRA structure and the channel hierarchy do not depend on the background being non-rotating. A concrete slowly-rotating Kerr extension is scoped in the companion research plan (Appendix I).

This derivation does not prove uniqueness, but it demonstrates that the NHSB closure class has a natural home inside the response landscape of a dissipative near-horizon shell.

I Research plan: From brane viscosity to constitutive laws

This appendix outlines the companion study that would convert the demonstrative Schwarzschild-shell derivation of Appendix H into a quantitative first-principles calculation. We include it here so that the derivation gap between "demonstrative" (this paper) and "derived" (the companion) is visible as a concrete, scoped research program rather than a vague promise.

I.1 Thesis

The NHSB constitutive laws (Eqs. (17)–(19)) were designed as phenomenological closures satisfying five design criteria. The companion study asks: does a concrete thin-shell model—a viscous brane at $r_m = r_+(1+\epsilon)$, matched to a Schwarzschild exterior and a regular interior via the Israel junction conditions [31]—naturally produce constitutive laws of the NHSB form? If so, what is the mapping between microphysical variables (surface density σ, viscosity $\eta(\omega)$, relaxation timescale) and effective waveform parameters $(\Lambda_\star, \mathcal{A}, \delta)$? If not, what form does emerge, and what does it imply for the NHSB phenomenology?

The answer, whatever it is, advances the program: confirmation validates the phenomenological structure; deviation identifies where the closures need refinement; failure modes reveal which physical assumptions are load-bearing.

I.2 Physical setup

The background is a static thin shell carrying a viscous fluid:

$$S_{ab} = \sigma\, u_a u_b + p_s\, \gamma_{ab} - 2\eta\, \sigma_{ab} - \zeta\, \theta\, \gamma_{ab}, \quad (59)$$

where σ and p_s are the surface energy density and pressure, η and ζ are shear and bulk viscosities, σ_{ab} is the shear tensor, and θ is the expansion scalar. This is the standard Israel–Lanczos–Sen framework extended to viscous shells [49]. The exterior is Schwarzschild; the interior is Minkowski or de Sitter.

I.3 Three derivation targets

Target 1: Conservative tidal response. Apply a static $\ell = 2$ tidal field from infinity. Solve the linearized Israel conditions $[\delta K_{ab}] - [\delta K]\, h_{ab} - [K]\, \delta h_{ab} = -8\pi\, \delta S_{ab}$ for the ratio of the response (decaying) to the applied (growing) tidal harmonic. Test whether $\Lambda_{\rm eff}(\omega) = [\Lambda_\star/|\ln\epsilon|]\, \Xi(\omega)$, where Ξ is a monotonic rolloff function. Sub-questions: does the $1/|\ln\epsilon|$ prefactor emerge from the background geometry? Does the rolloff have the microtexture-filter form $[1+(\omega\tau)^\delta]^{-1}$? What determines δ?

Target 2: Dissipation and reflectivity. Switch from quasistatic tidal perturbations to oscillatory wave scattering. The viscous shell provides a partially absorbing boundary condition; the reflectivity $R(\omega)$ is the ratio of outgoing to incoming wave amplitudes at the brane. Test whether $R(\omega) = \sqrt{1-\mathcal{A}}\, \exp[-(\omega\tau_c)^\delta]$, and whether the same δ that controls the tidal sector also controls the reflectivity.

Target 3: The single-δ question. Report whether the tidal and scattering sectors share an exponent. If they do, the single-δ NHSB model is justified as the leading-order closure. If they do not, the $(\delta_\Lambda, \delta_R)$ extension is not just possible but necessary. If it depends on the viscosity model, the NHSB parameterization is agnostic to the right question.

I.4 Possible outcomes

Three scenarios are possible and each advances the program:

Outcome A (NHSB confirmed): The derived $\Lambda_{\rm eff}$ and R match the NHSB forms at leading order. Report the mapping $\Lambda_\star = f(\sigma, p_s, \eta, \epsilon)$, $\mathcal{A} = g(\eta, \zeta, \epsilon)$, $\delta = h(\text{viscosity model})$.

Outcome B (Partial match): One sector matches but the other requires a different exponent or functional form. Report what matches, what differs, and what it implies for the single-δ assumption and the BRA channel hierarchy.

Outcome C (Different structure): The derived forms are qualitatively different (e.g., resonance structure rather than a simple filter). Report the actual forms and discuss which NHSB predictions survive and which need revision.

I.5 Scope boundary

The companion study is restricted to Schwarzschild background, linear perturbations, and constant or power-law viscosity models. Kerr background, nonlinear merger-regime effects, and fractional/anomalous viscosity models (connecting $\delta \neq 1$ to anomalous diffusion on the brane [39]) are identified as subsequent extensions. The calculation is analytic/semi-analytic throughout; no numerical relativity is required.

References

[1] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), Observation of Gravitational Waves from a Binary Black Hole Merger, Phys. Rev. Lett. **116**, 061102 (2016).

[2] S. W. Hawking, Breakdown of predictability in gravitational collapse, Phys. Rev. D **14**, 2460 (1976).

[3] P. O. Mazur and E. Mottola, Gravitational vacuum condensate stars, Proc. Nat. Acad. Sci. **101**, 9545 (2004).

[4] C. Cattoen, T. Faber, and M. Visser, Gravastars must have anisotropic pressures, Class. Quantum Grav. **22**, 4189 (2005).

[5] I. Bena and D. R. Mayerson, Black holes, black rings, and their microstates, Lect. Notes Phys. **1005**, 1 (2022).

[6] A. Almheiri, D. Marolf, J. Polchinski, and J. Sully, Black holes: Complementarity vs. firewalls, J. High Energy Phys. **02**, 062 (2013).

[7] L. Susskind, L. Thorlacius, and J. Uglum, The stretched horizon and black hole complementarity, Phys. Rev. D **48**, 3743 (1993).

[8] V. Cardoso, E. Franzoni, and P. Pani, Is the gravitational-wave ringdown a probe of the event horizon?, Phys. Rev. Lett. **116**, 171101 (2016).

[9] V. Cardoso and P. Pani, Tests for the existence of black holes through gravitational wave echoes, Nat. Astron. **1**, 586 (2017).

[10] T. Damour and S. N. Solodukhin, Wormholes as black hole foils, Phys. Rev. D **76**, 024016 (2007).

[11] T. Binnington and E. Poisson, Relativistic theory of tidal Love numbers, Phys. Rev. D **80**, 084018 (2009).

[12] A. Le Tiec, M. Casals, and E. Franzini, Tidal Love Numbers of Kerr Black Holes, Phys. Rev. D **103**, 084021 (2021).

[13] S. Datta, R. Brito, S. Bose, P. Pani, and S. A. Hughes, Tidal heating as a discriminator for horizons in extreme mass ratio inspirals, Phys. Rev. D **101**, 044004 (2020).

[14] H. S. Chia, Tidal deformation and dissipation of rotating black holes, Phys. Rev. D **109**, 024030 (2024).

[15] S. Mukherjee, S. Datta, S. Bose, and K. S. Phukon, Binary black holes in the heat of merger, arXiv:2506.22363 [gr-qc] (2025).

[16] M. Silvestrini, E. Maggio, S. Chakraborty, and P. Pani, One membrane to Love them all: Tidal deformations of compact objects from the membrane paradigm, Phys. Rev. D **112**, 124021 (2025); arXiv:2506.16516.

[17] Z. Mark, A. Zimmerman, S. M. Du, and Y. Chen, A recipe for echoes from exotic compact objects, Phys. Rev. D **96**, 084002 (2017).

[18] E. Maggio, A. Testa, V. Cardoso, and P. Pani, Analytical model for gravitational-wave echoes from spinning remnants, Phys. Rev. D **100**, 064056 (2019).

[19] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), GW170817: Measurements of Neutron Star Radii and Equation of State, Phys. Rev. Lett. **121**, 161101 (2018).

[20] H. S. Chia, Z. Zhou, and M. M. Ivanov, Tidal heating constraints for black holes and exotic compact objects from the LIGO-Virgo-KAGRA data, Phys. Rev. D **111**, 063002 (2025); arXiv:2404.14641 [gr-qc].

[21] J. Abedi, H. Dykaar, and N. Afshordi, Echoes from the abyss: Tentative evidence for Planck-scale structure at black hole horizons, Phys. Rev. D **96**, 082004 (2017).

[22] J. Westerweck *et al.*, Low significance of evidence for black hole echoes in gravitational wave data, Phys. Rev. D **97**, 124037 (2018).

[23] R. Abbott *et al.* (LIGO Scientific Collaboration, Virgo Collaboration, and KAGRA Collaboration), Search for gravitational-wave signals associated with ultracompact objects, Phys. Rev. D **104**, 082002 (2021).

[24] E. Maggio, V. Cardoso, S. R. Dolan, and P. Pani, Ergoregion instability of exotic compact objects: ultracompact stars and quantum corrections at the horizon scale, Phys. Rev. D **99**, 064007 (2019).

[25] E. Maggio, P. Pani, and V. Ferrari, Exotic compact objects and how to quench their ergoregion instability, Phys. Rev. D **96**, 104047 (2017).

[26] M. V. S. Saketh and E. Maggio, Quasinormal modes of slowly-spinning horizonless compact objects, Phys. Rev. D **110**, 084030 (2024); arXiv:2406.10070.

[27] K. W. Tsang, A. Ghosh, A. Samajdar, K. Chatziioannou, S. Mastrogiovanni, M. Agathos, and C. Van Den Broeck, A morphology-independent search for gravitational wave echoes in data from the first and second observing runs of Advanced LIGO and Advanced Virgo, Phys. Rev. D **101**, 064012 (2020).

[28] K. S. Thorne, R. H. Price, and D. A. Macdonald, *Black Holes: The Membrane Paradigm* (Yale University Press, New Haven, 1986).

[29] G. E. Volovik, What can the quantum liquid say on the brane black hole, the entropy of an extremal black hole, and the vacuum energy?, Found. Phys. **33**, 349 (2003).

[30] M. Pavšič, On a unified theory of generalized branes coupled to gauge fields, including the gravitational and Kalb-Ramond fields, Found. Phys. **37**, 1197 (2007).

[31] W. Israel, Singular hypersurfaces and thin shells in general relativity, Nuovo Cim. B **44**, 1 (1966); Erratum: **48**, 463 (1967).

[32] Y. Sekino and L. Susskind, Fast scramblers, J. High Energy Phys. **10**, 065 (2008).

[33] J. Maldacena, S. H. Shenker, and D. Stanford, A bound on chaos, J. High Energy Phys. **08**, 106 (2016).

[34] P. Hayden and J. Preskill, Black holes as mirrors: Quantum information in random subsystems, J. High Energy Phys. **09**, 120 (2007).

[35] V. Cardoso, E. Franzoni, A. Maselli, P. Pani, and G. Raposo, Testing strong-field gravity with tidal Love numbers, Phys. Rev. D **95**, 084014 (2017); Erratum: **95**, 089901 (2017).

[36] P. Pani, I-Love-Q relations for gravastars and the approach to the black-hole limit, Phys. Rev. D **92**, 124030 (2015).

[37] S. Ghosh and M. Hannam, On the identification of exotic compact binaries with gravitational waves: a phenomenological approach, Phys. Rev. D **112**, 104017 (2025); arXiv:2505.16380.

[38] S. Mukherjee, S. Datta, S. Bose, and S. K. Phukon, IMRPhenomD_Horizon: A frequency-domain gravitational waveform model incorporating horizon fluxes, Phys. Rev. D **110**, 124027 (2024).

[39] R. Metzler and J. Klafter, The random walk's guide to anomalous diffusion: a fractional dynamics approach, Phys. Rep. **339**, 1 (2000).

[40] R. Abbott *et al.* (LIGO–Virgo–KAGRA Collaboration), GWTC-4.0: Compact binary coalescences observed by LIGO and Virgo during the first part of the fourth observing run, arXiv:2502.11885 [gr-qc] (2025).

[41] R. Abbott *et al.* (LIGO–Virgo–KAGRA Collaboration), Open data from LIGO, Virgo, and KAGRA through the first part of the fourth observing run, arXiv:2508.08423 [astro-ph.IM] (2025).

[42] S. Mukherjee, S. Datta, S. Tiwari, K. S. Phukon, and S. Bose, Toward establishing the presence or absence of horizons in coalescing binaries of compact objects by using their gravitational-wave signals, Phys. Rev. D **106**, 104032 (2022); arXiv:2202.08661 [gr-qc].

[43] S. Datta, R. Brito, S. Bernuzzi, and A. Nagar, Tidal heating as a discriminator for horizons in extreme-mass-ratio inspirals, Phys. Rev. D **110**, 024048 (2024).

[44] N. Uchikata *et al.*, Searching for gravitational-wave echoes in the third observing run of Advanced LIGO–Virgo, Phys. Rev. D **108**, 084016 (2023); arXiv:2309.01894 [gr-qc].

[45] M. Evans *et al.*, A horizon study for Cosmic Explorer: science, observatories, and community, arXiv:2109.09882 [astro-ph.IM] (2021).

[46] M. Isi, M. Giesler, W. M. Farr, M. A. Scheel, and S. A. Teukolsky, Testing the no-hair theorem with GW150914, Phys. Rev. Lett. **123**, 111102 (2019).

[47] K. S. Cole and R. H. Cole, Dispersion and absorption in dielectrics. I. Alternating current characteristics, J. Chem. Phys. **9**, 341 (1941).

[48] P. Kovtun, D. T. Son, and A. O. Starinets, Viscosity in strongly interacting quantum field theories from black hole physics, Phys. Rev. Lett. **94**, 111601 (2005).

[49] E. Poisson, *A Relativist's Toolkit: The Mathematics of Black-Hole Mechanics* (Cambridge University Press, 2004).