Article

Not peer-reviewed version

A Layered Information-Geometric Framework for Stable Estimation Under Missing and Heteroscedastic Observations

Chien Chih Chen *

Posted Date: 22 April 2026

doi: 10.20944/preprints202604.1573.v1

Keywords: stable estimation; missing data; heteroscedastic noise; PT-even gatekeeping; diagonal-Fisher refinement; coarse-to-fine seeding; safeguarded optimization; curvature-aware estimation

 doi:10.20944/preprints202604.1573.v1

Article

A Layered Information-Geometric Framework for Stable Estimation Under Missing and Heteroscedastic Observations

Chien Chih Chen

Chunghwa Telecom Laboratories, Taiwan; rocky@cht.com.tw

Abstract

We investigate the problem of stable signal estimation under irregular observation conditions characterized by missing samples, heteroscedastic noise, and reportability constraints. In such environments, the primary engineering challenge arises from the severe curvature distortion of the objective landscape, which renders traditional Euclidean gradient methods numerically unstable. To address this, we propose a layered information-geometric framework for stable estimation on a distorted statistical manifold. The architecture consists of three integrated components: (i) a deterministic basin-safe front-end that resolves global navigation on the non-convex information landscape; (ii) a diagonal-Fisher local refiner that performs local metric normalization to correct for geometric scaling under irregular weighting; and (iii) a PT-even gatekeeper that acts as an engineering feasibility constraint by restricting the optimization trajectory to the reportable subspace. As a concrete structured path within this layered estimator, the QFRT branch—a Quaternionic Fourier–Ramanujan representation used here as an arithmetic-periodicity anchor under shared whitening and weighting—serves as a structured source of robustness under gappy observations without altering the bounded scope of the present technical claim. Across a locked stress matrix, the resulting hybrid estimator exhibits a two-layer gain structure. First, a general refinement gain in RMSE_ω is achieved via curvature-aware updates in the classical-tangent regime. Second, a specific PT-sensitive gain emerges when nuisance-coupled sectors become observable, effectively suppressing non-reportable "ghost-mode" leakage ($|z| \approx 0$) where unprojected baselines suffer from substantial parameter drift. Mechanism diagnostics support a seed-path shielding interpretation: the classical front-end resolves the ω-dominant basin selection problem, shielding the downstream PT-aware refinement from unfavorable seed geometry. The resulting contribution is a technical methods framework for auditable stable estimation under missing and heteroscedastic observations.

Keywords: stable estimation; missing data; heteroscedastic noise; PT-even gatekeeping; diagonal-Fisher refinement; coarse-to-fine seeding; safeguarded optimization; curvature-aware estimation

---

1. Introduction

Many estimation problems of practical interest operate under hostile and irregular observation conditions, including missing samples, severe fading, and heteroscedastic noise. Standard estimation pipelines, typically analyzed under regular sampling and stationary noise, often falter in such regimes. From an information-geometric perspective, these irregularities do not merely increase variance; they induce *severe curvature distortion* of the underlying objective landscape. Such distortion renders traditional Euclidean gradient-based methods numerically unstable, leading to slow convergence, divergence, or entrapment in nuisance subspaces outside the declared reportable channel.

This paper proposes a layered information-geometric framework designed to restore stability and reportability to the estimation process. We treat the estimation problem not as a simple parameter search, but as a constrained optimization problem on a statistical manifold. By shifting the focus

 doi:10.20944/preprints202604.1573.v1

from raw numerical residuals to the intrinsic geometry of the information landscape, we provide a principled way to navigate distorted curvatures and enforce a projected-model reportability constraint under irregular observations.

For the reader, the architecture can be summarized in one line before the formal development becomes notation-heavy: a coarse-to-fine classical front-end first provides a basin-safe seed; a diagonal-Fisher local refinement then rescales the update according to local reliability-weighted sensitivity; and a PT-even reportability gate keeps the returned solution inside the declared reportable channel. Section 4 turns this reader map into the explicit layered estimator.

1.1. The Necessity of a Layered Architecture

To address the geometric challenges posed by irregular observations, we argue that a monolithic update rule is insufficient. Instead, a stable estimator requires a *three-layer architecture* where each component serves a distinct engineering function:

1. **Global Navigation (Basin-Safe Seeding):** In a non-convex information landscape, local refinement is only meaningful if it begins within a favorable basin. We therefore use a deterministic coarse-to-fine classical front-end to select a stable starting point before the local update is allowed to act.
2. **Local Metric Normalization (Diagonal-Fisher Refinement):** To counteract curvature distortion caused by heteroscedastic weighting, we employ a diagonal-Fisher preconditioner. In signal-processing terms, this layer rescales each parameter step by local reliability-weighted sensitivity rather than by raw Euclidean magnitude. This is the local refinement layer used throughout the paper.
3. **Subspace Projection (PT-Even Gatekeeping):** To keep the returned solution within the declared reportable channel, we introduce a PT-even reportability gate. PT-structured language is used here strictly as an *engineering feasibility constraint*. Within the projected model, this gate suppresses nuisance-mode leakage and prevents the optimizer from reporting estimates outside the declared reportable subspace.

1.2. Main Idea and Scope Lock

The central thesis of this work is that stable estimation under irregular observations is achieved by aligning the optimization path with the reportable subspace of the signal manifold. Our approach does not seek to universally replace classical FFT-family baselines; rather, it provides a second-layer gain. A strong classical front-end identifies a stable basin, while the PT-aware local refinement improves precision and suppresses leakage within that basin.

Empirical evidence validates a *two-layer gain structure*: a general refinement gain already present in classical-tangent regimes, and a specific PT-sensitive gain that emerges once nuisance-coupled sectors become observable. To ensure transparency, the framework is governed by a safeguarded acceptance mechanism that explicitly rejects trajectories when the numerical geometry becomes unreliable. The scope of this paper is thus locked to *stable signal estimation under missing and heteroscedastic observations*, with mechanism diagnostics used to audit where the gain resides and where its boundary lies.

In the present implementation, this layered architecture is instantiated through a QFRT-style two-path structure. Path A provides the deterministic basin-safe classical seed, while Path B is a QFRT-structured branch in which a Quaternionic Fourier–Ramanujan representation is used as an arithmetic-periodicity anchor under the same whitening and weighting discipline. The name Quaternionic Fourier–Ramanujan Transform is therefore justified only at the level of this structured branch: Path B uses a quaternionic Fourier representation together with a Ramanujan-type arithmetic-periodicity kernel to supply an additional periodicity cue inside the shared weighted model, and is not presented as a stand-alone estimator, as the framework itself, or as the full explanation of the observed gain.

Unless otherwise noted, result-level references in this paper point to the QFRT public technical anchor (`qfrt-evidence-public`), treated here as a curated reporting and traceability surface. Canoni-

 doi:10.20944/preprints202604.1573.v1

cal documentation, manifest, and manuscript-facing reporting files provide the audit trail for tables and figures without implying a public execution environment.

1.3. Contributions

The main contributions of this paper are as follows:

- We formulate a **layered information-geometric framework** for irregular observations, integrating basin-safe seeding with geometry-aware local refinement.
- We introduce the concept of **Local Metric Normalization** via diagonal-Fisher scaling to stabilize parameter scaling under severe curvature distortion.
- We define a **PT-even gatekeeping mechanism** as an operational reportability constraint, aligning optimization trajectories with the projected reportable channel and suppressing ghost-mode leakage.
- We empirically demonstrate a **two-layer gain structure** and localize the onset of PT-sensitive gain, providing mechanism-level evidence that supports a basin-safe seeding and subspace-projection interpretation.

1.4. Paper Organization

Section 2 reviews related work in information geometry, natural gradients, and robust synchronization. Section 3 defines the canonical forward model and the PT-even/odd residual decomposition. Section 4 details the layered estimator framework, including the diagonal-Fisher update and the safeguarded governance. Section 5 specifies the definition-locked experimental protocol. Section 6 presents primary results and mechanism-oriented findings. Section 7 discusses interpretation and limitations, and Section 8 concludes. Auxiliary proofs, including the *variational safety* of the PT-projector, are provided in the Appendices.

2. Related Work

The challenge of stable estimation under irregular observations sits at the intersection of classical synchronization, robust statistics, and information geometry. To position the proposed layered framework, we review the literature across four key dimensions: classical baselines, weighted inference, geometric optimization, and safeguarded numerical workflows.

2.1. Classical Synchronization and Irregularity

FFT-based synchronization remains the industry standard due to its computational efficiency and well-understood behavior under regular sampling [2,3]. However, in 6G NTN scenarios [1], the introduction of large-scale missingness and non-stationary noise alters the spectral concentration of these estimators. While Fourier-domain methods are robust to AWGN, they often exhibit increased sensitivity to implementation-specific masking and normalization choices when data are gappy [5, 6]. We use these classical procedures as *transparency anchors*: deterministic baselines that resolve global basin selection but lack the geometric awareness required for high-precision refinement under distortion.

2.2. Information Geometry and Natural Gradients

The foundational work of Amari [8,9] established that the parameter space of a statistical model is not a Euclidean flatland but a Riemannian manifold equipped with the Fisher information metric. In the presence of heteroscedastic noise and irregular weighting, the "distance" between iterates should be measured by the change in the underlying distribution rather than raw Euclidean coordinates. Natural gradient methods address this by preconditioning updates with the inverse Fisher Information Matrix (FIM), effectively performing a natural-gradient step that is invariant to reparameterization.

While theoretically optimal for resolving the instability caused by non-uniform observation quality, the full FIM is often computationally prohibitive. This has motivated the development of various approximations, such as Kronecker-factored curvature or structured surrogates [14–16]. Our

 doi:10.20944/preprints202604.1573.v1

work follows this lineage by treating the Diagonal-Fisher refinement as a **lightweight geometric approximation**. Unlike full information-geometric optimizers, our approach focuses on *local metric normalization* to suppress the specific curvature distortions induced by heteroscedastic weighting in the 6G NTN context.

2.3. Weighted Estimation and Curvature Distortion

Inference with incomplete observations has a long history in statistics, typically addressed via weighted least squares or Expectation-Maximization (EM) variants [10,11]. While these formulations are principled, they primarily focus on the *objective definition* rather than the *optimization trajectory*. Recent studies in robust statistics suggest that as missingness increases, the local conditioning of the objective becomes severely distorted, leading to imbalanced gradients across different parameter axes [12,13]. This paper bridges this gap by explicitly coupling the weighted objective with a geometry-aware update rule to ensure that the optimization path remains stable despite these distortions.

2.4. Safeguarded Optimization and Reportability Constraints

Modern numerical optimization emphasizes the separation of the *search direction* from the *governance mechanism* [17]. In communication systems, this is often manifested as projection-based methods that enforce constraints like power limits or spectral masks. The concept of PT-symmetry, originally from non-Hermitian physics [7], offers a useful language for defining such constraints. We re-interpret PT-gatekeeping as an operational reportability constraint within the projected model. By restricting iterates to a reportable subspace, the framework keeps the optimization trajectory aligned with the structural channel declared reportable by the signal model.

2.5. Positioning of the Present Work

Existing literature provides either robust global search (FFT) or rigorous local geometry (Natural Gradient). However, a unified approach that integrates basin-safe navigation with lightweight geometric approximation under reportability constraints is missing. The present work fills this void by developing a **layered estimator** that utilizes information-geometric principles to provide stable, auditable, and reportability-constrained refinement under the irregular observation conditions studied here.

Conceptually, the reportability declaration used here is consistent with PTQ viewpoints in which physical or reportable content is fixed only after projection onto an admissible observable sector, with PT-even selection providing the scalar observable branch retained for reporting [19,20]. In the present paper, however, that lineage is used only as a bounded engineering interpretation for projector-defined admissibility inside the projected model; it is not presented as a derivation of the full estimator or of the observed engineering gains.

3. Problem Setup & Canonical Model

This section establishes the analytical foundation for estimation on the signal manifold. We define a *projected complex model* in which the signal space is partitioned into reportable and nuisance sectors. This formulation allows us to treat synchronization not merely as a numerical fitting problem, but as a constrained search within a declared reportable channel under irregular observations.

3.1. Observation Model in 6G NTN Environments

Consider a complex baseband block $y \in \mathbb{C}^N$ observed over a Non-Terrestrial Network (NTN) link, where $s \in \mathbb{C}^N$ is the known reference template. The canonical parameter vector is defined as:

$$\theta \triangleq (\tau, \omega, z), \tag{1}$$

 doi:10.20944/preprints202604.1573.v1

where $\tau \in \mathbb{R}$ represents the circular timing offset (samples), $\omega \in \mathbb{R}$ is the carrier frequency offset (radians/sample), and $z \in \mathbb{C}$ is a **nuisance coefficient** coupled to the non-reportable sector of the signal model. The observation model is expressed as:

$$y[n] = H(\theta)[n] + \varepsilon[n], \qquad n = 0, 1, \ldots, N-1, \tag{2}$$

where $\varepsilon[n]$ denotes additive disturbance. In 6G NTN scenarios, the observations are typically *irregular*, characterized by missing samples and non-stationary noise power.

3.2. Irregularity Weights and Curvature Distortion

To unify the treatment of missingness and heteroscedasticity, we introduce a set of nonnegative weights $w[n]$ that represent the *local metric reliability* of the information manifold:

$$w[n] \triangleq \frac{m[n]}{\sigma^2[n]}, \tag{3}$$

where $m[n] \in \{0, 1\}$ is a binary mask (with 0 indicating a gappy/missing sample) and $\sigma^2[n]$ is the time-varying noise variance. From a geometric perspective, the distribution of $w[n]$ induces a **curvature distortion** on the objective landscape, as samples with higher reliability exert greater pull on the optimization trajectory.

3.3. PT Action and Reportable/Nuisance Decomposition

The core of our framework relies on an operational discrete complex involution, termed the PT action. For any signal $x \in \mathbb{C}^N$, the action is defined as:

$$(\mathrm{PT}\, x)[n] \triangleq \overline{x[N-1-n]}, \qquad n = 0, 1, \ldots, N-1. \tag{4}$$

Within the projected model used here, this involution defines a **PT-even structural criterion** for separating the reportable branch from nuisance directions. We decompose the template s into a reportable PT-even branch and a nuisance PT-odd branch:

$$s^+[n] \triangleq \frac{1}{2}\big(s[n] + \mathrm{PT}(s)[n]\big), \qquad s^-[n] \triangleq \frac{1}{2}\big(s[n] - \mathrm{PT}(s)[n]\big). \tag{5}$$

The composite template, $s_z = s^+ + z\, s^-$, allows the model to capture **nuisance modes** (ghost-modes) such as I/Q imbalance or non-reciprocal hardware artifacts through the scalar weight z. By enforcing $z \approx 0$ in reportable scenarios, we protect the structural integrity of the synchronized state.

3.4. Canonical Operators and Forward Model

The transformation of the template on the manifold is governed by the circular shift and CFO operators. The shift operator $\mathrm{Shift}_\tau(\cdot)$ is implemented via a DFT-domain phase ramp to ensure differentiability [4]:

$$\mathrm{Shift}_\tau(x) \triangleq \mathcal{F}^{-1}\big(\mathcal{F}(x)[k] \cdot \exp(-j\frac{2\pi k}{N}\tau)\big). \tag{6}$$

The frequency rotation is given by $\mathrm{CFO}_\omega(x)[n] \triangleq \exp(j\omega n)x[n]$. Combining these, the canonical forward model $H(\theta)$ is:

$$H(\theta) \triangleq \mathrm{CFO}_\omega\big(\mathrm{Shift}_\tau(s^+ + z\, s^-)\big). \tag{7}$$

Under the shared weighting model of Section 3.2, we use this forward model in a two-path instantiation: Path A supplies the classical basin-selection route, while Path B is a QFRT-structured path, where "QFRT" denotes a Quaternionic Fourier–Ramanujan representation used here as a structured arithmetic-periodicity anchor under the same whitening and weighting discipline. In the present manuscript, the role of Path B is functional and bounded: it contributes structured periodicity information that is

 doi:10.20944/preprints202604.1573.v1

particularly useful when the observations are gappy or irregular, but it is not taken to explain the full layered estimator by itself.

More technically, Path B can be specified as a weighted Quaternionic Fourier–Ramanujan branch that evaluates the PT-even template against arithmetic-periodicity atoms before the PT-aware local refinement is applied. Operationally, this branch acts as a structured periodicity cue: after the shared shift/CFO parameterization is applied, the weighted correlation against Ramanujan atoms highlights periodic components that remain informative even when masking and heteroscedastic weights weaken plain spectral concentration. In this manuscript, the Ramanujan component is used only for that bounded role inside the shared weighted model; basin selection, local metric normalization, acceptance control, and reportability remain properties of the full layered architecture.

$$\mathcal{A}_{\mathrm{QFR}}(\theta; q) \triangleq \sum_{n=0}^{N-1} w[n]\, c_q(n)\, \overline{\left(\mathrm{CFO}_{\omega}(\mathrm{Shift}_{\tau}(s^{+}))[n]\right)}, \qquad q \in \mathcal{Q}, \tag{8}$$

where $c_q(n)$ denotes a Ramanujan-sum atom at order q. For a signal-processing reader, $\mathcal{A}_{\mathrm{QFR}}(\theta; q)$ should be read as a weighted periodicity score computed on the same reportable branch and under the same reliability weights used by the rest of the estimator. Because the anchor is accumulated through those shared weights, it can retain structured periodicity information even when observations are gappy or irregular. Its role is therefore supportive rather than exhaustive: it helps define the QFRT-structured branch, but it does not by itself account for the layered gain reported later in the paper.

3.5. Weighted Objectives and the Reportability Declaration

The residual is defined as $r(\theta) \triangleq y - H(\theta)$, which we further project onto the PT-even subspace to obtain the **reportable residual**: $r^{+}(\theta) \triangleq \Pi_{+} r(\theta)$. We then define two canonical objectives:

1. **Raw Objective (S_{raw}):** Evaluates the total weighted residual power across all modes, used primarily for global basin selection.
2. **Reportable Objective (S_{base}):** Evaluates power strictly within the PT-even projected subspace:

$$S_{\mathrm{base}}(\theta) \triangleq \sum_{n=0}^{N-1} w[n]\, \left| r^{+}(\theta)[n] \right|^{2}. \tag{9}$$

The use of S_{base} as the primary optimization target constitutes a **reportability declaration**: the estimator is explicitly constrained to find solutions that align with the declared reportable signal channel, while treating z-coupled drift as non-reportable leakage.

4. Layered Estimator: PT-Even Gatekeeping, Local Metric Normalization, and Basin-Safe Seeding

To restore estimation stability under the severe curvature distortions induced by irregular observations, we propose a layered information-geometric framework. Unlike monolithic optimizers that treat parameter updates as Euclidean displacements, our architecture recognizes that missingness and heteroscedasticity distort the underlying information manifold. As summarized in Figure 1, the operational pipeline combines a deterministic coarse-to-fine classical front-end, a PT-even reportability projection, a diagonal-Fisher local refiner, and a safeguarded acceptance layer into one bounded estimation architecture.

For mathematical exposition, we introduce the layers in the order needed to define the update rule: first the reportable residual, then the local geometry-normalized step, and finally the classical seeding layer that places that step in a favorable basin.

The framework is structured into three functional layers: (i) a PT-even gatekeeper that defines the projected reportable subspace; (ii) a diagonal-Fisher local refiner that performs local metric normalization; and (iii) a deterministic basin-safe front-end that ensures stable navigation through the non-convex objective landscape.

 doi:10.20944/preprints202604.1573.v1

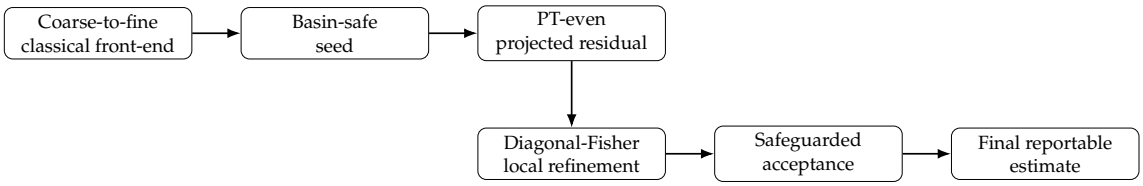

Figure 1. Layered architecture of the proposed estimator. A deterministic coarse-to-fine classical front-end first provides a basin-safe seed; the PT-even projected residual defines the reportable channel; diagonal-Fisher local refinement performs geometry-normalized updating; and safeguarded acceptance governs the final returned estimate.

4.1. PT-Even Gatekeeping: Enforcing Reportability

The first layer of our framework is the **PT-even gatekeeper**, which restricts the optimization trajectory to the projected reportable subspace of the signal manifold. Under the reportability declaration in Section 3, returned solutions are evaluated in the PT-even channel. Any parameter drift into the PT-odd branch ($z \neq 0$) is therefore operationally interpreted as the emergence of non-reportable "ghost-modes" caused by numerical artifacts or hardware non-idealities.

To enforce this constraint, the gatekeeper aligns the gradient path strictly with the reportable residual $r^{+}(\theta) \triangleq \Pi_{+} r(\theta)$. For the PT-aware local refiner (P0), the effective residual is defined as:

$$r_{\text{eff}}(\theta) \triangleq \Pi_{+} r(\theta). \tag{10}$$

By projecting the residual prior to gradient computation, the gatekeeper acts as an *engineering feasibility constraint*, preventing the optimizer from claiming convergence in non-reportable regions of the parameter space.

4.2. Diagonal-Fisher Refinement: Local Metric Normalization

The second layer addresses the geometric scaling imbalance caused by irregular observation weights $w[n]$. Under heteroscedastic noise, the Hessian matrix of the cost function becomes ill-conditioned, leading to unstable step sizes and divergent trajectories. We counteract this via **Local Metric Normalization** using a diagonal-Fisher preconditioner.

Let $J_{i,\text{eff}}(\theta) \triangleq \Pi_{+} \partial_{\theta_i} r(\theta)$ be the effective Jacobian projected onto the reportable subspace. The weighted gradient component $g_i(\theta)$ is given by:

$$g_i(\theta) \triangleq 2\,\Re\left\{ \sum_{n=0}^{N-1} w[n]\, \overline{r_{\text{eff}}(\theta)[n]}\, J_{i,\text{eff}}(\theta)[n] \right\}. \tag{11}$$

To correct the manifold curvature, we define a local metric surrogate $D_i(\theta)$ based on the diagonal elements of the Fisher Information Matrix (FIM):

$$D_i(\theta) \triangleq \sum_{n=0}^{N-1} w[n]\, \left| J_{i,\text{eff}}(\theta)[n] \right|^2 + \lambda_{\text{ridge}}. \tag{12}$$

Preconditioning the update with the inverse of this diagonal surrogate corresponds to taking a diagonal Fisher-scaled step on the statistical manifold [9]. We refer to this lightweight geometry-normalized step as the B-lite update:

$$\Delta\theta_i = -\eta\, \frac{g_i(\theta)}{\widetilde{D}_i(\theta)}, \tag{13}$$

where $\widetilde{D}_i(\theta) = \max(D_i(\theta), D_{\min})$. This step ensures that parameter updates are normalized by the local information density, stabilizing the Hessian conditioning even under sparse observation masks.

 doi:10.20944/preprints202604.1573.v1

4.3. Basin-Safe Seeding: Global Navigation in Non-Convex Landscapes

While local metric normalization ensures stable iterative steps, it cannot resolve the **global basin selection problem** inherent in non-convex synchronization landscapes. To address the ambiguity between timing offset τ and phase increment ω, we introduce a deterministic **Basin-Safe Seeding** layer via the coarse-to-fine classical front-end (`BL_OGTR_CF`).

This layer identifies a stable starting point θ_0 by minimizing the raw objective S_{raw} over a discretized grid:

$$(\hat{\tau}_{\text{cf}}, \hat{\omega}_{\text{cf}}) \approx \arg\min_{\tau,\omega} S_{\text{raw}}(\tau, \omega, z = 0). \tag{14}$$

This design provides a *shielding effect*: the classical front-end resolves the global navigation on the information landscape, placing the local refiner in a favorable basin where the PT-aware and geometry-normalized updates can converge within the reportable projected channel. The synergistic method, `PO_OGTR_CF_INIT`, thus integrates global robustness with local geometric precision.

4.4. Safeguarded Acceptance and Divergence Reporting

To ensure auditable outcomes, the update is governed by a safeguarded acceptance mechanism (Algorithm 1). This layer evaluates candidates on the reportable objective $S_{\text{acc}}(\theta) = S_{\text{base}}(\theta)$, allowing for backtracking or explicit rejection.

Algorithm 1 Safeguarded Acceptance Mechanism

Require: Current iterate θ_{cur}, proposed step $\Delta\theta$, score S_{acc}, factor β, max backtracks $B_{\max}$
$S_{\text{cur}} \leftarrow S_{\text{acc}}(\theta_{\text{cur}})$
for $b = 0$ to $B_{\max}$ **do**
  $\theta_{\text{cand}} \leftarrow \theta_{\text{cur}} + \beta^b \Delta\theta$
  if $S_{\text{acc}}(\theta_{\text{cand}})$ is finite **and** $S_{\text{acc}}(\theta_{\text{cand}}) \le S_{\text{cur}} + \tau_{\text{acc}}$ **then**
    return `ACCEPT`, θ_{cand}, b
  end if
end for
return `REJECTED_STEP`, θ_{cur}, $B_{\max}$

For consistent stability tracking, we define the **effective divergence indicator**:

$$\text{diverged}_{\text{eff}} \triangleq \mathbb{I}\big[(\texttt{status} = \texttt{REJECTED\_STEP}) \vee (\texttt{is\_diverged}) \vee (S_{\text{acc}} > S_{\text{cap}})\big], \tag{15}$$

which provides a unified reporting flag for numerical failures, explicit rejections, and cap-triggered instabilities. Finally, we implement a **Cross-Objective Fairness Panel** (4.5) to evaluate all methods across both S_{raw} and S_{base} objectives, ensuring operational transparency.

4.5. Cross-Objective Fairness Panel

Because P0-type methods optimize S_{base} while the classical baselines optimize S_{raw}, we adopt a cross-objective fairness panel. At the final returned estimate $\hat{\theta}$, we explicitly evaluate and log both objectives for all methods. This ensures operational transparency and prevents metric-switching ambiguity across reportable and raw paths.

5. Experimental Protocol: Core Benchmarks, Regime-Boundary Sweeps, and Mechanism Diagnostics

This section establishes the empirical protocol used throughout the paper. To ensure strict reproducibility and fair comparison under irregular observations, the problem instances, method roles, metrics, and sweep definitions are definition-locked. The evidence is organized in two layers:

1. a *core estimation benchmark* for stability and headline performance under locked scenarios; and
2. auxiliary *Phase I/II diagnostics* that explain the onset, robustness, and mechanism of the observed gain.

 doi:10.20944/preprints202604.1573.v1

5.1. Core Estimation Scenarios

We consider synthetic link-estimation instances with parameter vector $\theta = (\tau, \omega, z)$, where τ is a circular shift in samples, ω is a CFO-like phase increment in radians/sample, and $z \in \mathbb{C}$ is a PT-odd nuisance mixing coefficient. Unless stated otherwise, experiments use block length $N = 64$ and nominal SNR $= 10$ dB.

The definition-locked scenario set is summarized in Table 1. The first four scenarios serve as the primary core benchmark. In addition, a strong-coupling PT-sensitive scenario is used as a robustness axis in the regime-boundary sweeps.

Table 1. Definition-locked scenario set used in the core benchmark and regime-boundary analysis. All scenarios share $\tau_{\text{true}} = 0.05$ samples and $\omega_{\text{true}} = -0.03$ rad/sample.

Scenario	τ_{true}	ω_{true}	z_{true}	Generative description	Role in evaluation
S1_Ideal	0.05	-0.03	$0 + 0j$	Ideal AWGN baseline.	Nominal convergence reference under stationary noise and full sampling.
S2_Gappy	0.05	-0.03	$0 + 0j$	Bernoulli missingness with nominal sample dropout encoded through $w[n]$.	Primary irregular-sampling stress test.
S3_Hetero	0.05	-0.03	$0 + 0j$	Heteroscedastic noise with strong variance contrast across the frame.	Primary stress test for curvature normalization under irregular weighting.
S_PT_z0	0.05	-0.03	$0 + 0j$	PT-constrained scenario where the reportable channel excludes PT-odd contamination.	Mechanism-validation scenario for ghost-mode suppression and low-z onset analysis.
S_PT_zNZ_strong	0.05	-0.03	strong nonzero z	Strong PT-sensitive scenario with nontrivial nuisance coupling.	Robustness axis for SNR and missingness sweeps in the regime-boundary analysis.

5.2. Methods and Their Roles

The evaluated method family is summarized in Table 2. The table is organized by *role* rather than by a single ranking, because the central claim of the paper is layered: the classical front-end resolves coarse basin ambiguity, while the PT-aware local refiner contributes an additional gain inside the basin.

Table 2. Evaluated methods and their roles in the manuscript.

Method	Fisher-like preconditioner	PT gatekeeper	Primary optimized objective	Role in this paper
P0	✓	✓	S_{base}	PT-aware local refiner operating in the reportable channel.
B1	✓	×	S_{raw}	Ablation removing the gatekeeper while retaining diagonal scaling.
B2	none ($D = I$)	×	S_{raw}	Ablation removing both PT-even gating and diagonal curvature normalization.
BL_OGTR_CF	–	×	S_{raw}	Primary classical front-end reference: deterministic coarse-to-fine raw-objective search.
BL_OGTR	–	×	S_{raw}	Auxiliary coarse classical reference used as a transparency anchor.
P0_OGTR_CF_INIT	✓	✓	Front-end: S_{raw}; local refinement: S_{base}	Headline hybrid method: basin-safe classical seed followed by PT-aware local refinement.

For the core stability benchmark, iterative methods are evaluated under locked step-size and safeguard configurations, including conservative and aggressive settings. The front-end baselines remain deterministic. Unless otherwise stated, the default iterative policy in the core benchmark is the fixed-offset initialization inherited from the locked reporting protocol lineage; however, as explained in Section 5.5, this policy should not be interpreted as a truth-near initialization.

 doi:10.20944/preprints202604.1573.v1

5.3. Phase I Regime-Boundary Analysis

Beyond the core benchmark, a Phase I regime-boundary analysis is used to characterize where the hybrid gain turns on and how it behaves under degraded observation conditions. This phase consists of three locked sweeps:

- a z**-sweep**, implemented by overriding z_{true} on top of S_PT_z0 to localize the onset of PT-sensitive gain;
- a **PT-sensitive missing-rate sweep**, performed in S_PT_zNZ_strong to test robustness under increasing missingness;
- a **PT-sensitive SNR sweep**, also performed in S_PT_zNZ_strong to test robustness across degraded observation quality.

The central Phase I quantities are two operational gain indicators measured relative to the primary classical reference `BL_OGTR_CF`:

1. **material_rmse_gain**, indicating a practically nontrivial reduction in RMSE_ω; and
2. **pt_sensitive_gain**, indicating simultaneous improvement in the coupled-sector proxy abs_z_drift and the reportable base loss.

This distinction is important. Material RMSE_ω gain may already appear in the classical-tangent regime, whereas PT-sensitive gain appears only when the nuisance-coupled sector becomes weakly but nontrivially observable.

For the missing-rate axis, only *post-fix* mask-sweep results are treated as scientifically valid. Earlier pre-fix PT mask runs are excluded from manuscript-level conclusions because the PT scenarios did not correctly honor the requested mask override in that earlier implementation.

5.4. Phase II Mechanism Diagnostics

Phase II is devoted to explaining *why* the hybrid is stable. These diagnostics are not primary benchmark evidence; rather, they are mechanism-facing analyses designed to interpret the gain observed in the core benchmark and Phase I.

The Phase II suite includes:

- **init-policy ablation**, comparing fixed-offset and cold-zero starts;
- **seed-path diagnostics**, recording initial loss, first-step loss change, and seed-to-truth distances;
- **offset/seed ablation**, decomposing the fixed-offset policy into τ-only, ω-only, z-only, and combined components;
- ω**-threshold and** ω**-micro scans**, used to localize the harmful basin onset along the ω direction;
- **conditional** τ **scans**, used to test whether τ acts as an independent harmful direction or primarily as an amplifier inside the bad-ω region.

These diagnostics support the manuscript's mechanism claim that the hybrid's stability is best explained by *basin-safe seeding*, with the dominant harmful seed direction associated with ω rather than with generic nonzero initialization.

5.5. Initialization Policy Note

A terminology clarification is essential for interpreting the Phase II results. The default iterative start recorded in the reporting lineage has historically been referred to as `current_init`. In the present manuscript, however, that policy should be understood as a *fixed-offset initialization*, not as a truth-near initialization.

This distinction matters because the mechanism diagnostics show that the harmful behavior of the default start is not due to nonzero initialization per se, but is strongly associated with the direction of the imposed offset—most notably the ω component. Accordingly, throughout the manuscript we interpret this policy operationally as *fixed-offset init*, and we reserve diagnostic language such as *exact-truth init* and *cold-zero init* for the explicit mechanism-facing variants introduced only in Phase II.

 doi:10.20944/preprints202604.1573.v1

5.6. Metrics, Fairness, and Gain Definitions

For each run, we record absolute parameter errors, root-mean-square errors (especially RMSE_ω and RMSE_τ), step acceptance outcomes, rejection indicators, leakage diagnostics, coupled-sector drift proxies, and computational cost measured in machine-agnostic objective evaluations.

Because the PT-aware methods optimize the reportable score S_{base} while the classical baselines optimize S_{raw}, the cross-objective fairness panel described in Section 4.5 is enforced throughout: both objectives are explicitly evaluated and logged at the final returned estimate for all methods.

For Phase I, gain statements are made relative to `BL_OGTR_CF`. The manuscript uses the following operational language:

- **general refinement gain**: a practically nontrivial improvement in RMSE_ω relative to `BL_OGTR_CF`, even when no coupled-sector advantage is yet visible;
- **PT-sensitive gain**: a simultaneous improvement in coupled-sector diagnostics and reportable objective value, indicating that the hybrid is doing more than simply refining the classical tangent regime.

Whenever onset is discussed, it is reported as an *interval* induced by the sampled grid, not as a sharp critical threshold.

5.7. Reproducibility Note

All primary and diagnostic experiments use fixed seed lists and locked configurations to ensure strict numerical reproducibility [18]. The locked reporting protocol uses a fixed seed list

$$\{42, 123, 456, 7, 13, 202, 314, 2718, 4096, 999\},$$

and uses the QFRT public technical anchor repository (`qfrt-evidence-public`) as the primary public evidence surface. The anchor is cited here as a curated reporting and traceability layer, with canonical documentation, scope, manifest, and manuscript-facing reporting files providing the public audit trail. Full execution infrastructure and high-volume raw runtime outputs are not part of this public technical-anchor surface.

5.8. Public Evidence Availability

The public evidence package for this manuscript is the QFRT public technical anchor, `qfrt-evidence-public`. It should be cited as a curated reporting and traceability surface for the tables, figures, locked configuration snapshots, and artifact crosswalks used in this paper. It is not presented as the canonical execution repository for regenerating all runs. Availability and scope are defined in the anchor documentation, with the public inventory recorded in `artifacts/manifest.json` and manuscript-facing reporting rooted at `artifacts/reporting/`.

6. Results: Stability, Two-Layer Gain, and Seed-Path Mechanism

This section reports the main empirical findings under the definition-locked protocol introduced in Section 5. The evidence is organized in three layers. First, we establish that the PT-aware local refiner remains numerically stable under irregular observations, whereas the structural ablations fail in distinct and diagnosable ways. Second, we show that the headline hybrid method exhibits a two-layer gain structure relative to the primary classical reference `BL_OGTR_CF`: a general refinement gain that is already present in the classical-tangent regime, and a PT-sensitive gain that turns on only for small but nonzero z. Third, we summarize the regime-boundary and mechanism-oriented diagnostics that explain why the hybrid is stable, emphasizing the role of basin-safe seeding and the ω-dominant harmful seed direction.

The major result tables in this section are aligned with the curated reporting surface of the QFRT public technical anchor. Frozen summary tables serve as the paper's evidence layer, and

 doi:10.20944/preprints202604.1573.v1

`artifacts/reporting/paper_artifact_map.csv` provides the table-to-artifact mapping for traceability.

6.1. Stability Under Irregular Observations

The first requirement of the estimator is numerical stability under irregular observations. Table 3 summarizes the governance outcomes for the core stress matrix under the aggressive safeguarded configuration. The comparison is intentionally shown on P0, B1, and B2 because this isolates the effect of PT-even gatekeeping and local curvature normalization before the classical front-end seed is introduced.

Table 3. Governance outcomes across the four core benchmark scenarios. Entries show rejection rates over 10 fixed seeds under the aggressive safeguarded configuration.

Scenario	P0 reject rate	B1 reject rate	B2 reject rate
S1_Ideal	0%	100%	100%
S2_Gappy	0%	100%	100%
S3_Hetero	0%	100%	100%
S_PT_z0	0%	100%	100%

P0 is the only iterative estimator that remains fully accepted across all four irregular scenarios. The no-gatekeeper ablation B1 is systematically rejected, indicating that diagonal scaling alone does not prevent drift into non-reportable directions. The no-Fisher ablation B2 fails even more directly, confirming that Euclidean updates are numerically fragile under severe mask- and weight-induced curvature distortion.

To ensure that this apparent stability is not being artificially supplied by the safeguard, Table 4 reports the backtracking statistics for P0. Together, Tables 3 and 4 serve as the core governance summaries and should be read against the same curated reporting layer used for public review. The zero-backtrack outcome across the full core matrix shows that the safeguard acts as a protection layer rather than as a hidden performance booster.

Table 4. Safeguard transparency for P0. The stable geometry of P0 requires no backtracking intervention from the protection layer in the core benchmark.

Scenario	P0: median total_backtracks [IQR]	P0: OK termination rate
S1_Ideal	0 [0,0]	100%
S2_Gappy	0 [0,0]	100%
S3_Hetero	0 [0,0]	100%
S_PT_z0	0 [0,0]	100%

A representative mechanism-facing visualization is shown in Figure 2. In the designated S_PT_z0 setting, where the reportable channel excludes PT-odd contamination, persistent $\hat{z} \neq 0$ is interpreted as ghost-mode leakage. The figure is aligned with the public figure mapping and configuration materials provided through the technical anchor. It shows that P0 suppresses such drift, whereas the ablations enter non-reportable directions before either diverging or being rejected.

 doi:10.20944/preprints202604.1573.v1

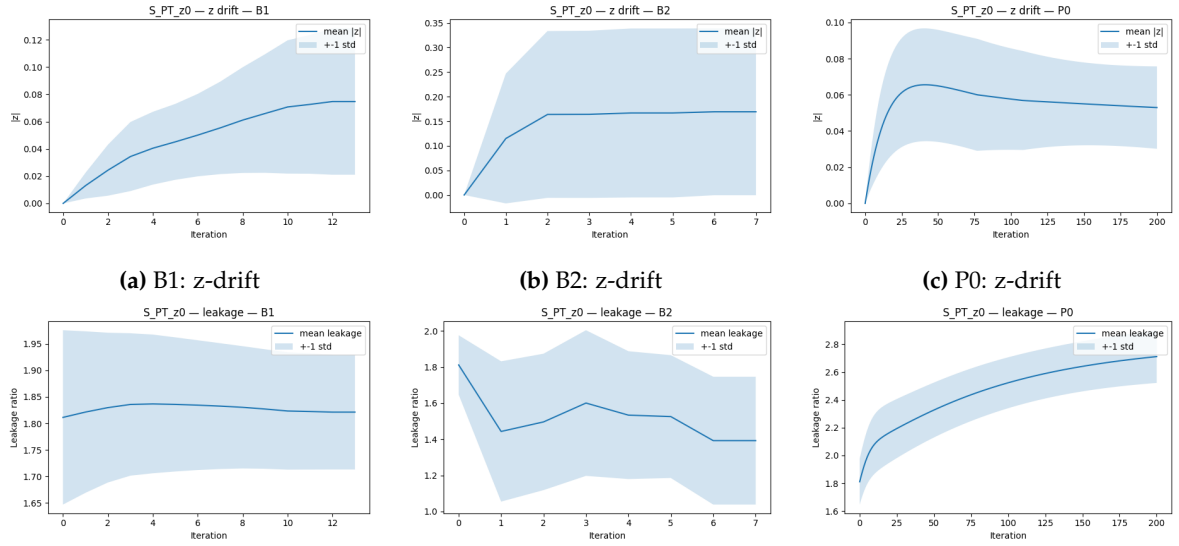

(a) B1: z-drift **(b)** B2: z-drift **(c)** P0: z-drift

(d) B1: leakage trace **(e)** B2: leakage trace **(f)** P0: leakage trace

Figure 2. Representative PT-gatekeeper diagnostics in S_PT_z0**.** Top: z-drift trajectories for B1, B2, and P0. Bottom: PT-odd leakage traces. The proposed PT-even gatekeeper suppresses non-reportable drift, whereas the ablations enter nuisance-dominated directions before rejection or numerical failure.

6.2. Two-Layer Gain Under the Fixed-Offset Policy

We now turn from core stability to the headline hybrid estimator `P0_OGTR_CF_INIT`, evaluated relative to the primary classical reference `BL_OGTR_CF`. Under the fixed-offset policy used in the locked reporting protocol, the hybrid exhibits a two-layer gain structure. Table 5 summarizes the key onset markers extracted from the z-axis regime-boundary sweep and matches the frozen public reporting summaries.

Table 5. Two-layer gain summary for the headline hybrid method relative to the primary classical reference. The onset quantities are reported as empirical grid-induced markers, not as sharp critical thresholds.

Headline method	Reference	First material RMSE_ω gain at $\lvert z\rvert$	Last no PT-sensitive gain at $\lvert z\rvert$	First PT-sensitive gain at $\lvert z\rvert$
P0_OGTR_CF_INIT	BL_OGTR_CF	0.000000	0.016763	0.021190

Table 5 highlights a practically nontrivial RMSE_ω improvement already present at $z = 0$, indicating that the hybrid provides a *general refinement gain* even in the classical-tangent regime. However, the stronger coupled-sector advantage—simultaneous improvement in the reportable base loss and the nuisance-coupled proxy—does not appear immediately. It turns on only after z becomes small but nonzero. This separation motivates the language used throughout the paper: the hybrid does not exhibit a single monolithic gain, but a layered gain structure with distinct operational meanings.

6.3. Onset of PT-Sensitive Gain in the Low-z Regime

The onset of PT-sensitive gain is summarized more explicitly in Table 6. The key point is that the onset should be interpreted as an *interval* determined by the sampled z grid, not as a sharp critical value.

Table 6. Low-z onset summary for PT-sensitive gain. The onset interval is grid-induced and should not be over-interpreted as a sharp threshold.

Headline method	Reference	PT-sensitive onset interval in $\lvert z\rvert$	Interpretation
P0_OGTR_CF_INIT	BL_OGTR_CF	$(0.016763,\ 0.021190]$	Small but nonzero-z onset; not a sharp critical point

The data therefore support the following operational picture. The hybrid already improves RMSE_ω at $z = 0$, but the specifically PT-sensitive component of the gain turns on only once the

 doi:10.20944/preprints202604.1573.v1

nuisance-coupled sector becomes weakly but nontrivially visible. This onset occurs early—well before the strong-coupling regime—but it is still nonzero. The manuscript therefore reports the gain onset conservatively as an interval rather than as a critical point estimate.

6.4. Robustness Across SNR and Missingness

Having established the two-layer structure and its low-z onset, we next ask whether the hybrid's gain survives degraded observation quality. Table 7 summarizes the PT-sensitive SNR sweep performed in the strong PT-sensitive regime and follows the same curated reporting convention as the other core summary tables.

Table 7. SNR-axis robustness summary in the strong PT-sensitive regime.

Headline method	Reference	SNR range (dB)	All material gain	All PT-sensitive gain	Loss-gain trend	Strongest / weakest gain SNR (dB)
P0_OGTR_CF_INIT	BL_OGTR_CF	0 – 20	True	True	attenuating	0 / 20

Across the tested SNR range, the hybrid retains both material RMSE_ω gain and PT-sensitive gain. At the same time, the gain amplitude is not constant: the base-loss improvement is strongest at the lowest tested SNR and weakest at the highest tested SNR. This is precisely the sense in which the SNR-axis result is best described as *robust but attenuating*.

A similar pattern appears along the missingness axis, summarized in Table 8. Only the *post-fix* PT mask sweep is used here; earlier pre-fix PT-mask outputs are excluded from manuscript-level conclusions by construction.

Table 8. Missingness-axis robustness summary in the strong PT-sensitive regime. Only post-fix validated mask sweeps are admissible for manuscript-level conclusions.

Metric / Attribute	**Value**
Headline method	P0_OGTR_CF_INIT
Reference	BL_OGTR_CF
Sweep validated	True
Requested mask range	0.000 – 0.650
Realized mask range	0.000 – 0.633
Observed-fraction range	1.000 – 0.367
All material gain	True
All PT-sensitive gain	True
Loss-gain trend	attenuating

The realized mask statistics confirm that the post-fix sweep is valid and monotone. Within that tested range, the hybrid retains both layers of gain relative to the primary classical reference. As with the SNR sweep, however, the gain is not constant; it weakens as the observation pattern becomes more hostile. Thus, the missingness result should also be read as robust but attenuating, not as a claim of unlimited resistance to arbitrary sample loss.

6.5. Init-Policy and Seed-Path Mechanism

The regime-boundary sweeps establish *that* the hybrid gains exist. The mechanism diagnostics explain *why*. The most important conclusion from the init-policy and seed-path analyses is that the hybrid's apparent initialization robustness should not be attributed to intrinsic insensitivity of P0 itself. Instead, the evidence supports a seed-path interpretation.

Under the fixed-offset policy, standalone P0 remains strongly seed-sensitive. By contrast, the headline hybrid `P0_OGTR_CF_INIT` is nearly invariant across the tested fixed-offset and cold-zero policies. This difference is best explained by the role of the classical front-end: `BL_OGTR_CF` effectively takes over the basin-selection problem and supplies a basin-safe seed, thereby shielding the downstream PT-aware local refinement from the unfavorable seed geometry that destabilizes standalone P0.

 doi:10.20944/preprints202604.1573.v1

A second clarification concerns the default public-policy start, historically referred to in code as `current_init`. In the manuscript it should not be read as a truth-near initialization. The Phase II diagnostics show that it is better understood as a *fixed-offset initialization*. This point matters because several mechanism-facing results depend on how that imposed offset interacts with the local geometry.

6.6. ω-Dominant Basin Sensitivity

The final mechanism result is that harmful seed geometry is not generic; it is directionally structured. In the offset/seed diagnostics, the dominant harmful direction is associated with the ω offset. A z-only offset is comparatively benign in the S_PT_z0 setting, and a τ-only offset has more limited standalone impact. By contrast, an ω-only offset produces a substantial degradation, and the combined (τ, ω) offset reproduces the poor behavior of the fixed-offset policy. This indicates that the harmful basin onset is primarily controlled by the ω direction.

More refined threshold and micro-scan diagnostics further localize this effect. The harmful transition is not centered at $\omega = 0$ itself; rather, it appears after a modest displacement away from the truth-side ω and is best described as a narrow onset interval rather than as a sharp cutoff. Conditional τ scans then suggest that τ acts mainly as an amplifier once the trajectory has entered the bad-ω region, rather than as the dominant source of instability on its own.

Taken together, these Phase II diagnostics support the manuscript's layered methodological interpretation. The primary function of `BL_OGTR_CF` is to provide a basin-safe seed. Once that seed is supplied, the PT-aware local refiner can deliver both the general refinement gain and the PT-sensitive gain without being dominated by the harmful seed geometry that affects standalone P0.

7. Discussion and Limitations

The empirical results presented in Section 6 demonstrate that the layered framework provides a measurable stability advantage in irregular observation environments. This section formalizes the link between the projected-model geometry and signal-processing constraints, providing a dictionary for interpreting the framework's gains and defining its operational boundaries.

7.1. Implementation Dictionary

The architecture of the layered estimator is motivated by information geometry and constrained optimization. Table 9 summarizes the correspondence between the conceptual language used in the paper and the functional roles played by each component.

Table 9. Implementation dictionary for the layered estimator.

Mathematical Object	Conceptual Role	Signal Processing Interpretation
Base Manifold ($\mathcal{M}$)	Information Landscape	The time-frequency resource grid under irregular weighting.
Metric Surrogate (D_i)	Riemannian Metric	Local Metric Normalization to correct for heteroscedastic curvature.
PT-Projector (Π_+)	Reportability Constraint	Matched subspace filter used to reject non-reportable "ghost-mode" leakage.
B-lite Update ($\Delta\theta$)	Curvature-Normalized Update	Preconditioned adaptive step for monotonic descent on a distorted objective landscape.
Nuisance Term (z)	Non-reportable Ghost Mode	Indicator for nuisance-coupled artifacts such as I/Q imbalance or non-reciprocal hardware effects.
Basin-Safe Seed (θ_0)	Global Navigation	Coarse-to-fine search to resolve non-convex ambiguity.

This dictionary is intended as an interpretive aid. It emphasizes how the method combines weighted geometry, explicit constraints, and seeded local refinement without claiming a broader physical theory beyond the estimation setting studied here.

 doi:10.20944/preprints202604.1573.v1

7.2. Interpreting B-lite as Curvature-Normalized Descent

The significant reduction in RMSE observed in heteroscedastic noise (Scenario S3) can be understood through the lens of curvature-normalized descent. Standard gradient descent assumes a Euclidean geometry, which leads to instability when the noise variance fluctuates across the frame. Our B-lite update uses the diagonal-Fisher matrix as a local metric surrogate, effectively smoothing the curvature of the loss landscape. In operational terms, this dampens the step in high-noise or high-curvature regions while allowing more effective progress in reliable sectors, helping explain why the preconditioned update remains stable where unpreconditioned methods diverge.

7.3. Future Outlook

While the current manuscript focuses on a *projected complex model*, the framework may be extendable to richer parameterizations and higher-dimensional state spaces. Such extensions are outside the scope of the present paper, whose evidence is limited to the stable-estimation setting studied here.

Future work can therefore focus on richer yet still auditable curvature surrogates, especially where limited off-diagonal coupling may be worth modeling without abandoning the present methods-paper scope. Beyond that, only carefully scoped extensions of the projected-complex formulation are justified when the same auditability and stability criteria can be maintained.

7.4. Limitations and Boundary of Claims

To ensure academic independence, we emphasize the specific boundaries of the current evidence:

1. **Geometric Approximation:** The diagonal-Fisher surrogate is a lightweight approximation. While sufficient for stabilizing local refinement, it does not capture full off-diagonal parameter coupling, which remains the responsibility of the basin-safe seeding layer.
2. **Dataset Scope:** The stability gains are validated within a definition-locked synthetic protocol. Although the stress tests emulate irregular observation conditions of practical interest, the present evidence remains a methods evaluation rather than a full system study.
3. **Initialization Sensitivity:** As shown in Phase II diagnostics, standalone local refinement remains seed-sensitive. The stability of the framework is a *layered effect* arising from the synergy between the front-end search and the geometry-aware refiner.
4. **Boundary of Claim:** The reported gains are acceptance-conditioned and mechanism-specific. They support a second-layer refinement role under missing and heteroscedastic observations, not a universal replacement claim for classical estimators.

8. Conclusion

This paper studied stable signal estimation under irregular observation conditions characterized by missing samples, heteroscedastic weighting, and a reportability constraint on nuisance-contaminated solution paths. In this setting, the main challenge is not only nominal estimation accuracy, but whether the update remains numerically stable, whether failures are auditable, and whether the returned solution stays aligned with the reportable channel.

The central result is that a *layered* estimator provides the most accurate interpretation of the observed gain. The proposed framework combines a PT-even gatekeeper, a diagonal-Fisher local refinement step, a deterministic basin-safe classical seed, and a safeguarded acceptance mechanism. Empirically, this yields three main findings. First, the PT-aware local refiner is stably governable under the locked irregular stress matrix, while structurally simplified ablations fail in distinct and diagnosable ways. Second, the headline hybrid method exhibits a two-layer gain structure relative to the primary classical reference: a general refinement gain already present in the classical-tangent regime, and a PT-sensitive gain that turns on only for small but nonzero z. Third, in the strong PT-sensitive regime, that gain remains robust—though attenuating—across the tested SNR and missingness ranges.

The mechanism diagnostics further clarify why the hybrid is stable. The evidence does not support the interpretation that P0 is intrinsically initialization-robust. Instead, it supports a seed-path

 doi:10.20944/preprints202604.1573.v1

explanation: the classical coarse-to-fine front-end resolves the coarse basin-selection problem and supplies a basin-safe seed, while the PT-aware local refinement contributes the second-layer gain once the trajectory is inside a favorable basin. Within the present diagnostic design, the dominant harmful seed direction is associated with the ω offset, while τ acts primarily as an amplifier in the bad-ω region.

Taken together, these results support a restrained but useful methodological conclusion. The paper does not show that a PT-aware estimator universally replaces classical baselines. Rather, it shows that under missing and heteroscedastic observations, a PT-aware local refiner can add a measurable and interpretable second-layer gain when paired with a strong classical front-end and governed by an explicit reportability constraint.

The most natural next steps are therefore algorithmic rather than system-level. Richer structured curvature surrogates may reduce the remaining dependence on the front-end seed by modeling cross-parameter coupling more directly, and additional definition-locked stress tests may clarify where the observed second-layer gain persists and where it does not.

Data Availability Statement: The public technical anchor for this work, including curated reporting artifacts, traceability metadata, and figure-facing evidence surfaces, is openly available at https://github.com/ice91/qfrt-evidence-public under the BSD 3-Clause License. This repository is provided as a curated reporting and traceability surface for the results discussed in this manuscript; it is not presented as the full execution environment for regenerating all runs.

Conflicts of Interest: The author is employed by Chunghwa Telecom Co., Ltd. This work was conducted in the author's personal research capacity, outside the scope of the author's company-assigned research duties, and without the use of company research funding, company computing infrastructure, company-confidential materials, or proprietary internal datasets. Chunghwa Telecom had no role in the design of the study; in the collection, analysis, or interpretation of data; in the writing of the manuscript; or in the decision to publish the results. The author declares no competing financial interest directly related to the results reported in this manuscript.

Appendix A. Variational Safety: Residual-Projector Commutativity

A critical requirement for the layered estimator is that the gatekeeping mechanism be applied consistently to the residuals and local derivatives used by the update rule. We refer to this implementation requirement as *Variational Safety*. In this appendix, we record the modest commutation property needed by the proposed projected complex model: because the PT-even projector is a fixed real-linear operator on signal space, it commutes with differentiation of the residual vector with respect to the continuous real-valued estimation parameters.

Appendix A.1. Definitions and Operators

Let the signal space be denoted by $\mathbb{C}^N$. For any synchronization parameter vector $\theta = (\tau, \omega, z)$, let $u(\theta) \in \mathbb{C}^N$ denote any differentiable signal-space vector, such as the residual $r(\theta)$ or the forward model $H(\theta)$. The operational PT-action, as defined in Section 3.3, is:

$$(\mathrm{PT}\, x)[n] \triangleq \overline{x[N-1-n]}. \tag{A16}$$

The PT-even projector is defined as the real-linear operator:

$$\Pi_+ \triangleq \frac{1}{2}(I + \mathrm{PT}), \tag{A17}$$

where I is the identity operator.

 doi:10.20944/preprints202604.1573.v1

Appendix A.2. Residual/Jacobian Commutativity

Proposition 1 (Residual-Projector Commutativity)**.** *For any differentiable signal-space vector $u(\theta) \in \mathbb{C}^N$ and any real-valued parameter component θ_i, the fixed PT-even projector satisfies:*

$$\partial_{\theta_i}\left(\Pi_+ u(\theta)\right) = \Pi_+ \, \partial_{\theta_i} u(\theta). \tag{A18}$$

Consequently, for the projected residual $r_{\text{eff}}(\theta) = \Pi_+ r(\theta)$, the effective Jacobian used by the local update can be written as

$$J_{i,\text{eff}}(\theta) \triangleq \partial_{\theta_i} r_{\text{eff}}(\theta) = \Pi_+ \, \partial_{\theta_i} r(\theta). \tag{A19}$$

Proof. Substituting the definition of Π_+ into the left-hand side (LHS) of (A18) gives

$$\begin{aligned} \text{LHS} &= \partial_{\theta_i}\left[\frac{1}{2}\left(u(\theta) + \text{PT}[u(\theta)]\right)\right] \\ &= \frac{1}{2}\left(\partial_{\theta_i} u(\theta) + \partial_{\theta_i}\text{PT}[u(\theta)]\right), \end{aligned} \tag{A20}$$

by linearity of differentiation.

Now consider the term $\partial_{\theta_i}\text{PT}[u(\theta)]$. The PT-action (A16) operates on the temporal indices n and the algebraic structure of the signal-space vector values. The mapping $n \mapsto N - 1 - n$ is a fixed permutation of the sample space, and the operation does not depend on the continuous synchronization parameters $\theta = (\tau, \omega, z)$. Therefore, the partial derivative with respect to θ_i commutes with this fixed signal-space operator:

$$\partial_{\theta_i}\text{PT}[u(\theta)] = \text{PT}[\partial_{\theta_i} u(\theta)]. \tag{A21}$$

Substituting this back into (A20) gives

$$\begin{aligned} \text{LHS} &= \frac{1}{2}\left(\partial_{\theta_i} u(\theta) + \text{PT}[\partial_{\theta_i} u(\theta)]\right) \\ &= \Pi_+ \, \partial_{\theta_i} u(\theta) = \text{RHS}. \end{aligned} \tag{A22}$$

Equation (A19) follows by choosing $u(\theta) = r(\theta)$. □

Remark 1. This result should be read as an implementation-consistency statement, not as a claim that a scalar objective can be projected in signal space or that all raw-objective gradients are equivalent after projection. It justifies computing the B-lite update in (13) from the projected residual and the consistently projected Jacobian, so that the local derivative information is evaluated in the same reportable channel as the acceptance objective.

Appendix B. Numerical Stability and Loss-Cap Sensitivity

To ensure that the stability conclusions of this work are not artifacts of a specific reporting threshold, we conduct a sensitivity check on the effective divergence indicator defined in (15).

The primary results in Section 6 utilize a reporting cap of $S_{\text{cap}} = 10^6$. In this appendix, we compare the divergence behavior against a relaxed cap of $S_{\text{cap}} = 10^8$ using the frozen run records from the aggressive safeguarded benchmark.

The analysis confirms that the qualitative stability profile of the layered estimator is invariant to the choice of S_{cap}. Loosening the cap does not rescue the unstable methods (B1 and B2), nor does tightening it artificially improve the performance of P0. This robustness supports the interpretation that the observed stability is consistent with the intended local metric normalization and PT-gatekeeping mechanism rather than only a numerical tuning effect.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 22 April 2026 doi:10.20944/preprints202604.1573.v1

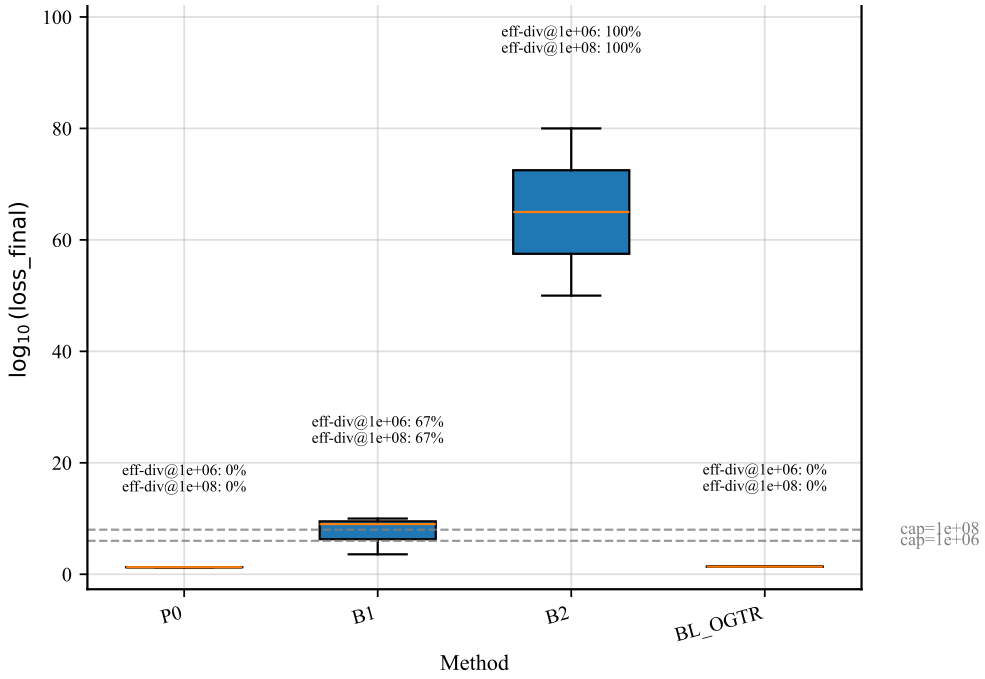

Figure 3. Loss-cap sensitivity of the effective divergence indicator across 10^6 and 10^8. The proposed PT-aware local refiner (P0) maintains 100% stability under both tested caps, while unstable ablations (B1, B2) remain systematically divergent.

Appendix C. Phase I Regime-Boundary Summary

This appendix collects the Phase I regime-boundary summaries that support the main-text claims about two-layer gain, low-z onset, and robustness across SNR and missingness. These summaries are reported relative to the primary classical reference `BL_OGTR_CF` and should be interpreted as compact summary tables drawn from the curated public reporting layer rather than as replacements for the underlying detailed sweep artifacts.

Appendix C.1. Two-Layer Gain and Low-z Onset

Table 10 summarizes the gain markers extracted from the z-axis regime-boundary sweep. The key distinction is between material RMSE_ω improvement and the onset of PT-sensitive gain. As with the main-text summary tables, this appendix table is aligned with the frozen reporting surface used for public review.

Table 10. Phase I z-axis gain summary relative to `BL_OGTR_CF`. The onset quantities are reported as grid-induced empirical markers rather than as sharp critical thresholds.

Headline method	Reference	First material gain at $\lvert z\rvert$	Last no PT-sensitive gain at $\lvert z\rvert$	First PT-sensitive gain at $\lvert z\rvert$	PT-sensitive onset interval
P0_OGTR_CF_INIT	BL_OGTR_CF	0.000000	0.016763	0.021190	(0.016763, 0.021190]

The interpretation of Table 10 is the same as in the main text: a general refinement gain is already present in the classical-tangent regime, while the PT-sensitive gain turns on only for small but nonzero z. The onset is therefore operationally localized to an interval rather than to a single critical value.

Appendix C.2. SNR-Axis Robustness Summary

Table 11 summarizes the Phase I SNR sweep in the strong PT-sensitive regime. The hybrid retains both layers of gain across the tested SNR range, while the gain amplitude attenuates as the observations become cleaner and the classical baseline catches up.

 doi:10.20944/preprints202604.1573.v1

Table 11. Phase I SNR-axis robustness summary in the strong PT-sensitive regime.

Headline method	Reference	SNR range (dB)	All material gain	All PT-sensitive gain	Loss-gain trend	Strongest / weakest gain SNR (dB)
P0_OGTR_CF_INIT	BL_OGTR_CF	0 – 20	True	True	attenuating	0 / 20

This table should be read as a compact robustness summary, not as a claim of invariant gain amplitude across all observation qualities.

Appendix C.3. Missingness-Axis Robustness Summary

Table 12 summarizes the post-fix PT mask sweep. Only post-fix outputs are admissible for manuscript-level interpretation, because earlier pre-fix PT mask runs did not correctly honor the requested mask override.

Table 12. Phase I missingness-axis robustness summary in the strong PT-sensitive regime. Only post-fix validated mask sweeps are admissible for manuscript-level conclusions.

Metric / Attribute	Value
Headline method	P0_OGTR_CF_INIT
Reference	BL_OGTR_CF
Sweep validated	True
Requested mask range	0.000 – 0.650
Realized mask range	0.000 – 0.633
Observed-fraction range	1.000 – 0.367
All material gain	True
All PT-sensitive gain	True
Loss-gain trend	attenuating

The realized-mask statistics confirm that the post-fix sweep is validity-aware and monotone. Within that tested range, the hybrid retains both layers of gain, though again with an attenuating trend.

Appendix C.4. Integrated Regime Interpretation

For completeness, Table 13 condenses the Phase I narrative into a single regime-facing summary. It is not needed to support the main text's core tables, but it is useful as an appendix-level interpretive map.

Table 13. Integrated Phase I regime interpretation.

Regime label	Operational condition	Interpretation
General refinement gain	$z = 0$, material gain present, PT-sensitive gain absent	The hybrid improves RMSE_ω even in the classical-tangent regime, but without yet exhibiting a coupled-sector advantage.
PT-sensitive onset	Small but nonzero z; onset interval $(0.016763,\ 0.021190]$	The coupled-sector benefit turns on early but not at exactly $z = 0$; the onset is interval-valued, not a sharp threshold.
PT-sensitive robust SNR regime	All tested SNR points retain both gain layers	The gain remains present throughout the tested SNR range, but with attenuating amplitude.
PT-sensitive robust mask regime	Post-fix validated mask sweep retains both gain layers	The gain remains present throughout the tested missingness range, but again with attenuating amplitude.

Appendix D. Phase II Mechanism Diagnostics

This appendix collects the mechanism-facing diagnostics that explain why the headline hybrid is stable. Unlike Appendix C, which summarizes regime-boundary behavior using frozen compact tables, the present appendix is interpretive: it organizes the Phase II diagnostics into a coherent seed-path

 doi:10.20944/preprints202604.1573.v1

and offset-geometry story. These diagnostics are important for explanation, but they should not be mistaken for universal geometric laws.

Appendix D.1. Init-Policy Ablation and Seed-Path Interpretation

The init-policy ablation shows that standalone P0 remains strongly sensitive to the starting seed, whereas the headline hybrid `P0_OGTR_CF_INIT` is nearly invariant across the tested fixed-offset and cold-zero policies. This observation is best understood as a seed-path effect rather than as intrinsic initialization robustness of the local PT-aware refiner itself.

The key interpretation is therefore layered:

- standalone P0 is a basin-sensitive local refiner;
- `BL_OGTR_CF` resolves the coarse basin-selection problem;
- the hybrid is stable because the front-end seed places the local refiner inside a favorable basin.

Appendix D.2. Offset/Seed Decomposition

The offset/seed ablation decomposes the fixed-offset start into τ-only, ω-only, z-only, and combined components. The resulting pattern is directionally structured rather than generic.

Operationally, the evidence supports the following ordering:

- a z-only offset is comparatively benign in S_PT_z0;
- a τ-only offset has limited standalone impact;
- an ω-only offset causes substantial degradation;
- a combined (τ, ω) offset reproduces the poor behavior of the default fixed-offset policy.

This supports the interpretation that the dominant harmful seed direction is associated with ω, not with nonzero initialization in general.

Appendix D.3. ω-Threshold and Micro-Scan Interpretation

The threshold and micro-scan diagnostics refine the mechanism story further. The harmful transition is not centered at $\omega = 0$ itself. Instead, the bad-basin onset appears after a modest displacement away from the truth-side ω, and is more accurately described as a narrow onset interval than as a single critical point.

This result is conceptually consistent with the main-text low-z onset language: both the coupled-sector gain onset and the seed-path instability onset are best reported as interval-valued operational findings induced by finite diagnostic grids.

Appendix D.4. Conditional τ Scans

The conditional τ scans ask whether τ is itself the dominant harmful direction, or whether it primarily amplifies instability once the trajectory has already entered a bad-ω region.

The evidence favors the second interpretation. In a safe-ω region, changing τ has comparatively limited effect. In threshold-near and clearly unsafe ω regions, however, the same τ variations become more damaging. Thus, within the present diagnostic design, τ is best interpreted as an *amplifier* of bad-ω geometry rather than as the primary source of the harmful basin.

Appendix D.5. Compact Mechanism Summary

For convenience, Table 14 summarizes the mechanism-facing conclusions that are consistent with the Phase II evidence used in the manuscript.

Taken together, the diagnostics in this appendix support the manuscript's layered methodological interpretation: the classical front-end solves the coarse basin problem, and the PT-aware local refiner then contributes both the general refinement gain and the PT-sensitive gain once the trajectory is safely seeded.

 doi:10.20944/preprints202604.1573.v1

Table 14. Compact Phase II mechanism summary.

Diagnostic block	Main observation	Interpretation
Init-policy ablation	P0 is seed-sensitive; hybrid is nearly invariant across tested policies	Hybrid stability is best explained by front-end seed shielding, not by intrinsic init-robustness of the local refiner.
Seed-path diagnostic	The classical front-end provides a stable seed trajectory before local PT-aware refinement	`BL_OGTR_CF` acts as a basin-safe seed provider.
Offset/seed ablation	ω-only offset is more harmful than τ-only or z-only offset	The dominant harmful seed direction is associated with ω.
ω threshold / micro scan	Bad-basin onset occurs after modest displacement away from truth-side ω	The instability onset is interval-valued, not a sharp cutoff at $\omega = 0$.
Conditional τ scan	τ becomes more harmful inside threshold-near or unsafe ω regions	τ acts mainly as an amplifier in the bad-ω region.

References

1. M. Giordani and M. Zorzi, "Non-Terrestrial Networks in the 6G Era: Challenges and Opportunities," *IEEE Network*, vol. 35, no. 2, pp. 244-251, 2021.
2. T. M. Schmidl and D. C. Cox, "Robust frequency and timing synchronization for OFDM," *IEEE Transactions on Communications*, vol. 45, no. 12, pp. 1613-1621, Dec. 1997.
3. J. W. Cooley and J. W. Tukey, "An algorithm for the machine calculation of complex Fourier series," *Mathematics of Computation*, vol. 19, no. 90, pp. 297-301, Apr. 1965.
4. A. V. Oppenheim and R. W. Schafer, *Discrete-Time Signal Processing*, 2nd ed. Upper Saddle River, NJ, USA: Prentice Hall, 1999.
5. S. M. Kay, *Fundamentals of Statistical Signal Processing, Volume I: Estimation Theory*. Englewood Cliffs, NJ, USA: Prentice Hall, 1993.
6. P. H. Moose, "A technique for orthogonal frequency division multiplexing frequency offset correction," *IEEE Transactions on Communications*, vol. 42, no. 10, pp. 2908-2914, Oct. 1994.
7. C. M. Bender and S. Boettcher, "Real Spectra in Non-Hermitian Hamiltonians Having PT Symmetry," *Physical Review Letters*, vol. 80, no. 24, pp. 5243-5246, Jun. 1998.
8. S.-I. Amari, *Information Geometry and Its Applications*. Tokyo, Japan: Springer, 2016.
9. S.-I. Amari, "Natural gradient works efficiently in learning," *Neural Computation*, vol. 10, no. 2, pp. 251-276, Feb. 1998.
10. R. J. A. Little and D. B. Rubin, *Statistical Analysis with Missing Data*, 2nd ed. Hoboken, NJ, USA: Wiley, 2002.
11. R. J. A. Little and D. B. Rubin, *Statistical Analysis with Missing Data*, 3rd ed. Hoboken, NJ, USA: Wiley, 2019.
12. P. J. Huber, *Robust Statistics*. New York, NY, USA: Wiley, 1981.
13. P. J. Huber and E. M. Ronchetti, *Robust Statistics*, 2nd ed. Hoboken, NJ, USA: Wiley, 2009.
14. J. Martens, "Deep learning via Hessian-free optimization," in *Proc. 27th Int. Conf. Machine Learning (ICML)*, Haifa, Israel, Jun. 2010, pp. 735-742.
15. J. Martens and R. Grosse, "Optimizing neural networks with Kronecker-factored approximate curvature," in *Proc. 32nd Int. Conf. Machine Learning (ICML)*, Lille, France, Jul. 2015, pp. 2408-2417.
16. R. Grosse and J. Martens, "A Kronecker-factored approximate Fisher matrix for convolution layers," *arXiv preprint arXiv:1602.01407*, 2016.
17. J. Nocedal and S. J. Wright, *Numerical Optimization*, 2nd ed. New York, NY, USA: Springer, 2006.
18. O. E. Gundersen and S. Kjensmo, "State of the art: Reproducibility in artificial intelligence," in *Proc. AAAI Conf. Artificial Intelligence (AAAI)*, New Orleans, LA, USA, Feb. 2018.
19. Chien-Chih Chen, "Projection-Defined Physicality in PT-Symmetric Quaternionic Spacetime," *Preprints*, no. 208283, 2026, doi: 10.20944/preprints202604.1054.v1.

 doi:10.20944/preprints202604.1573.v1

20. Chien-Chih Chen, "Guaranteed Tensor Luminality from Symmetry: A PT-Even Palatini Torsion Framework," *Symmetry*, vol. 18, no. 1, Art. no. 170, 2026, doi: 10.3390/sym18010170.