

Finite Observation

A Spectral Theory of Nonequilibrium Visibility

J. R. Dunkley

18th March 2026

Abstract

We study finite observation in irreducible finite-state continuous-time Markov nonequilibrium steady states through the local Donsker–Varadhan Hessian at the steady-state minimiser. In reduced Fisher coordinates, the nonequilibrium correction factors through a weighted signal operator and the detailed-balance backbone. For fixed-rank orthogonal observation in this reduced Euclidean Fisher geometry, the visible correction is controlled by a Ky Fan envelope, together with an associated retention hierarchy, hidden tail, and weighted stable rank. This hierarchy determines the spectral profile of the visible nonequilibrium correction and, equivalently, the corresponding singular-energy profile of the weighted signal operator, and therefore its nonzero singular-value magnitudes. We also identify the optimal orthogonal observer at fixed rank and show that a Haar-random rank-constrained observer retains, on average, exactly the corresponding dimensional fraction of the available signal. A separate statistical question arises after comparison with the observed detailed-balance backbone: in the compressed-Hessian Gaussian model, distinguishability is governed by a positive shadow operator whose spectrum can differ from the raw visibility spectrum except in special commuting or isotropic-backbone regimes. The same weighted geometry extends to abstract finite-dimensional metric settings, and hence to the linearised BKM-symmetrised operator setting.

Contents

1	Introduction	3
2	Setting and observable normal forms	5
3	The Donsker-Varadhan bridge	9
4	The observable envelope	17
5	Alignment, optimal observers, and generic observers	25
6	The shadow operator and Gaussian distinguishability	32
7	Spectroscopy and the anatomy of hidden nonequilibrium structure	50
8	Projection suppression and shared observed-Gaussian behaviour	56
9	Abstract weighted geometry and operator-space transfer	60
10	Linear quantum fluctuation extension	64
11	Worked Examples	66
12	Related-work connections	85
13	Scope	87
A	Reduced-coordinate identities	90
B	Linear algebra of weighted signal operators	92
C	Random projector calculations	95
D	Gaussian comparison formulas	96
E	Worked details for the literature examples	97
F	Operator-space technicalities	98
G	Interpretation and boundaries	99
H	Interpretive synthesis for the worked examples	100
I	An analytic note on the scope of the low-rank universality claim	100

1 Introduction

1.1 The finite observation problem

A nonequilibrium system may carry irreversibility across many microscopic directions while only a small observable sector is accessible. Hidden transitions, finite time resolution, finite spatial resolution, detector noise, and coarse graining all lead to the same structural question: what can a finite observer actually see of the nonequilibrium fluctuation structure?

In finite-dimensional Markov nonequilibrium steady states, the natural empirical-measure fluctuation object is the local Hessian of the Donsker–Varadhan rate function at its minimiser π [1–4]. A finite observer does not see the full Hessian. The observer only sees its compression to a rank-limited sector. The finite observation problem is therefore precise: given an irreducible nonequilibrium generator and an observable subspace of dimension d , how much of the nonequilibrium correction survives projection, and how does that visible share depend on the geometry of the system and the geometry of the observer? The detailed-balance geometry is the reference geometry adopted here for measuring this local Gaussian nonequilibrium correction.

We show that, in the finite-dimensional Markov Gaussian regime treated here, this question reduces to a spectral compression problem in reduced Fisher coordinates. The visible share is governed by a weighted singular-value hierarchy. The hidden share is its spectral tail. The effective complexity of visibility is controlled by weighted stable rank.

1.2 Motivation

Recent work across stochastic thermodynamics, coarse graining, and quantum nonequilibrium theory has made the finite observation problem sharper and more urgent. The literature contains lower bounds on entropy production under partial observation, hidden data, coarse temporal sampling, unresolved transitions, faulty observation maps, and coarse graining [5–12]. Closely related work studies estimator quality for dissipation under lumping, milestoneing, hidden cycles, current fluctuations, and finite statistics, and proposes random or optimised coarse-graining strategies for irreversibility detection [13–18].

Restricted-information and coarse-grained quantum thermodynamics make the same point from another direction: observation changes the effective thermodynamic description [19–23]. At a broader level, the thermodynamic uncertainty-relation programme has clarified how much of the field is organised around scalar fluctuation bounds rather than observer-resolved spectral structure [24, 25], and very recent counterexamples show that even two widely used model-free partial-observation bounds, the waiting-time and TUR bounds, admit no universal ordering on the same observed link [26].

What is missing from this landscape, as far as we are aware, is a finite-dimensional law for the visible share of the nonequilibrium correction itself in the Markov Gaussian setting treated here. The current literature is rich in estimators, bounds, and coarse-grained descriptions. The contribution of the present paper is to identify, for orthogonal rank-constrained observers in reduced Euclidean Fisher coordinates, the envelope of the projected Donsker–Varadhan correction, the associated observer-design law at fixed rank, the shadow criterion for observed Gaussian indistinguishability in the compressed-Hessian model, and the corresponding spectral readout of the hidden correction.

1.3 Main claim and conceptual picture

The central claim is that finite observation in the Donsker–Varadhan Gaussian regime has a finite-dimensional spectral normal form. In reduced Euclidean Fisher coordinates, the nonequilibrium correction factors through the reduced skew channel and the detailed-balance backbone via the weighted signal map A_{DV} . Finite observation therefore becomes a rank-constrained orthogonal compression problem for that map.

This leads to two distinct spectral questions. The first is raw visibility: how much nonequilibrium signal survives projection, how much is hidden in the spectral tail, and which observer of rank d is best aligned with the signal front. That question is governed by the singular spectrum of A_{DV} and the associated Ky Fan envelope hierarchy.

The second is detectability relative to the detailed-balance backbone: once the observed correction is compared to the observed backbone in its natural metric, a second positive operator appears, and its spectrum governs Gaussian distinguishability. These two spectra coincide only in special commuting or isotropic-backbone regimes. The raw and detectability spectra therefore read the same nonequilibrium deviation in two different equilibrium coordinate systems: the ambient detailed-balance backbone, and the backbone induced on the observed sector after projection. Raw nonequilibrium signal and statistically detectable nonequilibrium signal therefore need not live in the same observer directions. The theory therefore distinguishes directions that are structurally invisible under the observation map from directions that remain visible in principle but are weak or poorly aligned in the observable geometry.

Within the orthogonal-observer model used here, the paper gives a finite-dimensional theory of visibility, observer design, detectability, and spectral readout. The bridge theorem identifies the structure of the correction, the envelope gives the visible law, the shadow gives the distinguishability law, and the full hierarchy recovers the singular-energy profile of the correction, equivalently the nonzero singular-value magnitudes of A_{DV} .

1.4 Main results

The first result is the bridge theorem. We prove that the Donsker–Varadhan correction is quadratic in the reduced skew channel. This is the representation step that turns the fluctuation problem into a weighted operator problem.

The second result is the finite-observation law. For every orthogonal rank-constrained observer, the visible nonequilibrium content is given by a Ky Fan envelope of the weighted signal map. The retention hierarchy, hidden tail, and weighted stable rank follow from this representation.

The third result is the observer-design law. Visible signal is an alignment score against the left singular modes of the weighted signal map. This yields the optimal orthogonal observer, its gap-controlled uniqueness criterion, the worst observer, and the Haar-random rank- d baseline in expectation.

The fourth result is the two-spectrum theorem. The observed Gaussian law in the compressed-Hessian model is governed by a positive shadow operator, and backbone whitening yields a second spectral hierarchy for detectability relative to the detailed-balance backbone. In general this detectability spectrum is distinct from the raw visibility spectrum.

The fifth result is spectral completeness of the envelope hierarchy for magnitudes. The full envelope hierarchy recovers the ordered singular-energy profile $\frac{1}{4}s_k(A_{DV})^2$, equivalently the eigenvalue profile of Δ_{DV} , and hence the nonzero singular-value magnitudes of A_{DV} .

The sixth result is a transfer theorem. The weighted visibility, envelope, alignment, and spectral-readout structure transfer to abstract finite-dimensional metric settings and hence to linearised BKM operator geometry.

1.5 Contribution

The paper gives a finite-dimensional structural solution of the orthogonal-observer problem of finite observation in the Markov Gaussian regime treated here. Its contribution is to identify the operator geometry that organises what finite observers can still see within that model.

More specifically, it proves a bridge from the skew sector to the fluctuation correction, a rank-constrained envelope law, an observer-design law with a Haar baseline, a Gaussian shadow criterion, a second spectrum for detectability distinct from raw visibility, and an envelope hierarchy that determines the singular-energy profile of the correction and hence the singular-value magnitudes of A_{DV} .

1.6 Roadmap

Section 2 fixes the finite-dimensional Markov setting, the Fisher normal form, the notion of a finite observer, and the observable normal-form language. Section 3 proves the bridge theorem and identifies the nonequilibrium correction as a positive correction operator together with its weighted signal factorisation. Section 4 derives the finite-observation law: the Ky Fan envelope, retention hierarchy, hidden tail, stable-rank obstruction, and the direct-sum allocation reading of the envelope.

Section 5 develops the observer-design layer through the alignment theorem, optimal and worst observers, and the Haar-random baseline. Section 6 proves the shadow and detectability theory, making explicit the split between raw visibility and backbone-normalised distinguishability, and includes a fixed-channel Gaussian noisy-observation extension. Section 7 proves spectral completeness of the envelope hierarchy for singular-value magnitudes and develops the induced observability order between systems.

2 Setting and observable normal forms

2.1 Finite-dimensional Markov nonequilibrium steady states

Let $X = \{1, \dots, n\}$ be a finite state space, and let Q be the generator of an irreducible continuous-time Markov chain on X . Thus $Q_{ij} \geq 0$ for $i \neq j$, $\sum_j Q_{ij} = 0$ for each i , and irreducibility guarantees a unique stationary law

$$\pi = (\pi_1, \dots, \pi_n), \quad \pi_i > 0, \quad \sum_i \pi_i = 1, \quad (2.1)$$

satisfying $\pi Q = 0$. We are interested in the nonequilibrium steady-state case, meaning that Q need not satisfy detailed balance with respect to π .

Throughout the paper we work at the level of stationary empirical-measure fluctuations. The fluctuation object in this setting is the local Hessian of the Donsker–Varadhan rate function at its minimiser π , viewed in the Fisher geometry associated with the stationary law. This Fisher viewpoint is not merely convenient notation: in nonequilibrium stochastic thermodynamics it has a direct geometric and physical interpretation, and in the reversible Markov setting it sits on a well-developed information-geometric background [27, 28]. The detailed-balance geometry provides the reference unit system in which the Gaussian nonequilibrium correction is measured. The entire finite-observation theory developed below takes place in the tangent space at π , after passing to the Euclideanised Fisher coordinates adapted to that geometry.

Let

$$D_\pi := \text{diag}(\pi_1, \dots, \pi_n), \quad e_\pi := (\sqrt{\pi_1}, \dots, \sqrt{\pi_n})^T. \quad (2.2)$$

The Euclideanised Fisher tangent space is

$$T_\pi := \{u \in \mathbb{R}^n : \langle u, e_\pi \rangle = 0\}. \quad (2.3)$$

Equivalently, T_π is the codimension-one subspace orthogonal to the stationary direction in Fisher coordinates. We denote by

$$P_\pi := I - e_\pi e_\pi^T \quad (2.4)$$

the orthogonal projector onto T_π .

The Fisher metric in density coordinates becomes the standard Euclidean inner product after this transformation, so from this point onward all orthogonality, projection, adjoints, and singular values in the classical Markov setting are taken with respect to the Euclidean structure on T_π , unless stated otherwise.

2.2 Fisher geometry and the normal form

Define the Fisher-conjugated generator

$$L_\pi := D_\pi^{1/2} Q D_\pi^{-1/2}. \quad (2.5)$$

Its restriction to the tangent space carries the full linearised stationary fluctuation geometry. Since T_π is L_π -invariant, we may regard

$$K := L_\pi|_{T_\pi} \quad (2.6)$$

as an operator on T_π .

The central structural decomposition is the Fisher normal form

$$K = G + J, \quad (2.7)$$

where

$$G := \frac{K + K^T}{2}, \quad J := \frac{K - K^T}{2}. \quad (2.8)$$

Thus G is self-adjoint and J is skew-adjoint on T_π . In the present work, G is the dissipative backbone and J is the skew, antisymmetric sector carrying the nonequilibrium part relative

to the detailed-balance backbone. This symmetric/antisymmetric normal-form viewpoint is consonant with earlier steady-state fluctuation analyses of driven diffusions [29] and with the network-cycle description of nonequilibrium steady states originating in Schnakenberg theory [30]. The detailed-balance reference object is obtained by setting $J = 0$ while keeping the same stationary law and dissipative geometry.

The nonequilibrium correction to the Donsker-Varadhan Hessian will be shown in Section efsec:exact-dv-bridge to be carried entirely by the skew sector after the reduced-coordinate analysis, while the observed detailed-balance backbone is governed by the projected dissipative sector.

To work concretely, choose an orthonormal matrix

$$C \in \mathbb{R}^{n \times (n-1)} \quad (2.9)$$

whose columns form an orthonormal basis of T_π . Then

$$C^T C = I_{n-1}, \quad C C^T = P_\pi. \quad (2.10)$$

Reduced coordinates $x \in \mathbb{R}^{n-1}$ are identified with tangent vectors by $u = Cx$. In these coordinates the Fisher-conjugated generator becomes

$$\widehat{K} := C^T L_\pi C, \quad (2.11)$$

with corresponding reduced decomposition

$$\widehat{K} = \widehat{G} + \widehat{J}, \quad \widehat{G} = \frac{\widehat{K} + \widehat{K}^T}{2}, \quad \widehat{J} = \frac{\widehat{K} - \widehat{K}^T}{2}. \quad (2.12)$$

All finite-dimensional statements in the main text will be expressed on this reduced Euclidean space. For notational simplicity, we will often suppress the distinction between the operator on T_π and its matrix in reduced coordinates when no ambiguity arises. Where the reduced-coordinate nature matters, hats will be retained. This dissipative-plus-skew normal form already foreshadows the later operator-space transfer, where the same structural split reappears in metric operator geometry.

2.3 Observable sectors and finite observers

A finite observer is modelled by a rank-constrained orthogonal projector on the reduced Fisher space. Concretely, let

$$\Pi : \mathbb{R}^{n-1} \rightarrow \mathbb{R}^{n-1} \quad (2.13)$$

be a Euclidean orthogonal projector of rank d , and let

$$U := \text{Ran } \Pi. \quad (2.14)$$

The rank d is the observable dimension. It is the number of independent fluctuation directions actually accessible to the observer.

This choice of projector is the mathematical form of finite observation used here. It covers rank-constrained observable sectors, reduced coordinates, compressions to slow manifolds, and more abstract finite-resolution measurement ideals, provided the observation is represented by orthogonal compression in the underlying Euclideanised geometry. The exact envelope law proved below is for orthogonal rank- d observers in this sense. More

general coarse-graining maps require different machinery, with the fixed-channel Gaussian noisy-observation section giving one controlled extension. Later sections will explain how this structure extends to weighted operator-space settings, where orthogonality is taken with respect to a nontrivial metric.

Given Π , the observer does not see the full operator K , the full detailed-balance Hessian, or the full Donsker-Varadhan Hessian. The observer sees only their compressions to U . The finite observation problem is therefore a rank-constrained compression problem on a fluctuation operator whose full structure is initially hidden.

2.4 Projected observable normal forms

The projected observable normal form associated with Π is

$$N_\Pi(Q) := (U, g_\Pi, G_\Pi, J_\Pi), \quad (2.15)$$

where g_Π is the Euclidean metric restricted to U , and

$$G_\Pi := \Pi \widehat{G} \Pi|_U, \quad J_\Pi := \Pi \widehat{J} \Pi|_U. \quad (2.16)$$

This is the observable quadruple inherited by finite observation from the full Markov generator.

Observable equivalence is governed by such projected data once the finer nonequilibrium correction is sufficiently small. The work proves a stronger result. The projected quadruple is the observable backbone of a signal theory in which the nonequilibrium correction is placed in weighted spectral form and its visible share under Π is determined explicitly.

Three levels of observable structure will later reappear:

1. the purely metric or static level, carried by U and g_Π ,
2. the dissipative level, carried by G_Π ,
3. the full projected normal form, carried by (g_Π, G_Π, J_Π) .

The theory developed below refines this hierarchy by inserting one more object: the projected nonequilibrium correction to the Donsker-Varadhan Hessian. That correction is the signal of hidden irreversibility.

2.5 The Donsker-Varadhan Hessian as the fluctuation object

Let I_{DV} denote the Donsker-Varadhan rate function for the empirical measure of the irreducible Markov chain generated by Q . Since π is the unique stationary law, π is the minimiser of I_{DV} , and the Hessian of I_{DV} at π is positive definite on the reduced tangent space. We denote this Hessian in reduced Euclidean coordinates by

$$H_{DV}. \quad (2.17)$$

The canonical detailed-balance reference is obtained by Fisher symmetrisation. Let

$$L_\pi := D_\pi^{1/2} Q D_\pi^{-1/2}, \quad G_{\text{full}} := \frac{1}{2}(L_\pi + L_\pi^T), \quad Q_0 := D_\pi^{-1/2} G_{\text{full}} D_\pi^{1/2}. \quad (2.18)$$

Then Q_0 has the same stationary law π , satisfies detailed balance with respect to π , and has

reduced Fisher generator $\widehat{G}$. The corresponding detailed-balance Hessian on the reduced Euclidean Fisher space is therefore

$$H_0 := -\frac{1}{2} \widehat{G}. \quad (2.19)$$

Because the detailed-balance reference is irreducible, $-\widehat{G}$ is positive definite on the reduced space, so $H_0 \succ 0$ and is therefore invertible there. Under this sign convention, the positive conductance form is represented by $-\widehat{G}$. The nonequilibrium correction is

$$\Delta_{DV} := H_{DV} - H_0. \quad (2.20)$$

At this point, Δ_{DV} is only a definition. The bridge section then identifies it as a positive correction operator built from the skew sector and the detailed-balance backbone.

The object of interest is not entropy production by itself, nor the full rate function, nor a pathwise current observable. It is the fluctuation-side nonequilibrium correction encoded in H_{DV} . Once that correction is understood, finite observation becomes a problem of spectral retention.

3 The Donsker-Varadhan bridge

3.1 Reduced coordinates and the conductance identity

The bridge from the Donsker-Varadhan variational principle to the spectral theory developed later in the paper passes through a reduced-coordinate quadratic envelope. The full derivation is carried out in Appendix A. Here we isolate the structural cancellation that drives everything that follows.

Starting from the variational representation of the Donsker-Varadhan rate function, one introduces reduced logarithmic coordinates near stationarity and expands the envelope to second order. This produces a reduced quadratic fluctuation problem governed by two coupled ingredients. One is a current-like reduced operator carrying the skew dependence. The other is a reduced conductance matrix controlling the quadratic envelope itself. A priori, both the dissipative and skew sectors could have entered that conductance matrix.

They do not. The skew sector contributes no quadratic form of its own, so the conductance matrix of the full nonequilibrium envelope depends only on the dissipative backbone. This is the reason the final bridge is purely quadratic in the skew channel, with no mixed dissipative-skew leakage.

Let M_Q denote the reduced conductance matrix produced by the full nonequilibrium envelope construction, and let M_G denote the corresponding reduced conductance matrix for the detailed-balance reference problem obtained by suppressing the skew sector. Both objects are defined canonically by the same reduced-coordinate construction; the explicit formulas are recorded in Appendix A.

Proposition 3.1 (Reduced conductance identity). *In the finite-dimensional irreducible Markov setting described above, the reduced conductance matrix of the full Donsker-Varadhan quadratic envelope coincides with that of the detailed-balance reference problem:*

$$M_Q = M_G. \quad (3.1)$$

Equivalently, the quadratic conductance controlling the reduced Donsker-Varadhan Hessian depends only on the dissipative sector and is blind to the skew sector.

Proof. Write the Donsker-Varadhan variational formula as [1–3]

$$I_{DV}(\mu) = \sup_h \left(- \sum_{i,j} \mu_i q_{ij} (e^{h_j - h_i} - 1) \right). \quad (3.2)$$

Choose reduced Fisher coordinates $x \in \mathbb{R}^{n-1}$ and reduced logarithmic coordinates $y \in \mathbb{R}^{n-1}$ by

$$\mu(x) = \pi + D_\pi^{1/2} C x + O(|x|^2), \quad h(y) = D_\pi^{-1/2} C y. \quad (3.3)$$

Using the second-order reduced-coordinate expansion derived in Appendix A gives the reduced quadratic envelope

$$\mathcal{L}(x, y) = y^T \widehat{G} y - x^T \widehat{K} y + O(|(x, y)|^3), \quad (3.4)$$

where $\widehat{K} = C^T D_\pi^{1/2} Q D_\pi^{-1/2} C = \widehat{G} + \widehat{J}$. The key point is that the pure quadratic term in y is

$$y^T \widehat{K} y = y^T \widehat{G} y, \quad (3.5)$$

because $y^T \widehat{J} y = 0$ for every skew-symmetric matrix $\widehat{J}$. Hence the conductance block is determined entirely by $\widehat{G}$, and the nonequilibrium reduced conductance coincides with the detailed-balance one. Equivalently,

$$M_Q = M_G = -\widehat{G}. \quad (3.6)$$

The skew sector survives only in the current block $-x^T \widehat{K} y$.

Proposition 3.1 records the key asymmetry used in the bridge argument. The skew sector drives the nonequilibrium correction, but the quadratic conductance beneath the reduced envelope is controlled by the dissipative geometry alone. This is why the later factorisation of the correction is so rigid. The skew sector enters through a weighted signal map, but the weight itself is fixed by the detailed-balance backbone. The first antisymmetry input here is the identity $y^T \widehat{J} y = 0$, but the bridge theorem uses it inside the Legendre contraction, where the supremum over y removes the mixed dissipative-skew contribution at Hessian order. This cancellation is a level-2 empirical-measure statement, and should not be read as a statement about the unreduced level-2.5 fluctuation Hessian.

The rest of Section 3 exploits this asymmetry. The next subsection converts the reduced envelope into a quadratic representation of the nonequilibrium correction, which in turn becomes the signal operator governing finite observation.

Remark (Dirac-analogous factorisation of the DV Hessian). The bridge identity also admits a canonical first-order factorisation, using the principal positive square root of H_0 .

Define

$$D := H_0^{1/2} + \frac{1}{2}H_0^{-1/2}\widehat{J}^T. \quad (3.7)$$

Since $\widehat{J}$ is skew-symmetric, one has

$$D^T = H_0^{1/2} + \frac{1}{2}\widehat{J}H_0^{-1/2}, \quad (3.8)$$

and therefore

$$D^T D = H_0 + \frac{1}{4}\widehat{J}H_0^{-1}\widehat{J}^T = H_{DV}. \quad (3.9)$$

Thus $H_{DV} = D^T D$ exactly. Since $H_{DV} \succ 0$ on the reduced space, $\ker D = \{0\}$, so D is invertible. The operator D packages the detailed-balance backbone and the reduced skew channel into a single first-order object whose adjoint square recovers the full nonequilibrium fluctuation Hessian.

3.2 Quadraticity of the nonequilibrium correction

We now progress to the central bridge theorem. Proposition 3.1 implies that the reduced Donsker-Varadhan quadratic envelope is built on the detailed-balance conductance matrix alone. The skew sector therefore enters only through the reduced current channel. This yields a quadratic representation of the nonequilibrium correction.

To state the result cleanly, recall from Section 2 that the canonical detailed-balance reference is the Fisher-symmetrised one, with reduced Hessian

$$H_0 = -\frac{1}{2}\widehat{G}. \quad (3.10)$$

Let $\widehat{J}$ denote the canonical reduced skew channel produced by the reduced-coordinate envelope construction. The next theorem is the finite-dimensional bridge identity that identifies the fluctuation correction as a positive correction operator with an explicit weighted signal factorisation.

Theorem 3.2 (Quadraticity of the DV correction). *In the finite-dimensional irreducible Markov setting of Section 2, the Donsker-Varadhan Hessian decomposes as*

$$H_{DV} = H_0 + \Delta_{DV}$$

where the nonequilibrium correction is

$$\Delta_{DV} = \frac{1}{4}\widehat{J}H_0^{-1}\widehat{J}^T$$

Equivalently, Δ_{DV} is positive semidefinite and is quadratic in the reduced skew channel with no mixed dissipative-skew term.

Proof. By Proposition 3.1 and the reduced-coordinate expansion recalled above, the second-

order reduced envelope has the form

$$\mathcal{L}(x, y) = y^T \widehat{G} y - x^T \widehat{K} y + O(|(x, y)|^3), \quad \widehat{K} = \widehat{G} + \widehat{J}. \quad (3.11)$$

Set

$$S := -\widehat{G}. \quad (3.12)$$

Then S is positive definite on the reduced space and

$$\mathcal{L}(x, y) = -y^T S y + x^T (S - \widehat{J}) y + O(|(x, y)|^3) = -y^T S y + y^T (S + \widehat{J}) x + O(|(x, y)|^3), \quad (3.13)$$

using $\widehat{J}^T = -\widehat{J}$. Completing the square in y gives

$$\mathcal{L}(x, y) = -\left(y - \frac{1}{2} S^{-1} (S + \widehat{J}) x\right)^T S \left(y - \frac{1}{2} S^{-1} (S + \widehat{J}) x\right) + \frac{1}{4} x^T (S + \widehat{J})^T S^{-1} (S + \widehat{J}) x + O(|(x, y)|^3). \quad (3.14)$$

Taking the supremum over y , exactly as in the reduced variational contraction derived in Appendix A, yields

$$I_{DV}(\mu(x)) = \frac{1}{4} x^T (S + \widehat{J})^T S^{-1} (S + \widehat{J}) x + O(|x|^3). \quad (3.15)$$

Since $\widehat{J}^T = -\widehat{J}$,

$$(S + \widehat{J})^T S^{-1} (S + \widehat{J}) = (S - \widehat{J}) S^{-1} (S + \widehat{J}) = S - \widehat{J} S^{-1} \widehat{J} = S + \widehat{J} S^{-1} \widehat{J}^T. \quad (3.16)$$

Therefore

$$I_{DV}(\mu(x)) = \frac{1}{4} x^T (S + \widehat{J} S^{-1} \widehat{J}^T) x + O(|x|^3). \quad (3.17)$$

By definition of the reduced Hessian,

$$I_{DV}(\mu(x)) = \frac{1}{2} x^T H_{DV} x + O(|x|^3), \quad (3.18)$$

so

$$H_{DV} = \frac{1}{2} S + \frac{1}{2} \widehat{J} S^{-1} \widehat{J}^T. \quad (3.19)$$

For the detailed-balance reference $\widehat{J} = 0$, the same reduced envelope gives

$$H_0 = \frac{1}{2} S = -\frac{1}{2} \widehat{G}. \quad (3.20)$$

Hence $H_0^{-1} = 2S^{-1}$, and substitution gives

$$H_{DV} = H_0 + \frac{1}{4} \widehat{J} H_0^{-1} \widehat{J}^T. \quad (3.21)$$

Equivalently,

$$\Delta_{DV} = H_{DV} - H_0 = \frac{1}{4} \widehat{J} H_0^{-1} \widehat{J}^T. \quad (3.22)$$

The Gram form

$$\Delta_{DV} = \frac{1}{4} (\widehat{J} H_0^{-1/2}) (\widehat{J} H_0^{-1/2})^T \quad (3.23)$$

implies positive semidefiniteness.

Remark (Orthogonal invariance of the quadratic correction). The identity of Theorem 3.2 is invariant under orthogonal changes of reduced Euclidean Fisher coordinates. In particular, the spectrum of Δ_{DV} , and therefore every envelope, retention, and spectral-readout quantity derived from it, is coordinate-independent within the reduced Euclidean Fisher coordinates fixed by the paper.

Theorem 3.2 identifies Δ_{DV} as an explicit positive correction operator with a weighted signal factorisation at the level of the quadratic fluctuation correction. Once the detailed-balance backbone H_0 is fixed, the reduced skew channel determines the Gaussian nonequilibrium excess at Hessian order, with no additional mixed dissipative-skew term surviving in this reduced setting.

Two immediate consequences are worth recording.

Corollary 3.3 (Positivity and rank control). *The nonequilibrium correction Δ_{DV} is positive semidefinite. Moreover,*

$$\text{rank}(\Delta_{DV}) = \text{rank}(\hat{J}), \quad (3.24)$$

and the range of Δ_{DV} is exactly the signal subspace generated by the reduced skew channel.

Proof. Positivity follows from Theorem 3.2. Since H_0^{-1} is invertible,

$$\text{rank}(\hat{J}H_0^{-1}) = \text{rank}(\hat{J}). \quad (3.25)$$

Now set $D := \hat{J}H_0^{-1/2}$. Then

$$\hat{J}H_0^{-1}\hat{J}^T = DD^T, \quad (3.26)$$

and $\text{rank}(DD^T) = \text{rank}(D)$. Hence

$$\text{rank}(\hat{J}H_0^{-1}\hat{J}^T) = \text{rank}(\hat{J}). \quad (3.27)$$

The range statement follows from $\Delta_{DV} = \frac{1}{4}DD^T$, since $\text{Ran}(DD^T) = \text{Ran}(D)$, and $\text{Ran}(D) = \text{Ran}(\hat{J})$ because $H_0^{-1/2}$ is invertible.

Corollary 3.4 (Detailed-balance characterisation). *The following are equivalent:*

1. *the generator is detailed-balance in the Fisher normal form, that is, $J = 0$,*
2. *the reduced skew channel vanishes, that is, $\hat{J} = 0$,*
3. *the Donsker-Varadhan Hessian equals its detailed-balance reference, that is, $\Delta_{DV} = 0$.*

Proof. The equivalence of the first two statements is by construction of the reduced skew channel. The equivalence of the second and third follows from Theorem 3.2 and the positive definiteness of H_0 .

Corollaries 3.3 and 3.4 already show why Δ_{DV} is the right object for finite observation. It is the positive fluctuation correction carried by the skew sector and vanishes precisely at detailed balance.

3.3 Basis transport and the weighted signal operator

The exact quadraticity theorem immediately suggests the right signal object. Since H_0 is positive definite, define

$$A_{DV} := \widehat{J} H_0^{-1/2}. \quad (3.28)$$

Then

$$\Delta_{DV} = \frac{1}{4} A_{DV} A_{DV}^T. \quad (3.29)$$

This is the weighted signal operator of the paper, and Δ_{DV} is its induced positive correction operator.

The factorisation is conceptually simple but extremely strong. It says that the nonequilibrium correction is a covariance-type signal built from the skew channel after whitening by the detailed-balance Hessian. The dissipative sector therefore enters twice.

First, it provides the detailed-balance backbone H_0 . Second, it defines the metric in which the skew contribution is whitened and measured. At the level of currents and affinities, the cycle origin of this skew contribution is the same identified in Schnakenberg network theory and its fluctuation-theorem extension [30, 31].

To put the operator in its most transparent form, we diagonalise the detailed-balance backbone. Since H_0 is real symmetric positive definite, there exists an orthogonal matrix U and positive eigenvalues

$$0 < \mu_1 \leq \mu_2 \leq \dots \leq \mu_{n-1} \quad (3.30)$$

such that

$$H_0 = U \operatorname{diag}(\mu_1, \dots, \mu_{n-1}) U^T. \quad (3.31)$$

In this basis the weighted signal operator is

$$A_{DV} = U (U^T \widehat{J} U) \operatorname{diag}(\mu_1^{-1/2}, \dots, \mu_{n-1}^{-1/2}) U^T. \quad (3.32)$$

Thus the visible nonequilibrium signal is obtained by taking the skew sector in the eigenbasis of the detailed-balance Hessian and weighting each channel by the inverse square root of the corresponding detailed-balance stiffness.

This observation is the structural source of the stable-rank and envelope theory developed in Section 4. The relevant quantity is not the raw skew operator by itself. It is the skew operator after whitening by the detailed-balance geometry.

The factorisation also makes clear how the theory refines the unweighted signal picture. In the unweighted Euclidean setting, singular-energy retention is attached directly to an input operator. Here the physically relevant input operator is not $\widehat{J}$ alone but its Hessian-whitened version A_{DV} . The unweighted theory therefore appears as the unweighted shadow of the final weighted theory.

3.4 Slow observable block and converter structure

We now make precise the slow-block structure that underlies projection suppression and explains the intrinsic-slow and mixed-return channels.

Fix a rank- d observable projector Π and identify its range $U = \operatorname{Ran} \Pi$ with the slow observable

sector. For the block formula below we assume that U is H_0 -invariant, equivalently that Π commutes with H_0 . In physical terms, this means that the observed sector is aligned with detailed-balance relaxation modes, so the backbone is block diagonal across observed and hidden directions. Let $U^\perp$ denote its orthogonal complement in the reduced Euclidean space. Relative to the decomposition

$$\mathbb{R}^{n-1} = U \oplus U^\perp, \quad (3.33)$$

write

$$H_0 = \begin{pmatrix} H_{0,S} & 0 \\ 0 & H_{0,F} \end{pmatrix}, \quad \widehat{J} = \begin{pmatrix} J_{SS} & J_{SF} \\ J_{FS} & J_{FF} \end{pmatrix}. \quad (3.34)$$

The notation S stands for slow or observed, while F stands for fast or hidden. The block diagonal form of H_0 is the content of the H_0 -invariance assumption. Since $\widehat{J}$ is skew-symmetric, its blocks satisfy

$$J_{SS} = -J_{SS}^T, \quad J_{FF} = -J_{FF}^T, \quad J_{FS} = -J_{SF}^T. \quad (3.35)$$

The quadraticity theorem then gives a decomposition of the observed nonequilibrium correction on such spectral observable sectors.

Theorem 3.5 (Exact block formula on H_0 -invariant observed sectors). *Assume that the observable sector $U = \text{Ran } \Pi$ is H_0 -invariant. With notation as above, the slow observable block of the nonequilibrium correction is*

$$(\Delta_{DV})_{SS} = \frac{1}{4} J_{SS} H_{0,S}^{-1} J_{SS}^T + \frac{1}{4} J_{SF} H_{0,F}^{-1} J_{SF}^T. \quad (3.36)$$

Equivalently, relative to the chosen H_0 -adapted block decomposition $U \oplus U^\perp$, the observed correction splits into an intrinsic slow block and a hidden-return block from the complementary sector.

Proof. Block-decompose the quadraticity formula of Theorem 3.2 relative to the orthogonal splitting $U \oplus U^\perp$. Because U is H_0 -invariant, H_0 is block diagonal in the chosen basis, hence so is H_0^{-1} . The (S, S) -block of

$$\widehat{J} H_0^{-1} \widehat{J}^T \quad (3.37)$$

is therefore exactly

$$J_{SS} H_{0,S}^{-1} J_{SS}^T + J_{SF} H_{0,F}^{-1} J_{SF}^T. \quad (3.38)$$

Multiplying by $1/4$ gives the result.

Theorem 3.5 is the precise algebraic form of the observed block structure on H_0 -adapted sectors. The invariant observable object is the compressed correction $\Pi \Delta_{DV} \Pi$. Relative to the chosen H_0 -adapted splitting, the first term is the intrinsic slow correction generated entirely within the observed sector, while the second term is the hidden return from complementary directions into the observed sector, weighted by the inverse stiffness of the fast detailed-balance backbone. For a general non- H_0 -invariant projector one still has the compressed correction $\Pi \Delta_{DV} \Pi$, but not this simple block formula.

This formula is already enough to recover the structural content of projection suppression under slow-fast structural hypotheses. If the slow skew block is weak and the mixed return channel is either weak or heavily penalised by a stiff fast backbone, then the observed correction is small. But now this is no longer a theorem about disappearance under an external list of assumptions. It is the direct block-level anatomy of the correction.

The converter formula also reveals the correct physical interpretation of the fast backbone. A large fast detailed-balance stiffness suppresses the mixed return channel because the hidden sector becomes expensive as a route for observable nonequilibrium correction. This is the direct finite-dimensional mechanism behind fast-block coercivity.

We record the most basic consequence.

Corollary 3.6 (Fast-block coercive bound). *Assume that the fast detailed-balance block satisfies*

$$H_{0,F} \succeq \mu_F I \quad (3.39)$$

for some $\mu_F > 0$. Then

$$\frac{1}{4} J_{SF} H_{0,F}^{-1} J_{SF}^T \preceq \frac{1}{4\mu_F} J_{SF} J_{SF}^T. \quad (3.40)$$

In particular, the mixed return channel is suppressed by increasing fast-sector coercivity.

Proof. Since $H_{0,F} \succeq \mu_F I$, one has $H_{0,F}^{-1} \preceq \mu_F^{-1} I$. Conjugating by J_{SF} preserves positive-semidefinite order and yields the claim.

Corollary 3.6 is intentionally elementary. The point is not to optimise constants here. The point is to expose the mechanism. The fast hidden sector suppresses observable nonequilibrium not by magic, but because the bridge weights hidden return channels by the inverse fast detailed-balance stiffness.

3.5 Interpretation

Section 3 identifies the nonequilibrium correction as a positive weighted correction operator carried by the skew channel and generated by the weighted signal operator whitened by the detailed-balance backbone.

This changes the logic of the entire theory. Once

$$\Delta_{DV} = \frac{1}{4} A_{DV} A_{DV}^T \quad (3.41)$$

is known, finite observation becomes a problem of signal retention under orthogonal compression. The observer does not interact directly with the full generator, nor with the full rate function, but with a positive signal operator whose singular-energy content can be tracked.

Structural suppression and unweighted comparison can now be re-read from this point of view. Observable structure collapses to a projected normal form once the nonequilibrium correction becomes negligible under compression, and the singular-value mechanism behind optimal retention and rank-limited visibility becomes exact once the bridge theorem identifies the correct weighted signal operator.

The next section turns that signal into the observable envelope, the retention hierarchy, and the stable-rank complexity measure governing finite observation within the orthogonal-observer class.

In this form the finite-observation anatomy can be read cleanly in four layers. The algebraic source of the correction is the reduced skew channel $\widehat{J}$. The signal actually compressed

by finite observation is the backbone-weighted operator $A_{DV} = \widehat{J}H_0^{-1/2}$. On H_0 -adapted observed sectors, the hidden contribution appears explicitly through the mixed return block $\frac{1}{4}J_{SF}H_{0,F}^{-1}J_{SF}^T$. At fixed observable rank d , the optimally hidden budget is the singular-energy tail of A_{DV} , while statistical detectability after whitening is governed by a second spectrum and should not be conflated with raw visible mass.

4 The observable envelope

4.1 Visible nonequilibrium signal of a finite observer

Section 3 identified the nonequilibrium correction as the positive signal operator

$$\Delta_{DV} = \frac{1}{4} A_{DV} A_{DV}^T, \quad A_{DV} := \widehat{J}H_0^{-1/2}. \quad (4.1)$$

We now solve the finite observation problem for this signal. This section is about *raw visibility*: how much nonequilibrium signal survives rank-constrained orthogonal compression, how much is hidden in the spectral tail, and how the observation budget should be allocated across decoupled sectors. Statistical distinguishability relative to the detailed-balance backbone is a different question and is deferred deliberately to Section 6.

Let Π be a rank- d orthogonal projector on the reduced Euclidean Fisher space. The observer sees only the compression of Δ_{DV} to the d -dimensional sector $U = \text{Ran } \Pi$. The natural scalar measure of visible nonequilibrium content is therefore

$$V_{\Pi}^{DV} := \text{tr}(\Pi \Delta_{DV} \Pi). \quad (4.2)$$

This quantity is the observable nonequilibrium signal retained by Π . It is the total projected singular energy of the DV correction that survives finite observation.

The first theorem of the section shows that V_{Π}^{DV} has an exact projector formula.

Theorem 4.1 (Exact projector formula for orthogonal observers). *For every orthogonal projector Π on the reduced Euclidean Fisher space,*

$$V_{\Pi}^{DV} = \text{tr}(\Pi \Delta_{DV} \Pi) = \frac{1}{4} \|\Pi A_{DV}\|_F^2. \quad (4.3)$$

Equivalently, the visible nonequilibrium content retained by Π is the projected Frobenius energy of the weighted signal operator A_{DV} within the orthogonal-observer class fixed in this paper.

Proof. Using the exact weighted Gram form from Theorem 3.2,

$$\Delta_{DV} = \frac{1}{4} A_{DV} A_{DV}^T. \quad (4.4)$$

Hence

$$V_{\Pi}^{DV} = \text{tr}(\Pi \Delta_{DV} \Pi) = \frac{1}{4} \text{tr}(\Pi A_{DV} A_{DV}^T \Pi). \quad (4.5)$$

Since Π is an orthogonal projector,

$$\text{tr}(\Pi A_{DV} A_{DV}^T \Pi) = \text{tr}((\Pi A_{DV})^T (\Pi A_{DV})) = \|\Pi A_{DV}\|_F^2. \quad (4.6)$$

This gives the claimed formula.

Remark (Orthogonal-basis invariance of visible signal). Under an orthogonal change of reduced Euclidean Fisher basis, $A_{DV} \mapsto O A_{DV}$ and $\Pi \mapsto O \Pi O^T$, the value V_Π^{DV} is unchanged. Thus the visible signal is an invariant of the orthogonal observer class, not of a particular coordinate representative.

Theorem 4.1 is the first answer to the finite observation question. Once the bridge has been closed, the visible nonequilibrium signal of an orthogonal observer is no longer a vague statement about how much correction survives. It is a projected energy.

This theorem also turns finite observation into an optimisation problem over rank-constrained subspaces. The next subsection solves that problem.

4.2 Optimal retained signal at fixed observable dimension

Fix an observable dimension d . Among all rank- d orthogonal projectors, some will retain very little of the DV signal, while others will retain as much as possible. The optimal value is given by the Ky Fan top-singular-energy envelope of A_{DV} [32, 34].

Let

$$s_1(A_{DV}) \geq s_2(A_{DV}) \geq \dots \geq s_{n-1}(A_{DV}) \geq 0 \quad (4.7)$$

denote the singular values of A_{DV} , counted with multiplicity and extended by zeros if needed. Define

$$K_d^{DV} := \sup_{\text{rank } \Pi = d} V_\Pi^{DV}. \quad (4.8)$$

Theorem 4.2 (Exact weighted Ky Fan envelope for orthogonal rank- d observers). For each $d \in \{0, 1, \dots, n-1\}$, among orthogonal rank- d projectors on the reduced Euclidean Fisher space,

$$K_d^{DV} = \frac{1}{4} \sum_{k=1}^d s_k(A_{DV})^2. \quad (4.9)$$

The supremum is attained by the orthogonal projector onto the span of the top d left singular vectors of A_{DV} . This is a global optimum over the orthogonal rank-constrained observer class fixed in the paper.

Proof. By Theorem 4.1,

$$V_\Pi^{DV} = \frac{1}{4} \|\Pi A_{DV}\|_F^2. \quad (4.10)$$

Write the singular value decomposition

$$A_{DV} = U \Sigma V^T, \quad (4.11)$$

where $\Sigma = \text{diag}(s_1, \dots, s_{n-1})$. Since left multiplication by an orthogonal matrix pre-

serves projector rank and Frobenius norm, maximising $\|\Pi A_{DV}\|_F^2$ over rank- d projectors is equivalent to maximising $\|\tilde{\Pi}\Sigma\|_F^2$, where $\tilde{\Pi} = U^T \Pi U$ is again a rank- d orthogonal projector.

The quantity $\|\tilde{\Pi}\Sigma\|_F^2$ is the sum of the squares of the singular values weighted by the diagonal entries selected by the projector. By the Ky Fan maximum principle and its modern matrix-analytic formulation [32, 34], the maximum is obtained by projecting onto the span of the first d coordinate directions, giving

$$\sup_{\text{rank } \Pi=d} \|\Pi A_{DV}\|_F^2 = \sum_{k=1}^d s_k(A_{DV})^2. \quad (4.12)$$

Multiplying by 1/4 yields the result.

Proposition 4.3 (Lipschitz stability of the envelope values). *Let Δ and Δ' be positive semidefinite corrections on the same reduced Euclidean Fisher space, and let $K_d(\Delta)$ and $K_d(\Delta')$ denote the corresponding orthogonal rank- d envelopes. Then for every d ,*

$$|K_d(\Delta) - K_d(\Delta')| \leq \sum_{k=1}^d |\lambda_k(\Delta) - \lambda_k(\Delta')| \leq d \|\Delta - \Delta'\|_{op}. \quad (4.13)$$

In particular, the envelope values are stable under operator perturbation.

Proof. By the Ky Fan maximum principle [32, 34], $K_d(\Delta) = \sum_{k=1}^d \lambda_k(\Delta)$ and similarly for Δ' . The first inequality is immediate from termwise comparison of the ordered eigenvalue lists, and the second follows from Weyl's bound $|\lambda_k(\Delta) - \lambda_k(\Delta')| \leq \|\Delta - \Delta'\|_{op}$.

Theorem 4.2 is the finite-dimensional visibility law at fixed observable rank. It says that the best possible orthogonal d -dimensional observer captures exactly the top d singular energies of the weighted signal operator.

This identifies the optimal rank- d orthogonal observer within the reduced Euclidean Fisher geometry. It is the observer aligned with the top singular front of the weighted signal operator.

Corollary 4.4 (Worst retained signal). *Let $m = n - 1$ be the dimension of the reduced Euclidean Fisher space, and extend the singular-value list of A_{DV} to length m by zeros. Then the minimum visible nonequilibrium signal over rank- d orthogonal projectors is*

$$\inf_{\text{rank } \Pi=d} V_{\Pi}^{DV} = \frac{1}{4} \sum_{k=m-d+1}^m s_k(A_{DV})^2. \quad (4.14)$$

In particular, if $\text{rank}(A_{DV}) \leq m - d$, then the minimum is zero.

Proof. Apply the Ky Fan minimum principle to the positive semidefinite operator Δ_{DV} , or equivalently to the squared singular values of A_{DV} padded by zeros to the ambient reduced dimension m . The worst rank- d projector retains the bottom d singular energies.

Corollary 4.4 will later be useful when we discuss generic observers, bad alignment, and visibility intervals.

4.3 Retention fraction and hidden fraction

The envelope determines not only the best retained amount, but also the fraction of the total signal that survives at observable dimension d .

Define the total nonequilibrium signal by

$$K_{\text{tot}}^{DV} := \text{tr}(\Delta_{DV}) = \frac{1}{4} \|A_{DV}\|_F^2. \quad (4.15)$$

For each d , define the retention fraction

$$\Phi_d^{DV} := \frac{K_d^{DV}}{K_{\text{tot}}^{DV}}, \quad (4.16)$$

whenever $A_{DV} \neq 0$. Then

$$\Phi_d^{DV} = \frac{\sum_{k=1}^d s_k(A_{DV})^2}{\sum_{k=1}^{n-1} s_k(A_{DV})^2}. \quad (4.17)$$

The complementary hidden fraction is

$$1 - \Phi_d^{DV} = \frac{\sum_{k>d} s_k(A_{DV})^2}{\sum_{k=1}^{n-1} s_k(A_{DV})^2}. \quad (4.18)$$

Corollary 4.5 (Exact retention and hidden fractions for the optimal orthogonal rank- d envelope). Assume $A_{DV} \neq 0$. Then for every observable dimension d ,

$$\Phi_d^{DV} = \frac{\sum_{k=1}^d s_k(A_{DV})^2}{\sum_{k=1}^{n-1} s_k(A_{DV})^2}, \quad 1 - \Phi_d^{DV} = \frac{\sum_{k>d} s_k(A_{DV})^2}{\sum_{k=1}^{n-1} s_k(A_{DV})^2}. \quad (4.19)$$

In particular, the hidden nonequilibrium content is the singular-energy tail of the weighted signal operator.

Proof. This is immediate from Theorem 4.2 and the definition of K_{tot}^{DV} .

Corollary 4.5 shows that finite observation of nonequilibrium is a cumulative singular-energy problem. The visible share is the cumulative front, and the hidden share is the tail. The retention fraction is the fraction of the weighted DV signal retained by the best rank- d orthogonal observer.

4.4 Weighted stable rank as the complexity invariant

The retention hierarchy of Corollary 4.5 gives rise to a natural complexity proxy for the paper.

Definition 4.6 (Weighted stable rank). *The weighted stable rank of the nonequilibrium signal is*

$$\text{sr}(A_{DV}) := \frac{\|A_{DV}\|_F^2}{\|A_{DV}\|_{op}^2} = \frac{\sum_{k=1}^{n-1} s_k(A_{DV})^2}{s_1(A_{DV})^2}, \quad (4.20)$$

for nonzero A_{DV} .

This is the stable-rank or numerical-rank functional applied to the weighted signal operator in the geometry [35]. This quantity is the effective dimensionality of the weighted nonequilibrium signal. It measures how many singular directions are genuinely carrying signal, after weighting by the detailed-balance geometry.

If $A_{DV} = 0$, then $\Delta_{DV} = 0$, every envelope value K_d^{DV} vanishes, and the finite-observation problem is trivial because there is no nonequilibrium signal to retain.

A low stable rank means that the nonequilibrium correction is concentrated. A small number of observable directions can recover a large fraction of it. A high stable rank means that the correction is diffuse and spread across many directions. In that case, finite observers are forced to miss most of the signal unless their rank is correspondingly large.

The next theorem makes this precise.

Theorem 4.7 (Stable-rank obstruction). *Assume $A_{DV} \neq 0$. Then for every d ,*

$$\Phi_d^{DV} \leq \frac{d}{\text{sr}(A_{DV})}. \quad (4.21)$$

Equivalently, to retain at least fraction η of the nonequilibrium signal, any observer must satisfy

$$d \geq \eta \text{ sr}(A_{DV}). \quad (4.22)$$

Proof. By definition of stable rank,

$$s_1(A_{DV})^2 = \frac{\sum_{k=1}^{n-1} s_k(A_{DV})^2}{\text{sr}(A_{DV})}. \quad (4.23)$$

Since the singular values are non-increasing,

$$\sum_{k=1}^d s_k(A_{DV})^2 \leq d s_1(A_{DV})^2. \quad (4.24)$$

Therefore

$$\Phi_d^{DV} = \frac{\sum_{k=1}^d s_k(A_{DV})^2}{\sum_{k=1}^{n-1} s_k(A_{DV})^2} \leq \frac{d s_1(A_{DV})^2}{\sum_{k=1}^{n-1} s_k(A_{DV})^2} = \frac{d}{\text{sr}(A_{DV})}. \quad (4.25)$$

Rearranging gives the necessary rank condition for retaining at least fraction η .

Theorem 4.7 is the scalar obstruction governing finite observation. Weighted stable rank is a useful nonequilibrium complexity proxy, but it is not a complete descriptor of the envelope hierarchy. It gives a first-order visibility bound rather than a sufficient statistic for the full retention curve.

The literature on partial observation often organises itself around estimator performance, signal-to-noise ratio, missing transitions, or lower bounds on hidden entropy production.

Those are all useful, but they do not by themselves isolate the intrinsic dimensional difficulty of the weighted nonequilibrium signal. Weighted stable rank supplies a coarse measure of that difficulty.

4.5 Sharp consequences

The envelope and stable-rank law immediately give several structural consequences in finite dimensions.

Corollary 4.8 (Observability criterion). *Assume $A_{DV} \neq 0$. Then the following are equivalent for a given observable dimension d :*

1. *there exists a rank- d orthogonal observer in the reduced Euclidean Fisher space that sees the entire nonequilibrium signal,*
2. $\Phi_d^{DV} = 1$,
3. $d \geq \text{rank}(A_{DV})$.

Proof. The equality $\Phi_d^{DV} = 1$ means that all nonzero singular energies are contained in the first d modes, which is equivalent to $d \geq \text{rank}(A_{DV})$. In that case, projection onto the span of all left singular vectors with nonzero singular value retains the entire signal.

Corollary 4.9 (Blind-observer threshold). *Let*

$$r := \text{rank}(A_{DV}) = \text{rank}(\Delta_{DV}). \quad (4.26)$$

Then a rank- d observer can be completely blind to the nonequilibrium correction, meaning

$$V_{\Pi}^{DV} = 0, \quad (4.27)$$

if and only if

$$d \leq (n - 1) - r. \quad (4.28)$$

Equivalently, a d -dimensional blind observer exists if and only if the observable sector can be placed entirely inside $\ker(\Delta_{DV})$.

Proof. By Corollary 4.4 and the visibility interval,

$$\inf_{\text{rank } \Pi = d} V_{\Pi}^{DV} = \sum_{k=(n-1)-d+1}^{n-1} \tau_k, \quad (4.29)$$

where τ_k are the eigenvalues of Δ_{DV} in nonincreasing order, padded by zeros. This minimum vanishes if and only if the bottom d eigenvalues are zero, which is equivalent to

$$d \leq \dim \ker(\Delta_{DV}) = (n - 1) - r. \quad (4.30)$$

Corollary 4.10 (One-mode visibility). Assume $A_{DV} \neq 0$. Then the maximal rank-one visible fraction is

$$\Phi_1^{DV} = \frac{1}{\text{sr}(A_{DV})}. \quad (4.31)$$

Thus the inverse stable rank is the best one-mode visibility of the nonequilibrium signal.

Proof. By definition,

$$\Phi_1^{DV} = \frac{s_1(A_{DV})^2}{\sum_{k=1}^{n-1} s_k(A_{DV})^2} = \frac{1}{\text{sr}(A_{DV})}. \quad (4.32)$$

Corollary 4.10 gives perhaps the most immediate interpretation of the weighted stable rank. It tells us how visible the nonequilibrium signal is to the best possible one-dimensional observer.

Corollary 4.11 (Envelope monotonicity and concavity). The sequence $d \mapsto K_d^{DV}$ is non-decreasing and discretely concave:

$$K_{d+1}^{DV} - K_d^{DV} \leq K_d^{DV} - K_{d-1}^{DV}. \quad (4.33)$$

Equivalently, the marginal gain from increasing the observer dimension decreases monotonically with rank.

Proof. This follows immediately from the representation

$$K_d^{DV} = \frac{1}{4} \sum_{k=1}^d s_k(A_{DV})^2, \quad (4.34)$$

since the increments are the non-increasing sequence $\frac{1}{4}s_d(A_{DV})^2$.

Corollary 4.11 has a useful interpretation. The first few observable directions are the most valuable. Beyond that, one enters the diminishing-returns regime governed by the singular-energy tail.

We also note the residual form.

Corollary 4.12 (Optimal hidden tail). The minimal hidden nonequilibrium signal after the best rank- d observation is

$$K_{\text{tot}}^{DV} - K_d^{DV} = \frac{1}{4} \sum_{k>d} s_k(A_{DV})^2. \quad (4.35)$$

Proof. Subtract the envelope from the total signal.

Corollary 4.12 shows that finite observation has an optimal truncation error. This will later become the basis of the spectroscopy theorem.

Proposition 4.13 (Equality in the stable-rank obstruction). Assume $A_{DV} \neq 0$, and let

$$a_k := s_k(A_{DV})^2, \quad a_1 \geq a_2 \geq \dots \geq 0. \quad (4.36)$$

Then

$$\Phi_d^{DV} \leq \frac{d}{\text{sr}(A_{DV})} \quad (4.37)$$

is an equality if and only if

$$a_1 = a_2 = \dots = a_d. \quad (4.38)$$

The analogous statement holds for the detectability retained fraction after replacing A_{DV} by $C_{DV} := H_0^{-1/2} \widehat{J} H_0^{-1/2}$.

Proof. Since

$$\text{sr}(A_{DV}) = \frac{\sum_{k=1}^{n-1} a_k}{a_1}, \quad (4.39)$$

the stable-rank bound is equivalent to

$$\sum_{k=1}^d a_k \leq d a_1. \quad (4.40)$$

Equality holds if and only if $a_1 = \dots = a_d$. The detectability statement is identical after replacing the raw signal operator by the backbone-whitened signal operator.

4.6 Comparison with unweighted signal geometry

A prior unweighted analysis of rank-limited visibility in a Euclidean setting did not yet include the Donsker-Varadhan bridge or the Hessian whitening. The theory contains that picture as an unweighted shadow.

The point is simple. In the final weighted theory, the physically relevant signal is

$$A_{DV} = \widehat{J} H_0^{-1/2}, \quad (4.41)$$

not $\widehat{J}$ by itself. The role of the unweighted Euclidean hierarchy is therefore to capture the same signal-retention geometry before the dissipative weighting is inserted. Once the bridge is known, the correct signal is no longer the raw skew channel but the skew channel whitened by the detailed-balance Hessian.

This relationship can be quantified directly. Let

$$\mu_{\min} := \lambda_{\min}(H_0), \quad \mu_{\max} := \lambda_{\max}(H_0). \quad (4.42)$$

Since

$$\mu_{\max}^{-1} I \preceq H_0^{-1} \preceq \mu_{\min}^{-1} I, \quad (4.43)$$

one obtains the operator inequalities

$$\frac{1}{4\mu_{\max}} \widehat{J} \widehat{J}^T \preceq \Delta_{DV} \preceq \frac{1}{4\mu_{\min}} \widehat{J} \widehat{J}^T. \quad (4.44)$$

Consequently, if $K_d(\widehat{J})$ denotes the unweighted rank- d singular-energy envelope attached to $\widehat{J}$, equivalently the Ky Fan envelope of $\frac{1}{4}\widehat{J}\widehat{J}^T$, then

$$\frac{1}{4\mu_{\max}} K_d(\widehat{J}) \leq K_d^{DV} \leq \frac{1}{4\mu_{\min}} K_d(\widehat{J}). \quad (4.45)$$

Proposition 4.14 (Comparison with the unweighted envelope). *The weighted DV envelope is sandwiched between the old unweighted skew singular-energy envelope, scaled by the spectral window of the detailed-balance Hessian:*

$$\frac{1}{4\mu_{\max}} K_d(\widehat{J}) \leq K_d^{DV} \leq \frac{1}{4\mu_{\min}} K_d(\widehat{J}). \quad (4.46)$$

Proof. From

$$A_{DV}A_{DV}^T = \widehat{J}H_0^{-1}\widehat{J}^T \quad (4.47)$$

and the positive-semidefinite order bounds on H_0^{-1} , one gets

$$\mu_{\max}^{-1} \widehat{J}\widehat{J}^T \preceq A_{DV}A_{DV}^T \preceq \mu_{\min}^{-1} \widehat{J}\widehat{J}^T. \quad (4.48)$$

Multiplying by $1/4$ gives the corresponding comparison for Δ_{DV} . Taking the optimal rank- d compressed traces on both sides yields the stated envelope bounds.

Proposition 4.14 is important for the internal logic of the theory. It shows that the unweighted Euclidean envelope is the unwhitened skeleton of the final weighted signal geometry, and it quantifies its precise relationship to the full theory.

At this point the finite observation problem has been solved at the level of retained amount, retained fraction, and complexity obstruction. The next section turns to the geometry of observer quality. Visible nonequilibrium becomes a question of alignment between the observer and the singular directions of the weighted signal operator.

5 Alignment, optimal observers, and generic observers

5.1 The alignment theorem

Section 4 solved the rank-constrained visibility problem at the level of optimal retained signal. We now refine that result by identifying the geometric mechanism behind observer quality.

Let

$$A_{DV} = \sum_{k=1}^r s_k u_k v_k^T \quad (5.1)$$

be a singular value decomposition of the weighted signal operator, where

$$s_1 \geq s_2 \geq \dots \geq s_r > 0 \quad (5.2)$$

are the nonzero singular values, the vectors u_k form an orthonormal family of left singular directions in the reduced Euclidean Fisher space, and the vectors v_k form an orthonormal family of right singular directions. The finite observation problem concerns how a rank-

constrained projector Π couples to the left singular directions $\{u_k\}$, since these are the visible directions of the positive signal operator

$$\Delta_{DV} = \frac{1}{4} A_{DV} A_{DV}^T. \quad (5.3)$$

The answer is given by the following theorem.

Theorem 5.1 (Alignment theorem for orthogonal observers). *For every orthogonal projector Π ,*

$$V_{\Pi}^{DV} = \frac{1}{4} \sum_{k=1}^r s_k (A_{DV})^2 \|\Pi u_k\|^2. \quad (5.4)$$

Equivalently, the visible nonequilibrium signal retained by Π is a weighted alignment score between the observer and the left singular directions of the weighted signal operator.

Proof. Starting from the projector formula of Theorem 4.1,

$$V_{\Pi}^{DV} = \frac{1}{4} \|\Pi A_{DV}\|_F^2. \quad (5.5)$$

Insert the singular value decomposition:

$$\Pi A_{DV} = \sum_{k=1}^r s_k (\Pi u_k) v_k^T. \quad (5.6)$$

Since the v_k form an orthonormal family, the Frobenius norm squares add:

$$\|\Pi A_{DV}\|_F^2 = \sum_{k=1}^r s_k^2 \|\Pi u_k\|^2. \quad (5.7)$$

Multiplying by 1/4 gives the claim.

Remark (Degenerate singular subspace ambiguity in the alignment theorem). If the singular spectrum of A_{DV} has degeneracies, the left singular directions are defined only up to orthogonal rotation inside each degenerate singular subspace. The theorem should therefore be read at the level of the corresponding invariant subspace. What is canonical is the invariant subspace itself, together with the corresponding weighted sum on that subspace, not any individual basis vector inside a degenerate block.

Theorem 5.1 is one of the central structural identities of the paper. It shows that visible nonequilibrium is not determined by rank alone. Rank fixes the size of the observer, but the actual retained signal is controlled by how the observer aligns with the singular-energy front of A_{DV} .

This point deserves emphasis. A poor observer need not be too small; it can be misaligned. A coarse observer can outperform a finer one if it captures a larger weighted overlap with the dominant singular directions. In the language of the partial-observation literature, observer geometry matters as much as observer size [6–8, 13, 36].

The alignment theorem also makes the finite-observation picture more precise. The singular-

energy levels $\frac{1}{4}s_k(A_{DV})^2$ measure the available weighted signal energy in each mode, while the factors $\|\Pi u_k\|^2$ measure the observer's access to those modes. The visible nonequilibrium signal is the product of these two ingredients, summed over modes.

5.2 Optimal observer design

The Ky Fan envelope theorem already identified the maximal retained signal at fixed rank. We now restate that result in the language of observer design. This is an immediate design corollary of the envelope law and the alignment formula, and it is central to the final paper.

Corollary 5.2 (Optimal orthogonal rank- d observer). *Fix an observable dimension d . Let*

$$U_d^* := \text{span}\{u_1, \dots, u_d\}, \quad (5.8)$$

where $u_1, \dots, u_d$ are the top d left singular directions of A_{DV} . Let Π_d^ be the orthogonal projector onto U_d^* . Then*

$$V_{\Pi_d^*}^{DV} = K_d^{DV}, \quad (5.9)$$

and Π_d^ is an optimal rank- d observer.*

If the singular value $s_d(A_{DV})$ is strictly larger than $s_{d+1}(A_{DV})$, then U_d^ is uniquely determined. More generally, let $\sigma_* := s_d(A_{DV})$. Every optimal orthogonal rank- d observer must contain the full direct sum of left-singular eigenspaces with singular value strictly larger than σ_* , and the remaining dimension may be chosen arbitrarily inside the eigenspace for σ_* . This is the only freedom coming from degeneracy.*

Proof. By Theorem 4.2, the optimal retained signal at rank d is

$$K_d^{DV} = \frac{1}{4} \sum_{k=1}^d s_k(A_{DV})^2. \quad (5.10)$$

For the projector Π_d^* , one has

$$\Pi_d^* u_k = \begin{cases} u_k, & 1 \leq k \leq d, \\ 0, & k > d \end{cases} \quad (5.11)$$

when there is no degeneracy across the d -th cutoff, and the alignment theorem then gives

$$V_{\Pi_d^*}^{DV} = \frac{1}{4} \sum_{k=1}^d s_k(A_{DV})^2 = K_d^{DV}. \quad (5.12)$$

The uniqueness statement is standard from the Ky Fan maximum principle [32, 34]. If the cutoff singular value is isolated, the top d -dimensional left singular subspace is unique. If there is degeneracy, every optimal d -plane must contain the full direct sum of left-singular eigenspaces with singular value strictly larger than the cutoff, and only the remaining dimension is free inside the eigenspace at the cutoff. This is exactly the equality case in the Ky Fan principle [32, 34].

Corollary 5.2 identifies the best orthogonal finite observer for nonequilibrium detection in the Gaussian fluctuation setting.

Most recent work asks how much irreversibility can be inferred from a *given* partial observation scheme [5–7, 13, 36]. The corollary answers the stronger question within the orthogonal rank-constrained class: for a fixed observational budget d , which observable sector retains the most nonequilibrium signal? The answer is at the level of optimal value, and unique at the subspace level only when the cutoff spectral gap is open.

Corollary 4.4 gives the complementary minimum retained signal at fixed rank. Finite observation therefore comes with both a best observer and a worst observer, and the spread between them is the first visibility interval for the DV correction.

Remark (Gap-dependent subspace stability). The optimal value is stable under perturbation by Proposition 4.3, but the optimal subspace is stable only when the cutoff spectral gap $s_d(A_{DV}) - s_{d+1}(A_{DV})$ is positive. In the gapless case the optimal value remains stable, while the optimal d -plane need not be unique.

5.3 Worst-case and generic observation

The best and worst rank- d observers define the two extremes of finite observation at fixed observable dimension. Between those extremes lies the generic case. We now identify the natural baseline for generic observation.

Let

$$m := n - 1 \quad (5.13)$$

be the dimension of the reduced Euclidean Fisher space. A rank- d orthogonal projector Π is called Haar-random if its range is distributed uniformly on the Grassmannian $\text{Gr}(d, m)$. Such a projector is the natural model for an unstructured observer of dimension d .

The following proposition gives the generic baseline in expectation.

Proposition 5.3 (Haar-random orthogonal-observer baseline in expectation). *Let Π be Haar-random among rank- d orthogonal projectors on the m -dimensional reduced Euclidean Fisher space. Then*

$$\mathbb{E}[\Pi] = \frac{d}{m} I, \quad (5.14)$$

and therefore

$$\mathbb{E}[V_{\Pi}^{DV}] = \frac{d}{m} K_{\text{tot}}^{DV} = \frac{d}{m} \frac{1}{4} \|A_{DV}\|_F^2. \quad (5.15)$$

Equivalently, the expected retained fraction of a random rank- d observer is exactly

$$\mathbb{E} \left[\frac{V_{\Pi}^{DV}}{K_{\text{tot}}^{DV}} \right] = \frac{d}{m}. \quad (5.16)$$

Moreover, for each singular direction u_k ,

$$\mathbb{E}[\|\Pi u_k\|^2] = \frac{d}{m}. \quad (5.17)$$

Proof. By orthogonal invariance of the Haar measure on the Grassmannian, $\mathbb{E}[\Pi]$ commutes with every orthogonal transformation of $\mathbb{R}^m$. Therefore $\mathbb{E}[\Pi] = cI$ for some scalar c . Taking

traces gives

$$cm = \text{tr}(\mathbb{E}[\Pi]) = \mathbb{E}[\text{tr}(\Pi)] = d, \quad (5.18)$$

so $c = d/m$. Hence

$$\mathbb{E}[\Pi] = \frac{d}{m}I. \quad (5.19)$$

Using the projector formula from Theorem 4.1,

$$\mathbb{E}[V_{\Pi}^{DV}] = \frac{1}{4} \mathbb{E}[\text{tr}(\Pi A_{DV} A_{DV}^T)] = \frac{1}{4} \text{tr}(\mathbb{E}[\Pi] A_{DV} A_{DV}^T). \quad (5.20)$$

Substituting $\mathbb{E}[\Pi] = \frac{d}{m}I$ gives

$$\mathbb{E}[V_{\Pi}^{DV}] = \frac{1}{4} \frac{d}{m} \text{tr}(A_{DV} A_{DV}^T) = \frac{d}{m} K_{\text{tot}}^{DV}. \quad (5.21)$$

Dividing by K_{tot}^{DV} yields the expected retained fraction. Finally, since

$$\|\Pi u_k\|^2 = u_k^T \Pi u_k, \quad (5.22)$$

one has

$$\mathbb{E}[\|\Pi u_k\|^2] = u_k^T \mathbb{E}[\Pi] u_k = \frac{d}{m} \|u_k\|^2 = \frac{d}{m}. \quad (5.23)$$

Proposition 5.3 gives the generic visibility benchmark in expectation. A random rank- d observer retains a fraction d/m on average. That number is purely geometric. It does not depend on the particular generator or the singular spectrum of A_{DV} . It is the baseline visibility of unstructured finite observation. No concentration or high-probability claim is being made here.

This baseline will be central for interpretation. If an actual observer performs only at the d/m level, then it is no better than random geometry. If it performs substantially above that level, then it is aligned with the singular-energy front of the signal. If it performs substantially below the optimum, then there is room for observer redesign.

The literature has already encountered phenomena of this type indirectly. In coarse-graining and milestone studies, some observation maps outperform others in ways that can seem paradoxical if one thinks only in terms of coarser versus finer observation [13]. The random-observer baseline clarifies the structural issue: generic observation has a fixed geometric retention level, and exceptional performance reflects alignment rather than raw resolution.

5.4 Alignment gain

The optimal retained fraction from Section 4 and the random-observer baseline from Theorem 5.3 together define the central quality measure of finite observation.

Definition 5.4 (Alignment gain). Assume $A_{DV} \neq 0$. For observable dimension d , define the alignment gain

$$\mathcal{A}_d := \frac{\Phi_d^{DV}}{d/m}. \quad (5.24)$$

Equivalently,

$$\mathcal{A}_d = \frac{m}{d} \frac{\sum_{k=1}^d s_k(A_{DV})^2}{\sum_{k=1}^m s_k(A_{DV})^2}. \quad (5.25)$$

Thus $\mathcal{A}_d$ measures how much better the optimal rank- d observer is than a generic rank- d observer. It is the multiplicative benefit of observer design over random finite observation.

The first basic bounds are immediate.

Proposition 5.5 (Alignment-gain bounds). Assume $A_{DV} \neq 0$. Then for every observable dimension d ,

$$1 \leq \mathcal{A}_d \leq \frac{m}{\text{sr}(A_{DV})}. \quad (5.26)$$

In particular, the optimal observer always performs at least as well as the random baseline, and the maximal possible advantage is controlled by the inverse weighted complexity of the signal.

Proof. For the lower bound, let

$$a_k := s_k(A_{DV})^2, \quad a_1 \geq a_2 \geq \dots \geq a_m \geq 0. \quad (5.27)$$

Then the average of the top d terms is at least the average of all m terms:

$$\frac{1}{d} \sum_{k=1}^d a_k \geq \frac{1}{m} \sum_{k=1}^m a_k. \quad (5.28)$$

Rearranging gives

$$\frac{\sum_{k=1}^d a_k}{\sum_{k=1}^m a_k} \geq \frac{d}{m}, \quad (5.29)$$

that is,

$$\Phi_d^{DV} \geq \frac{d}{m}. \quad (5.30)$$

Hence $\mathcal{A}_d \geq 1$.

For the upper bound, Theorem 4.7 gives

$$\Phi_d^{DV} \leq \frac{d}{\text{sr}(A_{DV})}. \quad (5.31)$$

Dividing by d/m yields

$$\mathcal{A}_d \leq \frac{m}{\text{sr}(A_{DV})}. \quad (5.32)$$

Proposition 5.5 shows that alignment gain is a true quality measure. It is always at least one, because design can never be worse than random choice on average. It is large when the signal is concentrated and therefore strongly rewardable by alignment. It collapses toward

one when the signal is diffuse and therefore little is gained by design over generic observation. A large post-cutoff singular-value gap therefore makes the optimal observer robust, while a flat spectrum pushes $\mathcal{A}_d$ toward one and makes observer choice less consequential.

This gives another interpretation of weighted stable rank. A high stable-rank signal is not only difficult to see. It is also difficult to exploit by observer design. There is little alignment advantage available because the signal is spread broadly. A low stable-rank signal offers both strong visibility and large design gains.

It is also useful to record the rank-one case.

Corollary 5.6 (Rank-one alignment gain). Assume $A_{DV} \neq 0$. For $d = 1$,

$$\mathcal{A}_1 = \frac{m}{\text{sr}(A_{DV})}. \quad (5.33)$$

Thus the one-dimensional design gain is the ambient reduced dimension divided by the weighted stable rank of the signal.

Proof. By Corollary 4.10,

$$\Phi_1^{DV} = \frac{1}{\text{sr}(A_{DV})}. \quad (5.34)$$

Since the random baseline at rank one is $1/m$, the claim follows immediately.

Corollary 5.6 is especially sharp. It says that the whole one-mode design problem collapses to a single ratio: ambient reduced dimension divided by weighted stable rank.

5.5 Interpretation for the literature

The results of this section explain something that has remained somewhat implicit in the recent partial-observation literature. Many papers study finite resolution, hidden transitions, unresolved events, faulty coarse graining, or partial access, and ask why some observation schemes detect irreversibility better than others [5–8, 13]. The usual language is framed in terms of better estimators, lower bounds, waiting-time information, or resolution trade-offs. Those are real and important. But they do not yet isolate the geometric core.

The theory says that the core issue is alignment with the weighted signal operator. A coarse-graining can outperform a finer observation if it is better aligned with the dominant left singular directions of A_{DV} . Conversely, increasing raw observational resolution may fail to help if the additional directions do not couple strongly to the singular front of the nonequilibrium signal. In that sense, the relevant distinction is not simply coarse versus fine. It is alignment versus misalignment.

This perspective also clarifies the meaning of generic observation. Absent design or structural knowledge, a rank- d observer retains on average only the baseline fraction d/m . Anything above that reflects nontrivial alignment. The gap between Φ_d^{DV} and d/m is therefore not just a mathematical curiosity. It quantifies the advantage of informed observation over blind compression.

The section also changes the status of the finite observation programme itself. Up to this point, one could read the theory mainly as a structural suppression law. After the alignment theorem and the random-observer baseline, that is no longer adequate. The theory has become a design theory. It tells us not only what finite observers lose, but which finite observers are

best, how much better they are than generic observers, and what geometric property of the signal makes that possible.

The next section introduces the shadow operator, which compares the visible nonequilibrium correction with the projected detailed-balance backbone and yields the Gaussian distinguishability theory.

6 The shadow operator and Gaussian distinguishability

6.1 Projected backbone and projected correction

Sections 3, 4, and 5 identified the weighted nonequilibrium signal, solved the rank-constrained visibility problem, and showed how observer quality is controlled by spectral alignment. What remains is to turn that visibility law into a statistical statement about what a finite observer can and cannot distinguish.

This is the point where the paper splits raw visibility from detectability. A direction can carry substantial visible nonequilibrium mass and yet be comparatively weak as a discriminator once the detailed-balance backbone is used as the local metric. Conversely, a direction that is not optimal for raw visible mass can become optimal after backbone normalisation. The observer enters this second problem twice: once through the projected nonequilibrium correction and once through the projected backbone. Raw visibility is therefore already measured in global detailed-balance units through the $H_0^{-1/2}$ weighting, whereas detectability renormalises after projection in the local equilibrium units of the observed sector. The operator introduced below is the exact object that measures this second question.

Fix a rank- d orthogonal projector Π on the reduced Euclidean Fisher space, and let

$$U := \text{Ran } \Pi. \quad (6.1)$$

The observed detailed-balance backbone is the compression of H_0 to U , while the observed nonequilibrium correction is the compression of Δ_{DV} to the same sector. Concretely, define

$$B_\Pi := \Pi H_0 \Pi|_U, \quad E_\Pi := \Pi \Delta_{DV} \Pi|_U. \quad (6.2)$$

The observed Donsker-Varadhan Hessian is therefore

$$H_\Pi := \Pi(H_0 + \Delta_{DV})\Pi|_U = B_\Pi + E_\Pi. \quad (6.3)$$

Since H_0 is positive definite on the full reduced space, its restriction to any nonzero subspace is positive definite. Thus B_Π is positive definite on U . Since Δ_{DV} is positive semidefinite by Theorem 3.2, its compression E_Π is positive semidefinite on U . The observed nonequilibrium law therefore differs from the observed detailed-balance law by a positive fluctuation correction on the same finite-dimensional observable sector.

This is the correct starting point for the equilibrium-shadow theorem. The question is not simply whether E_Π is small in an absolute sense. The real question is how large it is relative to the observed detailed-balance backbone B_Π . That relative object is introduced next.

6.2 The shadow operator

The observed nonequilibrium correction should be compared to the observed detailed-balance backbone in the natural backbone metric. This leads to the central operator of the section.

Definition 6.1 (Shadow operator). *The shadow operator associated with the observer Π is*

$$S_{\Pi} := B_{\Pi}^{-1/2} E_{\Pi} B_{\Pi}^{-1/2}. \quad (6.4)$$

The shadow operator is positive semidefinite on U , and it is the observable measure of nonequilibrium fluctuation strength after normalising by the detailed-balance backbone. It is therefore the correct mathematical form of the equilibrium shadow.

Proposition 6.2 (Backbone factorisation). *For every rank- d orthogonal projector Π ,*

$$H_{\Pi} = B_{\Pi}^{1/2} (I + S_{\Pi}) B_{\Pi}^{1/2}. \quad (6.5)$$

In particular, the observed nonequilibrium Hessian is a positive perturbation of the observed detailed-balance Hessian by the shadow operator.

Proof. By definition,

$$E_{\Pi} = B_{\Pi}^{1/2} S_{\Pi} B_{\Pi}^{1/2}. \quad (6.6)$$

Therefore

$$H_{\Pi} = B_{\Pi} + E_{\Pi} = B_{\Pi} + B_{\Pi}^{1/2} S_{\Pi} B_{\Pi}^{1/2} = B_{\Pi}^{1/2} (I + S_{\Pi}) B_{\Pi}^{1/2}. \quad (6.7)$$

Proposition 6.2 is the algebraic form of the shadow picture. The observed nonequilibrium fluctuation law is obtained from the observed detailed-balance backbone by inserting the positive dimensionless perturbation S_{Π} .

We next connect this to Gaussian fluctuation laws.

6.3 Exact observed Gaussian laws

Throughout Sections 6 and 8, the observed Gaussian law on U means the centred Gaussian law induced by the compressed quadratic fluctuation form on U , with precision B_{Π} in the detailed-balance case and $H_{\Pi} = B_{\Pi} + E_{\Pi}$ in the nonequilibrium case. This is the local compressed-Hessian observation model fixed by the paper. It is not, unless explicitly stated otherwise, a theorem about arbitrary nonlinear measurement marginals or full path-space observations.

At the quadratic fluctuation level, the observed detailed-balance and observed nonequilibrium laws on U are centred Gaussian measures with precisions

$$H_{DB,\Pi} = B_{\Pi}, \quad H_{NESS,\Pi} = H_{\Pi} = B_{\Pi} + E_{\Pi}. \quad (6.8)$$

Their covariance operators are therefore

$$\Sigma_{DB,\Pi} = B_{\Pi}^{-1}, \quad \Sigma_{NESS,\Pi} = H_{\Pi}^{-1}. \quad (6.9)$$

Using Proposition 6.2, the nonequilibrium covariance can be written as

$$\Sigma_{NESS,\Pi} = B_{\Pi}^{-1/2} (I + S_{\Pi})^{-1} B_{\Pi}^{-1/2}. \quad (6.10)$$

Thus every Gaussian distinguishability quantity on the observed sector is a spectral function of the shadow operator. This is the decisive statistical simplification. The observed detailed-balance backbone has been factored out completely, and the remaining difference between the two observable Gaussian laws is encoded entirely by S_{Π} .

We now write the principal divergences explicitly. Throughout, total variation is defined by

$$\text{TV}(P, Q) := \frac{1}{2} \|P - Q\|_1, \quad (6.11)$$

and the squared Hellinger distance is defined by

$$H^2(P, Q) := 1 - \int \sqrt{\frac{dP}{d\lambda} \frac{dQ}{d\lambda}} d\lambda, \quad (6.12)$$

for any common dominating measure λ .

Theorem 6.3 (Exact Gaussian comparison formulas). *Let $\lambda_1, \dots, \lambda_d$ be the eigenvalues of S_{Π} . Then the reverse Kullback-Leibler divergence is*

$$D_{KL}(\mathcal{N}_{DB,\Pi} \parallel \mathcal{N}_{NESS,\Pi}) = \frac{1}{2} \sum_{i=1}^d \left(\lambda_i - \log(1 + \lambda_i) \right), \quad (6.13)$$

and the forward Kullback-Leibler divergence is

$$D_{KL}(\mathcal{N}_{NESS,\Pi} \parallel \mathcal{N}_{DB,\Pi}) = \frac{1}{2} \sum_{i=1}^d \left(\log(1 + \lambda_i) - \frac{\lambda_i}{1 + \lambda_i} \right). \quad (6.14)$$

Moreover, the squared Hellinger distance is

$$H^2(\mathcal{N}_{DB,\Pi}, \mathcal{N}_{NESS,\Pi}) = 1 - \frac{\det(I + S_{\Pi})^{1/4}}{\det(I + \frac{1}{2} S_{\Pi})^{1/2}}. \quad (6.15)$$

These formulas are observable Gaussian comparison laws. They are naturally adjacent to trajectory-level time-reversal Kullback–Leibler comparisons in the style of Roldán–Parrondo, but they are not the same object [36].

Proof. For centred Gaussian laws with covariances Σ_1, Σ_0 ,

$$D_{KL}(\mathcal{N}(0, \Sigma_1) \parallel \mathcal{N}(0, \Sigma_0)) = \frac{1}{2} \left[\text{tr}(\Sigma_0^{-1} \Sigma_1) - d - \log \det(\Sigma_0^{-1} \Sigma_1) \right]. \quad (6.16)$$

Take $\Sigma_1 = \Sigma_{DB,\Pi} = B_{\Pi}^{-1}$ and $\Sigma_0 = \Sigma_{NESS,\Pi} = H_{\Pi}^{-1}$. Then

$$\Sigma_0^{-1} \Sigma_1 = H_{\Pi} B_{\Pi}^{-1} = B_{\Pi}^{1/2} (I + S_{\Pi}) B_{\Pi}^{-1/2}, \quad (6.17)$$

so its eigenvalues are $1 + \lambda_i$. Hence

$$D_{KL}(\mathcal{N}_{DB,\Pi} \parallel \mathcal{N}_{NESS,\Pi}) = \frac{1}{2} \sum_{i=1}^d \left((1 + \lambda_i) - 1 - \log(1 + \lambda_i) \right), \quad (6.18)$$

which simplifies to the first formula.

For the forward divergence, take $\Sigma_1 = \Sigma_{NESS,\Pi}$ and $\Sigma_0 = \Sigma_{DB,\Pi}$. Then

$$\Sigma_0^{-1} \Sigma_1 = B_{\Pi} H_{\Pi}^{-1} = B_{\Pi}^{1/2} (I + S_{\Pi})^{-1} B_{\Pi}^{-1/2}, \quad (6.19)$$

whose eigenvalues are $(1 + \lambda_i)^{-1}$. Substituting into the Gaussian KL formula yields

$$D_{KL}(\mathcal{N}_{NESS,\Pi} \parallel \mathcal{N}_{DB,\Pi}) = \frac{1}{2} \sum_{i=1}^d \left(\frac{1}{1 + \lambda_i} - 1 + \log(1 + \lambda_i) \right), \quad (6.20)$$

since

$$\frac{1}{1 + \lambda_i} - 1 = -\frac{\lambda_i}{1 + \lambda_i}. \quad (6.21)$$

This is the stated expression.

For the Hellinger distance, the standard centred-Gaussian formula gives

$$1 - H^2 = \frac{\det(\Sigma_{DB,\Pi})^{1/4} \det(\Sigma_{NESS,\Pi})^{1/4}}{\det\left(\frac{1}{2}(\Sigma_{DB,\Pi} + \Sigma_{NESS,\Pi})\right)^{1/2}}. \quad (6.22)$$

Using

$$\Sigma_{DB,\Pi} = B_{\Pi}^{-1}, \quad \Sigma_{NESS,\Pi} = B_{\Pi}^{-1/2} (I + S_{\Pi})^{-1} B_{\Pi}^{-1/2}, \quad (6.23)$$

one computes

$$\det(\Sigma_{DB,\Pi}) = \det(B_{\Pi})^{-1}, \quad \det(\Sigma_{NESS,\Pi}) = \det(B_{\Pi})^{-1} \det(I + S_{\Pi})^{-1}, \quad (6.24)$$

and

$$\frac{1}{2}(\Sigma_{DB,\Pi} + \Sigma_{NESS,\Pi}) = B_{\Pi}^{-1/2} \left(\frac{1}{2}(I + (I + S_{\Pi})^{-1}) \right) B_{\Pi}^{-1/2} = B_{\Pi}^{-1/2} (I + \frac{1}{2}S_{\Pi}) (I + S_{\Pi})^{-1} B_{\Pi}^{-1/2}. \quad (6.25)$$

Therefore

$$\det\left(\frac{1}{2}(\Sigma_{DB,\Pi} + \Sigma_{NESS,\Pi})\right) = \det(B_{\Pi})^{-1} \frac{\det(I + \frac{1}{2}S_{\Pi})}{\det(I + S_{\Pi})}, \quad (6.26)$$

so the common B_{Π} -factor cancels and one obtains

$$1 - H^2 = \frac{\det(I + S_{\Pi})^{1/4}}{\det(I + \frac{1}{2}S_{\Pi})^{1/2}}. \quad (6.27)$$

This gives the claimed expression.

Theorem 6.3 is the precise statistical content of the shadow construction. Every local centred-Gaussian distinguishability quantity between the observed detailed-balance law and the observed nonequilibrium law is an explicit function of the eigenvalues of S_{Π} .

Proposition 6.4 (Shadow spectrum as complete Gaussian control object). For a fixed observer Π , the observed Gaussian pair

$$\mathcal{N}_{DB,\Pi}, \quad \mathcal{N}_{NESS,\Pi} \quad (6.28)$$

is equivalent, under invertible linear whitening and orthogonal diagonalisation, to the product pair

$$\bigotimes_{i=1}^d \mathcal{N}(0, 1), \quad \bigotimes_{i=1}^d \mathcal{N}\left(0, \frac{1}{1 + \lambda_i}\right), \quad (6.29)$$

where $\lambda_1, \dots, \lambda_d$ are the eigenvalues of S_Π . Consequently, every invertibly invariant centred-Gaussian discrimination functional of the observed pair is determined by the shadow spectrum $\{\lambda_i\}_{i=1}^d$ alone.

Proof. Whiten by the detailed-balance covariance and diagonalise S_Π . The relation

$$\Sigma_{NESS,\Pi} = B_\Pi^{-1/2} (I + S_\Pi)^{-1} B_\Pi^{-1/2} \quad (6.30)$$

shows that, after whitening by $B_\Pi^{1/2}$, the observed pair becomes

$$\mathcal{N}(0, I), \quad \mathcal{N}\left(0, (I + S_\Pi)^{-1}\right). \quad (6.31)$$

Orthogonal diagonalisation of S_Π then yields the displayed product form. Any centred-Gaussian discrimination functional invariant under invertible changes of coordinates is preserved by these steps, so it depends only on the eigenvalues of S_Π .

This complements the current literature on partial observation, which largely provides lower bounds, estimators, or observation-dependent thermodynamic laws [5–8, 13, 19]. Here the local Gaussian visibility problem in the compressed-Hessian model is resolved in finite-dimensional form once the DV signal is written in the appropriate weighted coordinates.

6.4 Distinguishability package

The formulas of Theorem 6.3 immediately imply a collection of sharp and simple bounds. These are the estimates that turn the shadow operator into a practical observable criterion.

Proposition 6.5 (KL and total-variation bounds). For every observer Π ,

$$D_{KL}(\mathcal{N}_{DB,\Pi} \parallel \mathcal{N}_{NESS,\Pi}) \leq \frac{1}{4} \|S_\Pi\|_F^2. \quad (6.32)$$

Consequently,

$$\text{TV}(\mathcal{N}_{DB,\Pi}, \mathcal{N}_{NESS,\Pi}) \leq \frac{1}{2\sqrt{2}} \|S_\Pi\|_F. \quad (6.33)$$

For equal priors, the optimal Bayes classification error satisfies

$$P_e^* \geq \frac{1}{2} - \frac{1}{4\sqrt{2}} \|S_\Pi\|_F. \quad (6.34)$$

Proof. Consider the scalar function

$$f(\lambda) := \frac{\lambda^2}{2} - \lambda + \log(1 + \lambda), \quad \lambda \geq 0. \quad (6.35)$$

Then

$$f'(\lambda) = \lambda - 1 + \frac{1}{1 + \lambda} = \frac{\lambda^2}{1 + \lambda} \geq 0, \quad (6.36)$$

and $f(0) = 0$. Hence $f(\lambda) \geq 0$, or equivalently,

$$\lambda - \log(1 + \lambda) \leq \frac{\lambda^2}{2} \quad \text{for all } \lambda \geq 0. \quad (6.37)$$

Applying this estimate termwise in Theorem 6.3 gives

$$D_{KL}(\mathcal{N}_{DB,\Pi} \parallel \mathcal{N}_{NESS,\Pi}) \leq \frac{1}{2} \sum_{i=1}^d \frac{\lambda_i^2}{2} = \frac{1}{4} \|\mathcal{S}_\Pi\|_F^2. \quad (6.38)$$

Pinsker's inequality then yields

$$\text{TV}(\mathcal{N}_{DB,\Pi}, \mathcal{N}_{NESS,\Pi}) \leq \sqrt{\frac{1}{2} D_{KL}(\mathcal{N}_{DB,\Pi} \parallel \mathcal{N}_{NESS,\Pi})} \leq \frac{1}{2\sqrt{2}} \|\mathcal{S}_\Pi\|_F. \quad (6.39)$$

For equal priors, the optimal Bayes error satisfies

$$P_e^* = \frac{1 - \text{TV}}{2}, \quad (6.40)$$

so the total-variation bound gives the stated lower bound.

Proposition 6.5 shows that the shadow operator is not just a formal factorisation device. Its Frobenius norm directly controls statistical distinguishability at the Gaussian level.

We next bound the shadow operator in terms of the visible envelope and the observed backbone scale.

Proposition 6.6 (Backbone-normalised visibility bound). *For every observer Π ,*

$$\|\mathcal{S}_\Pi\|_{op} \leq \frac{V_\Pi^{DV}}{\lambda_{\min}(B_\Pi)}, \quad \|\mathcal{S}_\Pi\|_F \leq \frac{V_\Pi^{DV}}{\lambda_{\min}(B_\Pi)}. \quad (6.41)$$

In particular,

$$\|\mathcal{S}_\Pi\|_{op} \leq \frac{K_d^{DV}}{\lambda_{\min}(B_\Pi)}, \quad \|\mathcal{S}_\Pi\|_F \leq \frac{K_d^{DV}}{\lambda_{\min}(B_\Pi)}. \quad (6.42)$$

Proof. Since $\mathcal{S}_\Pi = B_\Pi^{-1/2} E_\Pi B_\Pi^{-1/2}$,

$$\|\mathcal{S}_\Pi\|_{op} \leq \|B_\Pi^{-1/2}\|_{op}^2 \|E_\Pi\|_{op} = \frac{\|E_\Pi\|_{op}}{\lambda_{\min}(B_\Pi)}. \quad (6.43)$$

Because E_Π is positive semidefinite,

$$\|E_\Pi\|_{op} \leq \text{tr}(E_\Pi) = V_\Pi^{DV}. \quad (6.44)$$

Thus

$$\|S_\Pi\|_{op} \leq \frac{V_\Pi^{DV}}{\lambda_{\min}(B_\Pi)}. \quad (6.45)$$

Similarly,

$$\|S_\Pi\|_F \leq \|B_\Pi^{-1/2}\|_{op}^2 \|E_\Pi\|_F \leq \frac{\|E_\Pi\|_F}{\lambda_{\min}(B_\Pi)}. \quad (6.46)$$

Again $E_\Pi \geq 0$, hence

$$\|E_\Pi\|_F \leq \text{tr}(E_\Pi) = V_\Pi^{DV}, \quad (6.47)$$

which gives

$$\|S_\Pi\|_F \leq \frac{V_\Pi^{DV}}{\lambda_{\min}(B_\Pi)}. \quad (6.48)$$

Since $V_\Pi^{DV} \leq K_d^{DV}$, the final bounds follow.

Proposition 6.6 is the missing link between the finite-observation envelope and the statistical shadow criterion. The visible envelope by itself is not enough to guarantee indistinguishability. What matters is the visible envelope relative to the observed detailed-balance backbone.

The correction is visible only through the ratio

$$\frac{K_d^{DV}}{\lambda_{\min}(B_\Pi)}. \quad (6.49)$$

This is the backbone-normalised observation budget. It is the quantity that turns retained signal into statistical closeness. This criterion is observer-specific unless one also imposes a uniform lower bound on $\lambda_{\min}(B_\Pi)$.

6.5 The equilibrium shadow criterion

We can now progress to the paper's main statistical theorem in its final form.

Theorem 6.7 (Gaussian equilibrium-shadow criterion). *Let Π be a rank- d observer and let S_Π be its shadow operator. Then the observed nonequilibrium and observed detailed-balance Gaussian laws are exactly identical if and only if $S_\Pi = 0$.*

More generally, if

$$\Theta_\Pi := \|S_\Pi\|_{op} \leq \varepsilon, \quad (6.50)$$

then

$$D_{KL}(\mathcal{N}_{DB,\Pi} \parallel \mathcal{N}_{NESS,\Pi}) \leq \frac{d \varepsilon^2}{4}, \quad (6.51)$$

$$\text{TV}(\mathcal{N}_{DB,\Pi}, \mathcal{N}_{NESS,\Pi}) \leq \frac{\sqrt{d} \varepsilon}{2\sqrt{2}}, \quad (6.52)$$

and

$$P_e^* \geq \frac{1}{2} - \frac{\sqrt{d} \varepsilon}{4\sqrt{2}}. \quad (6.53)$$

If

$$\frac{K_d^{DV}}{\lambda_{\min}(B_\Pi)} \leq \varepsilon \quad (6.54)$$

for some $\varepsilon > 0$, then the compressed Gaussian fluctuation law is quantitatively close to the corresponding detailed-balance Gaussian fluctuation law, with bounds given above in terms of ε . This is a statement about observed quadratic-Gaussian statistics, not about the full microscopic dynamics.

Proof. The identification follows from Proposition 6.2. Since B_Π is positive definite,

$$H_\Pi = B_\Pi \iff E_\Pi = 0 \iff S_\Pi = 0. \quad (6.55)$$

Now assume $\|S_\Pi\|_{op} \leq \varepsilon$. Since $S_\Pi \geq 0$, all eigenvalues satisfy $0 \leq \lambda_i \leq \varepsilon$. From Theorem 6.3,

$$D_{KL}(\mathcal{N}_{DB,\Pi} \parallel \mathcal{N}_{NESS,\Pi}) = \frac{1}{2} \sum_{i=1}^d (\lambda_i - \log(1 + \lambda_i)). \quad (6.56)$$

Using $\lambda - \log(1 + \lambda) \leq \lambda^2/2$, we obtain

$$D_{KL}(\mathcal{N}_{DB,\Pi} \parallel \mathcal{N}_{NESS,\Pi}) \leq \frac{1}{2} \sum_{i=1}^d \frac{\lambda_i^2}{2} \leq \frac{1}{2} \sum_{i=1}^d \frac{\varepsilon^2}{2} = \frac{d \varepsilon^2}{4}. \quad (6.57)$$

Pinsker then gives

$$\text{TV} \leq \sqrt{\frac{1}{2} D_{KL}} \leq \frac{\sqrt{d} \varepsilon}{2\sqrt{2}}. \quad (6.58)$$

The Bayes error bound follows as in Proposition 6.5.

Finally, Proposition 6.6 yields

$$\Theta_\Pi \leq \frac{K_d^{DV}}{\lambda_{\min}(B_\Pi)}. \quad (6.59)$$

Thus small backbone-normalised visible envelope implies small Θ_Π , and hence small divergence and small total variation.

Theorem 6.7 is the rigorous form of the equilibrium-shadow principle. Within the compressed

quadratic-Gaussian regime, finite orthogonal observation can erase visible nonequilibrium detail and leave only an equilibrium-shaped backbone. The theorem makes that precise.

There are two ingredients. Firstly, the visible nonequilibrium budget, controlled by the envelope K_d^{DV} . The second is the observed backbone stiffness, controlled by $\lambda_{\min}(B_{\Pi})$. The shadow criterion says that equilibrium-shaped observation emerges when the visible nonequilibrium budget is small relative to the stiffness of the observed detailed-balance backbone.

This corrects a point that would otherwise remain too loose. The inequality

$$d \ll \text{sr}(A_{DV}) \quad (6.60)$$

controls how much of the total nonequilibrium signal can be retained. It does not by itself imply statistical indistinguishability from detailed balance. The shadow theorem identifies the missing ingredient, namely the backbone scale.

6.6 Backbone-whitened detectability geometry

The shadow criterion identifies the correct backbone-normalised observable quantity, but it also admits a second spectral reformulation. Define the backbone-whitened operator

$$T := H_0^{-1/2} \Delta_{DV} H_0^{-1/2} = \frac{1}{4} C_{DV} C_{DV}^T, \quad C_{DV} := H_0^{-1/2} A_{DV} = H_0^{-1/2} \widehat{J} H_0^{-1/2}. \quad (6.61)$$

This operator encodes detectability relative to the detailed-balance backbone rather than raw visible signal. In general its spectral order need not coincide with the spectral order of A_{DV} , except in commuting or isotropic-backbone regimes treated later.

Theorem 6.8 (Backbone-whitened detectability equivalence). *Let $U \subset V$ be a d -dimensional observable subspace, and let $Q \in \mathbb{R}^{(n-1) \times d}$ have orthonormal columns spanning U . Define*

$$S_U := (Q^T H_0 Q)^{-1/2} Q^T \Delta_{DV} Q (Q^T H_0 Q)^{-1/2} \quad (6.62)$$

and

$$Y_U := H_0^{1/2} Q (Q^T H_0 Q)^{-1/2}. \quad (6.63)$$

Then

$$Y_U^T Y_U = I_d, \quad S_U = Y_U^T T Y_U. \quad (6.64)$$

Hence

$$\text{tr}(S_U) = \text{tr}(P_{W_U} T) = \frac{1}{4} \|P_{W_U} C_{DV}\|_F^2, \quad (6.65)$$

where $W_U := \text{Ran}(Y_U)$ and P_{W_U} is the Euclidean orthogonal projector onto W_U . Conversely, every Euclidean d -plane $W \subset V$ arises as $W = H_0^{1/2} U$ for a unique observable d -plane U .

Proof. Direct computation gives

$$Y_U^T Y_U = (Q^T H_0 Q)^{-1/2} Q^T H_0 Q (Q^T H_0 Q)^{-1/2} = I_d, \quad (6.66)$$

and

$$Y_U^T T Y_U = (Q^T H_0 Q)^{-1/2} Q^T \Delta_{DV} Q (Q^T H_0 Q)^{-1/2} = S_U. \quad (6.67)$$

For the converse, let $W \subset V$ be any Euclidean d -plane and let Y have orthonormal columns spanning W . Set $Z := H_0^{-1/2}Y$, and define

$$Q := Z(Z^T Z)^{-1/2}. \quad (6.68)$$

Then $Q^T Q = I_d$, so $U := \text{Ran}(Q)$ is an observable d -plane. Moreover,

$$Q^T H_0 Q = (Z^T Z)^{-1}, \quad H_0^{1/2} Q (Q^T H_0 Q)^{-1/2} = Y, \quad (6.69)$$

so indeed $W = W_U$.

Remark (Detectability versus raw visibility). Theorem 6.8 is a bijection between observable d -planes and whitened Euclidean d -planes. It is not a statement about arbitrary observation channels. Detectability here is defined relative to H_0 -normalised quadratic compression and should not be conflated with raw visible signal unless additional commuting or isotropic-backbone structure is imposed.

Theorem 6.8 shows that the backbone-normalised detectability problem is itself a Ky Fan problem, now for the whitened signal C_{DV} rather than the raw visible signal A_{DV} .

Corollary 6.9 (Exact detectability envelope). *Let*

$$\theta_1 \geq \theta_2 \geq \dots \geq \theta_{n-1} \geq 0 \quad (6.70)$$

be the eigenvalues of T , equivalently

$$\theta_k = \frac{1}{4} s_k(C_{DV})^2. \quad (6.71)$$

Define

$$\Theta_d := \sup_{\dim U=d} \text{tr}(S_U). \quad (6.72)$$

Then

$$\Theta_d = \sum_{k=1}^d \theta_k, \quad \inf_{\dim U=d} \text{tr}(S_U) = \sum_{k=(n-1)-d+1}^{n-1} \theta_k. \quad (6.73)$$

The maximiser is the observable subspace whose whitened image is the span of the top d eigendirections of T .

If

$$\tilde{\Phi}_d := \frac{\Theta_d}{\text{tr}(T)} = \frac{\sum_{k=1}^d s_k(C_{DV})^2}{\sum_{k=1}^{n-1} s_k(C_{DV})^2}, \quad (6.74)$$

then

$$\tilde{\Phi}_d \leq \frac{d}{\text{sr}(C_{DV})}, \quad \text{sr}(C_{DV}) = \frac{\|C_{DV}\|_F^2}{\|C_{DV}\|_{op}^2}. \quad (6.75)$$

Proof. By Theorem 6.8, every observable d -plane induces a Euclidean compression of T . The statement is therefore exactly the Ky Fan maximum principle for the positive semidefinite matrix T , together with the complement argument for the minimum and the usual stable-rank bound.

The operator T is not an arbitrary positive semidefinite matrix. Because it is generated by the whitened skew carrier C_{DV} , its positive spectrum is forced to organise into orthogonal current planes. The next proposition records this intrinsic structure and the detectability consequences needed later.

Proposition 6.10 (Canonical current-plane decomposition). *There exist pairwise orthogonal two-dimensional subspaces*

$$E_1, \dots, E_r \subset V, \quad (6.76)$$

positive numbers

$$\kappa_1 \geq \kappa_2 \geq \dots \geq \kappa_r > 0, \quad (6.77)$$

and an orthogonal decomposition

$$V = E_1 \oplus \dots \oplus E_r \oplus \ker C_{DV}, \quad (6.78)$$

such that for each j ,

$$C_{DV}|_{E_j} = \kappa_j R_j, \quad R_j^T = -R_j, \quad R_j^2 = -I_{E_j}, \quad (6.79)$$

and therefore

$$T|_{E_j} = \frac{\kappa_j^2}{4} I_{E_j}, \quad T|_{\ker C_{DV}} = 0. \quad (6.80)$$

Proof. Since

$$T = \frac{1}{4} C_{DV} C_{DV}^T = -\frac{1}{4} C_{DV}^2, \quad (6.81)$$

the operators T and C_{DV} commute. Decompose V into orthogonal eigenspaces of T :

$$V = \bigoplus_{\lambda \geq 0} E_\lambda. \quad (6.82)$$

Each E_λ is therefore C_{DV} -invariant. On E_λ with $\lambda > 0$,

$$C_{DV}^2 = -4\lambda I. \quad (6.83)$$

Hence

$$R_\lambda := \frac{1}{2\sqrt{\lambda}} C_{DV}|_{E_\lambda} \quad (6.84)$$

satisfies

$$R_\lambda^T = -R_\lambda, \quad R_\lambda^2 = -I, \quad R_\lambda^T R_\lambda = I. \quad (6.85)$$

Thus R_λ is orthogonal and skew. If $e \in E_\lambda$ is a unit vector and $f := R_\lambda e$, then

$$\langle e, f \rangle = \langle e, R_\lambda e \rangle = 0, \quad (6.86)$$

so $\text{span}\{e, f\}$ is an invariant two-plane on which $C_{DV} = 2\sqrt{\lambda} R_\lambda$. Repeating this construction yields an orthogonal decomposition of each positive eigenspace into invariant two-planes. Finally,

$$\langle v, Tv \rangle = \frac{1}{4} \|C_{DV}^T v\|^2, \quad (6.87)$$

so $\ker T = \ker C_{DV}$.

Corollary 6.11 (Paired detectability spectrum). *Let*

$$\theta_1 \geq \theta_2 \geq \dots \geq \theta_{n-1} \geq 0 \quad (6.88)$$

be the eigenvalues of T . Then every positive detectability eigenvalue has even multiplicity. More precisely,

$$\theta_{2j-1} = \theta_{2j} = \frac{\kappa_j^2}{4}, \quad j = 1, \dots, r, \quad (6.89)$$

and all remaining eigenvalues are zero. Consequently,

$$\text{rank}(\Delta_{DV}) = \text{rank}(T) = \text{rank}(C_{DV}) = 2r, \quad (6.90)$$

so the nonequilibrium correction can never have rank one.

Proof. This is immediate from Proposition 6.10 and the identity $T = \frac{1}{4} C_{DV} C_{DV}^T$.

Corollary 6.12 (Universal rank-one detectability ceiling). *Assume $\text{tr}(T) > 0$, so that the retained fractions from Corollary 6.9 are defined. Then*

$$\tilde{\Phi}_1 \leq \frac{1}{2}. \quad (6.91)$$

Equality holds if and only if $\text{rank}(\Delta_{DV}) = 2$, equivalently if and only if there is exactly one nonzero current plane.

Proof. By Corollary 6.11, the top eigenvalue is repeated:

$$\theta_1 = \theta_2. \quad (6.92)$$

Hence

$$\text{tr}(T) = \sum_{k=1}^{n-1} \theta_k \geq 2\theta_1, \quad (6.93)$$

which gives $\tilde{\Phi}_1 = \theta_1 / \text{tr}(T) \leq 1/2$. Equality holds exactly when $\theta_k = 0$ for all $k \geq 3$, that is, exactly when $\text{rank}(T) = 2$. By Corollary 6.11, this is equivalent to $\text{rank}(\Delta_{DV}) = 2$, or equivalently to the existence of exactly one nonzero current plane.

The backbone-whitened detectability spectrum is therefore constrained by skew-current geometry rather than being an arbitrary positive spectrum. In particular, the one-dimensional ceiling in the detectability problem is structural rather than model-specific.

Corollary 6.13 (Generalised eigenvalue formulation of detectability). *The nonzero eigenvalues θ_k of T are the generalised eigenvalues of the pair (Δ_{DV}, H_0) :*

$$\Delta_{DV}u = \theta H_0u. \quad (6.94)$$

The optimal detectable modes may be chosen as H_0 -orthonormal generalised eigenvectors

$$u_i^T H_0 u_j = \delta_{ij}, \quad (6.95)$$

and

$$\Theta_d = \sum_{k=1}^d \theta_k. \quad (6.96)$$

Proof. If $Tw = \theta w$, then setting $u := H_0^{-1/2}w$ gives the generalised eigenvalue relation $\Delta_{DV}u = \theta H_0u$. Conversely, any generalised eigenvector u yields a whitened eigenvector $w := H_0^{1/2}u$ of T with the same eigenvalue. The envelope identity follows from Corollary 6.9.

Corollary 6.14 (Detectability spectroscopy). *The cumulative detectability envelope recovers the detectability spectrum:*

$$\theta_d = \Theta_d - \Theta_{d-1}. \quad (6.97)$$

Hence the full backbone-normalised detectability spectrum is encoded by the hierarchy $\{\Theta_d\}_{d=0}^{n-1}$.

Proof. Immediate from

$$\Theta_d = \sum_{k=1}^d \theta_k. \quad (6.98)$$

Corollary 6.15 (Rank-one detectability and the top generalised eigenmode). *The optimal one-dimensional detectable observer solves*

$$\Theta_1 = \sup_{u \neq 0} \frac{u^T \Delta_{DV} u}{u^T H_0 u}. \quad (6.99)$$

Equivalently, Θ_1 is the top generalised eigenvalue of the pair (Δ_{DV}, H_0) , or the top eigenvalue of T . Thus

$$\Theta_1 = \theta_1 = \frac{1}{4} s_1(C_{DV})^2. \quad (6.100)$$

The optimal one-mode detectable direction is any H_0 -normalised generalised eigenvector satisfying

$$\Delta_{DV}u = \theta_1 H_0 u, \quad u^T H_0 u = 1. \quad (6.101)$$

Proof. This is the rank-one case of Corollary 6.9 and the Rayleigh–Ritz characterisation of the top generalised eigenvalue.

Theorem 6.16 (Two-envelope backbone spectroscopy). Assume

$$H_0 > 0, \quad \Delta_{DV} \geq 0, \quad (6.102)$$

and let

$$T = H_0^{-1/2} \Delta_{DV} H_0^{-1/2}. \quad (6.103)$$

Let

$$\tau_1 \geq \tau_2 \geq \dots \geq \tau_{n-1} \geq 0 \quad (6.104)$$

be the eigenvalues of Δ_{DV} , and let

$$\theta_1 \geq \theta_2 \geq \dots \geq \theta_{n-1} \geq 0 \quad (6.105)$$

be the eigenvalues of T . If

$$\mu_{\min} I \preceq H_0 \preceq \mu_{\max} I, \quad (6.106)$$

then for every i ,

$$\frac{\tau_i}{\mu_{\max}} \leq \theta_i \leq \frac{\tau_i}{\mu_{\min}}. \quad (6.107)$$

Consequently, for every d ,

$$\frac{K_d^{DV}}{\mu_{\max}} \leq \Theta_d \leq \frac{K_d^{DV}}{\mu_{\min}}. \quad (6.108)$$

Proof. The numbers θ_i are the generalised eigenvalues of the pair (Δ_{DV}, H_0) , equivalently the eigenvalues of T . By the Courant–Fischer principle,

$$\theta_i = \max_{\dim W=i} \min_{x \in W \setminus \{0\}} \frac{x^T \Delta_{DV} x}{x^T H_0 x}. \quad (6.109)$$

Since

$$\mu_{\min} \|x\|^2 \leq x^T H_0 x \leq \mu_{\max} \|x\|^2, \quad (6.110)$$

one has

$$\frac{1}{\mu_{\max}} \frac{x^T \Delta_{DV} x}{\|x\|^2} \leq \frac{x^T \Delta_{DV} x}{x^T H_0 x} \leq \frac{1}{\mu_{\min}} \frac{x^T \Delta_{DV} x}{\|x\|^2}. \quad (6.111)$$

Applying the Courant–Fischer principle again to Δ_{DV} yields

$$\frac{\tau_i}{\mu_{\max}} \leq \theta_i \leq \frac{\tau_i}{\mu_{\min}}. \quad (6.112)$$

Summing from $i = 1$ to d gives the envelope bounds.

Proposition 6.17 (Exact commuting-case backbone spectroscopy). Assume

$$[H_0, \Delta_{DV}] = 0, \quad H_0 > 0, \quad \Delta_{DV} \geq 0. \quad (6.113)$$

Then there exists an orthogonal basis in which

$$H_0 = \text{diag}(\mu_1, \dots, \mu_{n-1}), \quad \Delta_{DV} = \text{diag}(\tau_1, \dots, \tau_{n-1}), \quad (6.114)$$

with $\mu_i > 0$ and $\tau_i \geq 0$. In that basis,

$$T = H_0^{-1/2} \Delta_{DV} H_0^{-1/2} = \text{diag}\left(\frac{\tau_1}{\mu_1}, \dots, \frac{\tau_{n-1}}{\mu_{n-1}}\right). \quad (6.115)$$

Hence along each common eigenvector the corresponding generalised eigenvalue is the ratio τ_i/μ_i . Equivalently, the spectrum of T is the multiset of ratios $\{\tau_i/\mu_i\}_{i=1}^{n-1}$, with common eigenvectors of H_0 and Δ_{DV} furnishing the modes.

Proof. Real symmetric commuting matrices are simultaneously orthogonally diagonalisable. The displayed formula for T is then immediate.

Corollary 6.18 (Isotropic-backbone reduction). If H_0 is isotropic on $\text{Ran}(\Delta_{DV})$, meaning

$$H_0 = \mu I \quad \text{on } \text{Ran}(\Delta_{DV}), \quad (6.116)$$

then

$$T = \frac{1}{\mu} \Delta_{DV}. \quad (6.117)$$

Hence raw visibility and detectability choose the same eigendirections, and their spectra differ only by the constant factor $1/\mu$.

Theorem 6.16 gives the generic spectral-window sandwich between raw visibility and detectability, while Proposition 6.17 gives the exact commuting sharpened case. This commuting regime isolates the effect of the detailed-balance backbone: detectability is raw visibility reweighted by directional backbone penalties. It is therefore natural to speak of *backbone friction* when describing the difference between the raw signal spectrum and the detectability spectrum.

Corollary 6.19 (Raw-to-detectability discrimination bounds via the backbone window). Let U_d^* be an optimal detectable d -plane from Corollary 6.9, so that the shadow spectrum of $S_{U_d^*}$ is $\theta_1, \dots, \theta_d$. Then

$$D_{KL}(\mathcal{N}_{DB, U_d^*} \parallel \mathcal{N}_{NESS, U_d^*}) = \frac{1}{2} \sum_{i=1}^d (\theta_i - \log(1 + \theta_i)), \quad (6.118)$$

$$D_{KL}(\mathcal{N}_{NESS, U_d^*} \parallel \mathcal{N}_{DB, U_d^*}) = \frac{1}{2} \sum_{i=1}^d \left(\log(1 + \theta_i) - \frac{\theta_i}{1 + \theta_i} \right), \quad (6.119)$$

and

$$H^2(\mathcal{N}_{DB, U_d^*}, \mathcal{N}_{NESS, U_d^*}) = 1 - \prod_{i=1}^d \frac{(1 + \theta_i)^{1/4}}{(1 + \frac{1}{2}\theta_i)^{1/2}}. \quad (6.120)$$

If additionally $\mu_{\min} I \preceq H_0 \preceq \mu_{\max} I$, then

$$\begin{aligned} \frac{1}{2} \sum_{i=1}^d \left(\frac{\tau_i}{\mu_{\max}} - \log \left(1 + \frac{\tau_i}{\mu_{\max}} \right) \right) &\leq D_{KL}(\mathcal{N}_{DB, U_d^*} \parallel \mathcal{N}_{NESS, U_d^*}) \\ &\leq \frac{1}{2} \sum_{i=1}^d \left(\frac{\tau_i}{\mu_{\min}} - \log \left(1 + \frac{\tau_i}{\mu_{\min}} \right) \right), \end{aligned} \quad (6.121)$$

$$\begin{aligned} \frac{1}{2} \sum_{i=1}^d \left(\log \left(1 + \frac{\tau_i}{\mu_{\max}} \right) - \frac{\tau_i / \mu_{\max}}{1 + \tau_i / \mu_{\max}} \right) &\leq D_{KL}(\mathcal{N}_{NESS, U_d^*} \parallel \mathcal{N}_{DB, U_d^*}) \\ &\leq \frac{1}{2} \sum_{i=1}^d \left(\log \left(1 + \frac{\tau_i}{\mu_{\min}} \right) - \frac{\tau_i / \mu_{\min}}{1 + \tau_i / \mu_{\min}} \right), \end{aligned} \quad (6.122)$$

and

$$1 - \prod_{i=1}^d \frac{(1 + \tau_i / \mu_{\max})^{1/4}}{(1 + \frac{1}{2}\tau_i / \mu_{\max})^{1/2}} \leq H^2 \leq 1 - \prod_{i=1}^d \frac{(1 + \tau_i / \mu_{\min})^{1/4}}{(1 + \frac{1}{2}\tau_i / \mu_{\min})^{1/2}}. \quad (6.123)$$

Consequently,

$$H^2 \leq \text{TV} \leq \sqrt{\frac{1}{2} D_{KL}(\mathcal{N}_{DB, U_d^*} \parallel \mathcal{N}_{NESS, U_d^*})}, \quad P_e^* = \frac{1 - \text{TV}}{2}, \quad (6.124)$$

so the backbone window transfers raw visibility bounds directly into total-variation and Bayes-error bounds.

Proof. The exact formulas are Theorem 6.3 applied to the optimal detectable observer U_d^* , whose shadow spectrum is $\theta_1, \dots, \theta_d$ by Corollary 6.9. The eigenvalue bounds come from Theorem 6.16. The scalar functions

$$g(t) := \frac{1}{2}(t - \log(1 + t)), \quad h(t) := \frac{1}{2} \left(\log(1 + t) - \frac{t}{1 + t} \right) \quad (6.125)$$

are increasing on $[0, \infty)$, while

$$\psi(t) := \frac{(1+t)^{1/4}}{(1+\frac{1}{2}t)^{1/2}} \quad (6.126)$$

is decreasing on $[0, \infty)$. Substituting the eigenvalue sandwich into the exact formulas yields the displayed bounds, and the total-variation and Bayes relations are standard.

6.7 Noisy Gaussian observation extension

The deterministic-subspace theory solved in the main text extends at the Gaussian level to fixed linear noisy observation channels, although the resulting optimisation problem is no longer a Ky Fan problem. The deterministic projector model is recovered as the sharp-channel limit of this Gaussian observation setting.

Theorem 6.20 (Exact Gaussian noisy-channel shadow formulas). *Let the latent detailed-balance and NESS Gaussian fluctuation laws be*

$$x_{DB} \sim \mathcal{N}(0, \Sigma_0), \quad x_{NESS} \sim \mathcal{N}(0, \Sigma_1), \quad (6.127)$$

with

$$\Sigma_0 = H_0^{-1}, \quad \Sigma_1 = (H_0 + \Delta_{DV})^{-1}. \quad (6.128)$$

Let the observer act through a linear Gaussian channel

$$y = Cx + \eta, \quad \eta \sim \mathcal{N}(0, R), \quad (6.129)$$

where $C \in \mathbb{R}^{p \times (n-1)}$ and $R > 0$. Then

$$y_{DB} \sim \mathcal{N}(0, \Sigma_{0,C}), \quad y_{NESS} \sim \mathcal{N}(0, \Sigma_{1,C}), \quad (6.130)$$

with

$$\Sigma_{0,C} = C\Sigma_0 C^T + R, \quad \Sigma_{1,C} = C\Sigma_1 C^T + R. \quad (6.131)$$

Define the latent covariance defect

$$M := \Sigma_0 - \Sigma_1 \geq 0 \quad (6.132)$$

and the observed covariance defect

$$\Lambda_C := \Sigma_{0,C} - \Sigma_{1,C} = CM C^T \geq 0. \quad (6.133)$$

Define the noisy-channel shadow operator

$$\Xi_C := \Sigma_{0,C}^{-1/2} \Lambda_C \Sigma_{0,C}^{-1/2}. \quad (6.134)$$

Then

$$0 \leq \Xi_C < I, \quad \Sigma_{1,C} = \Sigma_{0,C}^{1/2} (I - \Xi_C) \Sigma_{0,C}^{1/2}. \quad (6.135)$$

If $\xi_1, \dots, \xi_p \in [0, 1)$ are the eigenvalues of Ξ_C , then

$$D_{KL}(y_{NESS} \parallel y_{DB}) = \frac{1}{2} \sum_{i=1}^p \left[-\xi_i - \log(1 - \xi_i) \right], \quad (6.136)$$

and

$$D_{KL}(y_{DB} \parallel y_{NESS}) = \frac{1}{2} \sum_{i=1}^p \left[\frac{\xi_i}{1 - \xi_i} + \log(1 - \xi_i) \right]. \quad (6.137)$$

Proof. The covariance factorisation is immediate from

$$\Sigma_{1,C} = \Sigma_{0,C} - \Lambda_C. \quad (6.138)$$

The KL formulas follow from the standard centred-Gaussian divergence formula applied to

$$\Sigma_{0,C}^{-1} \Sigma_{1,C} \sim I - \Xi_C. \quad (6.139)$$

No optimisation claim over C and R is included here. The theorem gives Gaussian comparison formulas for a fixed linear Gaussian observation channel. The bounded interval $[0, 1)$ is the

natural noisy analogue of the unbounded shadow spectrum: added observation noise forces the observed defect to remain strictly below the noisy baseline covariance.

6.8 Physical interpretation

Section 6 completes the passage from weighted signal geometry to observable statistics.

The shadow operator is the central object. It measures the visible nonequilibrium correction in the correct backbone metric. Large shadow means the observed nonequilibrium Gaussian law is visibly different from the observed detailed-balance Gaussian law. Small shadow means the observed law is close to the detailed-balance backbone even when the full system is strongly nonequilibrium microscopically.

This is the precise sense in which finite observation can hide irreversibility in the compressed quadratic-Gaussian regime. The hiding is not a metaphor. It is a spectral statement, and after backbone whitening it is governed by a second spectrum. The fixed-channel noisy extension above shows that the same Gaussian comparison structure survives beyond deterministic subspace observation, although the corresponding channel-design problem is open.

At this point the classical theorem package is now in place. The nonequilibrium correction is a weighted signal. Its visible share under rank constraint is computable. Observer quality is an alignment problem. Detectability is controlled by a second, backbone-whitened spectrum. Statistical indistinguishability from detailed balance is governed by the shadow operator. The next section shows that the envelope hierarchy also gives spectral readout of the hidden nonequilibrium structure.

7 Spectroscopy and the anatomy of hidden nonequilibrium structure

7.1 The cumulative envelope as a spectral readout

Up to this point, the observable envelope

$$K_d^{DV} = \frac{1}{4} \sum_{k=1}^d s_k (A_{DV})^2 \quad (7.1)$$

has been interpreted as the optimal retained signal at observable dimension d . That interpretation is already strong, but it is not yet the full story.

The envelope is a cumulative spectral readout of the weighted nonequilibrium signal. Once the bridge identity has identified

$$\Delta_{DV} = \frac{1}{4} A_{DV} A_{DV}^T, \quad (7.2)$$

the sequence

$$\{K_d^{DV}\}_{d=0}^{n-1} \quad (7.3)$$

contains the cumulative singular-energy content of A_{DV} mode by mode. The visible nonequilibrium signal retained by the best rank- d observer is therefore the d -mode cumulative spectrum of the signal.

This changes the status of the theory. Finite observation is not only a statement about what is lost under compression. The envelope hierarchy is also a diagnostic object: it reveals the spectral anatomy of what is hidden.

To make that precise, set

$$K_0^{DV} := 0. \quad (7.4)$$

The next theorem shows that the envelope increments recover the singular-energy profile.

Theorem 7.1 (Cumulative envelope and ordered spectral increments). *For each $d \in \{1, \dots, n-1\}$,*

$$K_d^{DV} - K_{d-1}^{DV} = \frac{1}{4} s_d(A_{DV})^2. \quad (7.5)$$

Equivalently, the ordered singular-energy profile of the weighted nonequilibrium signal is obtained by discrete differentiation of the envelope hierarchy.

Proof. By Theorem 4.2,

$$K_d^{DV} = \frac{1}{4} \sum_{k=1}^d s_k(A_{DV})^2. \quad (7.6)$$

Subtracting the corresponding expression for K_{d-1}^{DV} gives

$$K_d^{DV} - K_{d-1}^{DV} = \frac{1}{4} s_d(A_{DV})^2. \quad (7.7)$$

Theorem 7.1 is the decisive spectral statement of the paper. The observable envelope is an cumulative readout of the ordered singular-energy spectrum of the hidden nonequilibrium signal.

This is why the phrase “spectroscopy of nonequilibrium visibility” is not just rhetorical. The envelope hierarchy is a spectral measurement object. If one knows the full sequence K_d^{DV} , one knows the ordered singular-energy profile of the weighted signal.

7.2 Recovery of the singular-energy spectrum

The preceding theorem gives an immediate recovery result.

Corollary 7.2 (Recovery of the ordered singular-energy profile). *For each $d \in \{1, \dots, n-1\}$,*

$$s_d(A_{DV})^2 = 4(K_d^{DV} - K_{d-1}^{DV}). \quad (7.8)$$

Hence the full ordered singular-energy spectrum of A_{DV} is determined exactly by the envelope hierarchy.

Proof. Multiply the identity of Theorem 7.1 by 4.

Corollary 7.2 is one of the strongest hidden consequences of the finite-observation programme. The envelope hierarchy contains complete information about the ordered weighted singular spectrum of the DV signal. In particular, the envelope is not only enough to state retention

and hidden-fraction theorems. It is enough to reconstruct the ordered spectral anatomy of the signal.

This gives the theory a genuinely diagnostic character. A retained-signal curve indexed by observable dimension is often thought of as a cumulative performance measure. Here it is much more. It is a complete cumulative encoding of the signal spectrum itself. What it does *not* recover is the singular-direction frame, still less the microscopic generator.

It is convenient to record the equivalent statement directly at the level of the positive signal operator Δ_{DV} .

Corollary 7.3 (Recovery of the nonzero spectrum of the DV correction). *The nonzero eigenvalues of Δ_{DV} are*

$$\frac{1}{4} s_k(A_{DV})^2, \quad 1 \leq k \leq \text{rank}(A_{DV}), \quad (7.9)$$

and are recovered from the envelope increments:

$$\lambda_k(\Delta_{DV}) = K_k^{DV} - K_{k-1}^{DV}. \quad (7.10)$$

Proof. Since

$$\Delta_{DV} = \frac{1}{4} A_{DV} A_{DV}^T, \quad (7.11)$$

the nonzero eigenvalues of Δ_{DV} are $\frac{1}{4} s_k(A_{DV})^2$. The result then follows from Corollary 7.2.

Corollary 7.3 shows that the envelope hierarchy may equally well be read as the cumulative eigenvalue profile of the positive signal operator Δ_{DV} . The signal can therefore be viewed either at the singular-value level through A_{DV} or at the eigenvalue level through Δ_{DV} . The two viewpoints are equivalent.

7.3 Complexity measures from the envelope

Once the envelope hierarchy is known, several structural invariants of the nonequilibrium signal become immediately accessible. The first is the total signal strength:

$$K_{\text{tot}}^{DV} = K_{n-1}^{DV}. \quad (7.12)$$

The second is the top singular energy:

$$K_1^{DV}. \quad (7.13)$$

Together these determine the weighted stable rank.

Corollary 7.4 (Recovery of weighted stable rank). *If $A_{DV} \neq 0$, then*

$$\text{sr}(A_{DV}) = \frac{K_{n-1}^{DV}}{K_1^{DV}}. \quad (7.14)$$

Proof. By definition,

$$K_{n-1}^{DV} = \frac{1}{4} \sum_{k=1}^{n-1} s_k(A_{DV})^2 = \frac{1}{4} \|A_{DV}\|_F^2, \quad (7.15)$$

while

$$K_1^{DV} = \frac{1}{4} s_1(A_{DV})^2 = \frac{1}{4} \|A_{DV}\|_{op}^2. \quad (7.16)$$

Therefore

$$\frac{K_{n-1}^{DV}}{K_1^{DV}} = \frac{\|A_{DV}\|_F^2}{\|A_{DV}\|_{op}^2} = \text{sr}(A_{DV}). \quad (7.17)$$

The weighted stable rank was introduced in Definition 4.6 of Section 4 as the effective dimensionality of the nonequilibrium signal. Corollary 7.4 shows that it is also an envelope observable. It is not an auxiliary definition imposed from outside. It is already encoded in the retention hierarchy.

The envelope also determines the signal rank.

Corollary 7.5 (Recovery of signal rank). *The rank of the weighted signal operator is the number of strictly positive envelope increments:*

$$\text{rank}(A_{DV}) = \#\{d \in \{1, \dots, n-1\} : K_d^{DV} > K_{d-1}^{DV}\}. \quad (7.18)$$

Equivalently,

$$\text{rank}(\Delta_{DV}) = \#\{d \in \{1, \dots, n-1\} : K_d^{DV} > K_{d-1}^{DV}\}. \quad (7.19)$$

Proof. By Corollary 7.2,

$$K_d^{DV} > K_{d-1}^{DV} \iff s_d(A_{DV}) > 0. \quad (7.20)$$

The number of positive singular values is the rank of A_{DV} . Since $\Delta_{DV} = \frac{1}{4} A_{DV} A_{DV}^T$, the ranks coincide.

More generally, every unitarily invariant spectral quantity of A_{DV} , and every spectral quantity of Δ_{DV} , is determined by the envelope sequence. This includes Schatten norms, nuclear norm, operator norm, Frobenius norm, spectral entropy, and any other symmetric functional of the singular-energy profile.

Proposition 7.6 (Spectral completeness of the envelope). *The sequence $\{K_d^{DV}\}_{d=0}^{n-1}$ determines the full singular-value list of A_{DV} , and hence determines every spectral functional of A_{DV} and every spectral functional of Δ_{DV} .*

Proof. Corollary 7.2 recovers every singular value squared from the successive envelope increments. Any spectral functional is a symmetric function of the singular values, so it is determined by the recovered list. Since Δ_{DV} has nonzero spectrum $\frac{1}{4} s_k(A_{DV})^2$, its spectral functionals are also determined.

Proposition 7.6 is the most compact way to state the diagnostic strength of the theory. The envelope hierarchy determines the eigenvalue profile of Δ_{DV} , equivalently the singular-energy

profile $\frac{1}{4}s_k(A_{DV})^2$.

7.4 Observability order between systems

The envelope hierarchy now suggests a natural way to compare nonequilibrium systems by finite-observation visibility.

Let $Q^{(1)}$ and $Q^{(2)}$ be two finite-dimensional irreducible Markov nonequilibrium generators, with corresponding weighted signal operators $A_{DV}^{(1)}$ and $A_{DV}^{(2)}$, and envelope hierarchies $K_d^{(1)}$ and $K_d^{(2)}$. When the reduced dimensions differ, the shorter singular-energy list is understood to be zero-padded before comparison, so the quantifier "for every d " is unambiguous.

Definition 7.7 (Observability dominance). We say that $Q^{(1)}$ dominates $Q^{(2)}$ in the strong observability order if

$$K_d^{(1)} \geq K_d^{(2)} \quad \text{for every } d. \quad (7.21)$$

We say that $Q^{(1)}$ dominates $Q^{(2)}$ in the normalised observability order if

$$\Phi_d^{(1)} \geq \Phi_d^{(2)} \quad \text{for every } d. \quad (7.22)$$

The strong order compares absolute visible signal budgets. The normalised order compares how concentrated the signal is independently of overall scale.

The next proposition identifies the underlying spectral meaning of these two orders.

Proposition 7.8 (Spectral meaning of observability dominance). The following are equivalent:

1. $Q^{(1)}$ dominates $Q^{(2)}$ in the strong observability order,
2. for every d ,

$$\sum_{k=1}^d s_k(A_{DV}^{(1)})^2 \geq \sum_{k=1}^d s_k(A_{DV}^{(2)})^2. \quad (7.23)$$

In particular, strong observability dominance is weak majorisation of the squared singular-value vectors, equivalently dominance of the Ky Fan partial sums.

Similarly, if the total signals agree,

$$K_{\text{tot}}^{(1)} = K_{\text{tot}}^{(2)}, \quad (7.24)$$

then normalised observability dominance is equivalent to majorisation of the normalised singular-energy profiles.

Proof. The strong statement is immediate from the definition of the envelope:

$$K_d^{(a)} = \frac{1}{4} \sum_{k=1}^d s_k(A_{DV}^{(a)})^2, \quad a \in \{1, 2\}. \quad (7.25)$$

The factor $1/4$ is common to both systems and therefore irrelevant to the order.

For the normalised statement, divide both envelope hierarchies by the common total signal. The resulting cumulative fractions are the retention fractions $\Phi_d^{(a)}$.

Proposition 7.8 gives a clean comparative language that the existing literature lacks. Two systems can have the same overall nonequilibrium strength but very different visibility profiles under finite observation. One may be more observable at every observational budget because its signal is more concentrated in the leading modes. The other may be thermodynamically active but spectrally diffuse and therefore harder to access.

This is a central consequence of the theory. The finite-observation programme ranks systems not only by how much nonequilibrium they contain, but by how that nonequilibrium is distributed across observable directions.

7.5 Practical meaning

Section 7 changes the status of the envelope hierarchy in three ways.

First, it shows that the envelope is the cumulative spectral profile of the weighted nonequilibrium signal. The paper is therefore not only about how much finite observers miss, but about the spectral anatomy of what they miss.

Second, it shows that weighted stable rank is not simply a convenient scalar bound. It is an envelope-derived complexity invariant. The same retention hierarchy that determines the optimal visible fraction also determines the effective dimensionality of the signal. The theory is therefore internally closed: the object that controls visibility is extracted from the same spectral data that governs retention.

Third, it provides a new comparative language for nonequilibrium systems. Systems can now be compared by observability profile, not just by total dissipation or total correction size. A system with a highly concentrated signal may be easy to diagnose under severe observational constraints, while a system with a diffuse signal may remain effectively hidden even if its total nonequilibrium strength is large.

It reveals the diagnostic content of the theory. The envelope hierarchy is a spectroscopy of hidden nonequilibrium structure. The backbone-whitened detectability operator reveals a second spectrum, governing observer design against the detailed-balance backbone. Noisy observation channels admit a Gaussian shadow extension, even though the optimal channel-design problem remains open.

That shift has real consequences for the rest of the paper. Suppression and classification under structural assumptions, rank-constrained retention mechanisms, and the weighted signal operator are all now unified in one argument. Section 7 shows that once these ingredients are combined, one gets a cumulative spectral theory of finite nonequilibrium visibility.

The next section uses this spectral viewpoint to reinterpret suppression and shared observed-Gaussian behaviour under structural assumptions as corollaries of the deeper signal theory. That is where the classical part of the paper closes.

8 Projection suppression and shared observed-Gaussian behaviour

8.1 Projection suppression revisited

The preceding sections show that the finite-observation problem is governed by two objects: the visible envelope K_d^{DV} , which controls how much of the nonequilibrium correction can survive rank- d observation, and the shadow operator S_Π , which compares that visible correction to the projected detailed-balance backbone. This recasts structural projection suppression in a cleaner and more structural way.

A suppression theorem can be established under an explicit family of structural assumptions on the slow and fast sectors, together with a classification of observable normal forms once the projected nonequilibrium correction becomes asymptotically negligible. The key point of the work is that those assumptions are no longer the primitive centre of the theory. They are one sufficient route to a more basic conclusion.

The new primitive statement is this: if the visible envelope is small relative to the observed detailed-balance backbone, then the observed nonequilibrium Gaussian law is close to the observed detailed-balance law. In the notation of Section 6, suppression means that

$$\frac{K_d^{DV}}{\lambda_{\min}(B_\Pi)} \quad (8.1)$$

is small. Everything else is a mechanism for forcing this ratio to be small.

It is therefore useful to state the relationship to the structural suppression result explicitly.

Corollary 8.1 (Gaussian projection suppression criterion). *Let Π be a rank- d observable projector. Suppose that a family of generators $Q^{(\alpha)}$ satisfies*

$$\frac{K_d^{DV}(Q^{(\alpha)})}{\lambda_{\min}(B_\Pi^{(\alpha)})} \rightarrow 0 \quad \text{as } \alpha \rightarrow \infty. \quad (8.2)$$

Then the compressed Gaussian fluctuation laws converge, in Kullback-Leibler divergence, total variation, and Hellinger distance, to the corresponding compressed detailed-balance Gaussian laws. In particular, the projected Donsker-Varadhan Hessian becomes asymptotically indistinguishable from the projected detailed-balance backbone at the observed Gaussian level.

Proof. This follows immediately from Theorem 6.7 and Proposition 6.6. The shadow operator satisfies

$$\|S_\Pi^{(\alpha)}\|_{op} \leq \frac{K_d^{DV}(Q^{(\alpha)})}{\lambda_{\min}(B_\Pi^{(\alpha)})}, \quad (8.3)$$

so the hypothesis forces $\|S_\Pi^{(\alpha)}\|_{op} \rightarrow 0$. The Gaussian divergence bounds of Section 6 then imply convergence of the observed nonequilibrium laws to the observed detailed-balance laws.

Corollary 8.1 is the precise way in which the structural suppression result is absorbed. The suppression result does not disappear. It is strengthened and reinterpreted. A structural list of hypotheses is now seen as one way of forcing a simpler and more primitive visible-envelope

ratio to be small.

8.2 The role of the structural assumptions

A suppression analysis based on five structural assumptions governing the slow-fast geometry of the generator proves asymptotic suppression of the projected nonequilibrium correction. From this viewpoint, those assumptions should be read as engineering conditions that drive the backbone-normalised visible envelope to zero.

This is clearest at the block level. By Theorem 3.5, the observed nonequilibrium correction admits, relative to the chosen H_0 -adapted block decomposition, an intrinsic slow contribution and a hidden-return contribution from the complementary sector. The structural assumptions act on these two blocks in precisely the way now suggested by the finite-observation theory.

First, conditions enforcing small slow-skew action reduce the intrinsic slow contribution. Second, conditions forcing delocalisation or decay of the mixed channel reduce the hidden return contribution. Third, fast-block coercivity penalises any remaining hidden return by weighting it against a stiff fast detailed-balance backbone. Together, these mechanisms drive the visible correction downward.

These assumptions are therefore not arbitrary. They are sufficient conditions for keeping the visible nonequilibrium signal small in comparison with the observed detailed-balance backbone. From this viewpoint, they do not define the suppression phenomenon. They instantiate it.

This reinterpretation is important because it changes the logical status of the suppression result. In a purely structural formulation, the assumptions must be stated as primitive hypotheses because the signal form of the nonequilibrium correction is not yet available. Once the bridge is closed and the weighted envelope is known, the assumptions take their proper place as family-level mechanisms forcing the basic spectral ratio to be small.

One can therefore summarise the relationship as follows:

The structural suppression theorem says that certain slow-fast hypotheses imply projection suppression.

The present theory says that projection suppression occurs whenever the visible envelope is small relative to the observed backbone, and that those hypotheses are one way of ensuring exactly that.

That is a genuine simplification. It compresses a list of family-dependent conditions into a single visibility criterion with geometric and statistical meaning.

8.3 Shared observed-Gaussian behaviour

Once the projected nonequilibrium correction becomes negligible, observable behaviour collapses to the projected normal form $(U, g_{\Pi}, G_{\Pi}, J_{\Pi})$. We now reveal why that collapse occurs.

The reason is spectral. The nonequilibrium correction is a positive weighted correction operator. Finite observation can only retain the singular-energy front of that signal. When the visible envelope is small relative to the observed backbone, the observer sees only the

projected metric and projected generator sectors to leading order. The finer nonequilibrium structure may persist microscopically, but it is spectrally invisible in the observed Gaussian law.

This yields the following shared-normal-form statement.

Corollary 8.2 (Asymptotic shared observed-Gaussian behaviour under a common projected normal form). *Fix an observable rank d . Let $\{Q_\alpha^{(1)}\}$ and $\{Q_\alpha^{(2)}\}$ be two families of finite-dimensional irreducible Markov generators with the same projected normal form for each α :*

$$N_\Pi(Q_\alpha^{(1)}) = N_\Pi(Q_\alpha^{(2)}). \quad (8.4)$$

Let $B_{\Pi,\alpha}$ denote the common projected detailed-balance backbone determined by that shared normal form, and suppose that for $a \in \{1, 2\}$,

$$\frac{K_d^{DV}(Q_\alpha^{(a)})}{\lambda_{\min}(B_{\Pi,\alpha})} \rightarrow 0 \quad \text{as } \alpha \rightarrow \infty. \quad (8.5)$$

Then each observed Gaussian fluctuation law converges, in Kullback-Leibler divergence, total variation, and Hellinger distance, to the common projected detailed-balance Gaussian law determined by $B_{\Pi,\alpha}$. Consequently, the two observed Gaussian laws converge to one another in total variation and Hellinger distance.

In particular, the observed Gaussian equivalence class is governed asymptotically by the projected normal form once the visible nonequilibrium correction becomes negligible relative to the projected backbone.

Proof. Since the projected normal forms coincide, the projected detailed-balance backbones agree and are equal to $B_{\Pi,\alpha}$. For each family, Corollary 8.1 implies convergence of the observed nonequilibrium Gaussian law to the common detailed-balance Gaussian law because

$$\|S_{\Pi,\alpha}^{(a)}\|_{op} \leq \frac{K_d^{DV}(Q_\alpha^{(a)})}{\lambda_{\min}(B_{\Pi,\alpha})} \rightarrow 0. \quad (8.6)$$

The pairwise total-variation and Hellinger conclusions then follow from the triangle inequality for those distances. We do not invoke a triangle inequality for Kullback-Leibler divergence.

Corollary 8.2 gives the final version of this shared-normal-form idea. Different generators can collapse onto the same observed Gaussian class once the projected correction is suppressed relative to the projected backbone. The paper explains the mechanism: finite observation removes most of the weighted nonequilibrium signal, and once the visible remainder is small compared with the backbone, only the projected normal form survives statistically.

In this sense, the shared behaviour is asymptotic and Gaussian. It is a consequence of a spectral bottleneck on nonequilibrium visibility.

8.4 Internal hierarchy of observational equivalence

The internal hierarchy of projected observables now sits naturally inside the final theorem package. It is useful to restate it compactly.

At the coarsest level there is *static equivalence*: two systems agree on the observed metric

structure of the accessible sector. This is the level of U and g_{Π} .

At the next level there is *dissipative equivalence*: two systems share the same projected dissipative backbone G_{Π} , and therefore the same observed detailed-balance Hessian B_{Π} .

At the finest level there is *projected normal-form equivalence*: two systems share the same projected triple $(g_{\Pi}, G_{\Pi}, J_{\Pi})$, and therefore the same observable linear normal form in the absence of visible nonequilibrium correction.

The theory inserts one more statement into this hierarchy. The shadow operator tells us when the observed nonequilibrium correction is too small to disturb this projected hierarchy at the observed Gaussian level. Thus the hierarchy is now:

1. static observable structure,
2. projected dissipative backbone,
3. projected normal form,
4. projected nonequilibrium shadow.

The classification theorem can now be read as saying that once the fourth layer is spectrally negligible, the first three determine the observed class completely.

This is a substantial conceptual improvement over an isolated taxonomy of projected data. The hierarchy is no longer just a nested classification of projected observables. It is the hierarchy of what survives once finite observation has truncated the weighted nonequilibrium signal.

8.5 Parameter compression

The full observable normal form on a rank- d sector is determined by a finite parameter count, leading to a strong compression of the observable description. That compression result also has a more structural meaning.

The visible nonequilibrium correction is not an arbitrary additional object of the same formal size as the backbone. It is a weighted signal that finite observation can only access through a rank-constrained singular-energy front. Once that visible front is negligible relative to the projected backbone, the observer is left with the projected normal form alone, with its already-known compressed parameter count.

Thus the parameter compression theorem can now be read as the endpoint of a two-step reduction:

1. the weighted nonequilibrium signal is spectrally thinned by finite observation,
2. the remaining observable law collapses onto the projected normal form with its finite parameter budget.

That is the sharper conceptual picture. The finite parameter count is not just a convenient descriptive simplification. It is what is left after the visible nonequilibrium sector has been spectrally exhausted.

9 Abstract weighted geometry and operator-space transfer

9.1 Abstract metric setting

The classical theory developed so far lives on the reduced Euclidean Fisher space, where the detailed-balance Hessian has already been used to whiten the nonequilibrium correction into the signal operator A_{DV} . The next step is to recognise that the linear algebra of visible envelopes, retention fractions, and stable-rank control is not inherently classical. It belongs to a more general weighted Hilbert-space framework.

Let V be a finite-dimensional real or complex vector space equipped with a positive-definite metric operator M , so that

$$\langle x, y \rangle_M := \langle x, My \rangle \quad (9.1)$$

defines an inner product. We denote the corresponding norm by $\|x\|_M$. An operator Π on V is called M -orthogonal if it is idempotent and self-adjoint with respect to $\langle \cdot, \cdot \rangle_M$. Equivalently,

$$\Pi^2 = \Pi, \quad \Pi^\# = \Pi, \quad (9.2)$$

where $T^\# := M^{-1}T^\dagger M$ is the M -adjoint.

Suppose now that a positive semidefinite correction operator Δ on V is represented in the weighted geometry through a symmetrised signal factorisation

$$\Delta = BB^\# \quad (9.3)$$

for some operator B . This is the abstract analogue of the classical identity

$$\Delta_{DV} = \frac{1}{4} A_{DV} A_{DV}^T \quad (9.4)$$

on the reduced Euclidean Fisher space.

The key observation is that the visible-envelope theory depends only on this weighted signal factorisation and the geometry of M -orthogonal compression. Once the metric is Euclideanised, the whole classical Ky Fan theory transfers directly.

9.2 Weighted envelopes and retention in the abstract setting

Let Π be an M -orthogonal projector of rank d . The natural visible signal retained by Π is

$$V_\Pi(\Delta) := \text{tr}(\Pi \Delta \Pi), \quad (9.5)$$

where the trace is the algebraic trace on the finite-dimensional space and the Euclideanisation is used only to express the projected Frobenius energy. Because $\Delta = BB^\#$, this quantity admits the same projected-energy interpretation as in the classical case.

To state the transfer theorem cleanly, choose the positive-definite self-adjoint square root $M^{1/2}$, and define the Euclideanised signal operator

$$\widetilde{B} := M^{1/2} B M^{-1/2}. \quad (9.6)$$

Likewise, if Π is M -orthogonal, define its Euclideanised projector

$$\widetilde{\Pi} := M^{1/2} \Pi M^{-1/2}. \quad (9.7)$$

Then $\widetilde{\Pi}$ is an ordinary Euclidean orthogonal projector of the same rank, and the visible signal may be computed entirely in Euclidean form. Because M is positive definite, this similarity map is invertible and preserves rank, so rank- d M -orthogonal observers correspond exactly to rank- d Euclidean orthogonal projectors. This is the content of the operator-space transfer theorem proved in this section.

Theorem 9.1 (Abstract weighted envelope theorem). *Let V be a finite-dimensional metric space with positive-definite metric M , and let $\Delta = BB^\sharp$ be a positive semidefinite correction operator on V . Define*

$$\widetilde{B} := M^{1/2} B M^{-1/2}, \quad \widetilde{\Pi} := M^{1/2} \Pi M^{-1/2}. \quad (9.8)$$

Then for every M -orthogonal projector Π ,

$$V_\Pi(\Delta) = \text{tr}(\Pi \Delta \Pi) = \|\widetilde{\Pi} \widetilde{B}\|_F^2. \quad (9.9)$$

Consequently, the optimal rank- d visible signal is given by the Ky Fan top-singular-energy envelope of $\widetilde{B}$, the retention fraction is the corresponding cumulative singular-energy ratio, and the effective complexity invariant is the stable rank of $\widetilde{B}$.

Proof. Conjugation by $M^{1/2}$ sends M -orthogonal projectors to Euclidean orthogonal projectors. Moreover,

$$M^{1/2} \Delta M^{-1/2} = M^{1/2} B B^\sharp M^{-1/2} = M^{1/2} B M^{-1} B^\dagger M^{1/2} = \widetilde{B} \widetilde{B}^\dagger, \quad (9.10)$$

so the weighted signal factorisation becomes an ordinary Euclidean Gram form. Therefore

$$\text{tr}(\Pi \Delta \Pi) = \text{tr}(\widetilde{\Pi} \widetilde{B} \widetilde{B}^\dagger \widetilde{\Pi}) = \|\widetilde{\Pi} \widetilde{B}\|_F^2, \quad (9.11)$$

and the Euclidean projector formula, Ky Fan envelope theorem, retention hierarchy, and stable-rank obstruction apply verbatim.

Proposition 9.2 (Factorisation-gauge invariance). *Suppose $\Delta = B B^\sharp = \widetilde{B} \widetilde{B}^\sharp$ with $\widetilde{B} = B U$ for an M -unitary operator U satisfying $U^\sharp U = I$. Then the visible signal, envelope hierarchy, retention fractions, stable rank, and ordered singular-energy profile obtained after Euclideanisation are identical for B and $\widetilde{B}$.*

Proof. Euclideanisation sends U to an ordinary unitary factor on the right of $\widetilde{B}$. Right multiplication by a unitary does not change singular values, and all stated quantities depend only on those singular values.

This shows that the visibility spectrum depends only on the correction operator Δ itself, not on the particular square-root factor used to represent its Gram structure.

The content of Theorem 9.1 is not merely a change of coordinates: Euclideanisation preserves the rank-constrained trace objective and transports the Gram correction into ordinary Euclidean form, so the weighted visibility problem becomes exactly a Ky Fan singular-value problem.

Theorem 9.1 is important for two reasons. First, it shows that the visible-envelope machinery is not a curiosity of the classical Fisher reduction. It is a theorem of weighted signal geometry. Second, it makes the quantum continuation of the paper almost automatic at the level of weighted-geometry transfer. Once the physically relevant quantum correction object is identified in a BKM-type metric geometry, the same envelope and retention theory follows by Euclideanisation.

9.3 Classical Fisher geometry as a special case

The finite-dimensional Markov theory developed in Sections 2 to 8 is the first special case of the abstract metric framework.

Indeed, after passing to the reduced Euclidean Fisher coordinates and whitening by the detailed-balance Hessian, the correction takes the ordinary Gram form

$$\Delta_{DV} = \frac{1}{4} A_{DV} A_{DV}^T. \quad (9.12)$$

This is a special case of the abstract factorisation $\Delta = BB^\sharp$, with the metric already Euclideanised and the signal operator chosen as

$$B = \frac{1}{2} A_{DV}, \quad (9.13)$$

so that $BB^\sharp = \frac{1}{4} A_{DV} A_{DV}^T = \Delta_{DV}$.

Thus every theorem of the earlier sections can be re-read as an instance of the abstract weighted geometry theorem, specialised to the finite-dimensional Markov nonequilibrium setting. The classical results are therefore not separate from the theory. They are its first concrete realisation.

This point clarifies what is fundamental in the earlier sections. The bridge theorem is classical and model-specific because it identifies the signal operator A_{DV} . Once that signal is identified, the visible-envelope, retention, alignment, and spectroscopy machinery is generic weighted linear algebra.

9.4 BKM operator geometry as a special case

The second special case is the BKM-symmetrised operator geometry. There the role of the ambient metric is played by the Bogoliubov-Kubo-Mori inner product, or more generally the corresponding positive-definite metric matrix on the finite-dimensional operator space under consideration.

In that setting, projector geometry is no longer Euclidean in the naive coordinate basis. However, after Euclideanisation by the metric square root, M -orthogonal projectors become ordinary Euclidean projectors, and the weighted signal map becomes an ordinary matrix to which the same singular-value theory applies.

This means that once a linear quantum fluctuation correction is identified in BKM geometry, the following objects transfer:

1. the visible signal retained by a rank-constrained observable sector,
2. the optimal rank- d envelope,

3. the retention fraction and hidden fraction,
4. the stable-rank complexity invariant,
5. the alignment theorem and random-observer baseline,
6. the spectroscopy hierarchy.

This is a much stronger statement than a vague analogy between classical Fisher geometry and quantum BKM geometry. It is a transfer of the projector-level visibility law, and it should be read at the level of conditional visibility geometry rather than as a complete shadow or detectability transfer.

It also explains why the operator-space step matters even before the physical quantum correction object is identified in full generality. Once such an object is written in weighted Gram form, the whole finite-observation geometry transfers unchanged.

9.5 What transfers and what does not

The abstract metric theorem makes the scope of the operator-space extension very clear.

What transfers is the weighted linear algebra of visibility. If a correction operator in a finite-dimensional metric geometry admits a weighted signal factorisation, then the projector formula, Ky Fan envelope, retention hierarchy, stable-rank invariant, alignment theorem, and spectroscopy corollary all transfer by Euclideanisation. The backbone-whitened detectability geometry of Section 6 does not yet transfer, because no identified quantum correction operator has yet been shown to support the corresponding shadow construction.

What does not automatically transfer is the physical identification of the correction operator itself. In the classical Markov setting, the bridge theorem provides this identification through the reduced-coordinate DV Hessian analysis. In the quantum setting, one still needs a physically justified linear fluctuation object in the appropriate BKM geometry before the abstract weighted theory can be applied. At the algebraic level, the missing step is a BKM-antisymmetry or completion-of-the-square cancellation for the quantum skew channel, which is not automatic for the commutator map in BKM geometry.

This distinction is essential and must remain visible in the final paper. The work proves a classical bridge and an abstract weighted geometry theorem. It then uses the latter to formulate a disciplined linear quantum extension. It does not claim a fully general nonlinear quantum Donsker-Varadhan theorem.

It shows that the classical and quantum parts of the paper are related in the right order. The classical theory is exact because both the physical bridge and the visibility geometry are known. The operator-space part is exact at the geometric level and conditional at the physical fluctuation-identification level.

The next section turns to the linear quantum fluctuation extension itself. It shows how the weighted visibility law fits naturally into recent work on BKM geometry, restricted information, and quantum nonequilibrium thermodynamics.

10 Linear quantum fluctuation extension

10.1 Quantum fluctuation object in BKM coordinates

The abstract weighted geometry of Section 9 shows that the finite-observation machinery transfers once a physically meaningful correction object is identified in a finite-dimensional metric geometry. In the quantum nonequilibrium setting, the natural candidate metric is the Bogoliubov-Kubo-Mori metric, which already appears in recent work on slow quantum nonequilibrium steady-state thermodynamics and geometric dissipation [21–23]. BKM is used here because of its fluctuation and linear-response relevance, not because it is the unique monotone quantum Fisher metric in the Petz family; different monotone choices would induce different weighted observer geometries [33].

The guiding idea is the following. In the classical Markov setting, Section 3 identifies a fluctuation correction object whose visible share under finite-rank observation is governed by a weighted signal operator. The quantum linear extension seeks the corresponding object in BKM-symmetrised operator coordinates. Once such an object is written in weighted Gram form, the visibility law of Sections 4 to 7 follows from the abstract metric theorem.

We therefore work in a finite-dimensional operator space $\mathcal{V}$, equipped with a positive-definite BKM metric matrix M_{BKM} . The relevant observable sectors are M_{BKM} -orthogonal projectors, and the physically meaningful fluctuation object is a linearised correction operator associated with a quantum nonequilibrium steady-state generator. In the paper, we do not attempt to construct the full nonlinear quantum analogue of the classical bridge. Instead, we identify the precise linear fluctuation setting in which the weighted visibility theory transfers.

This distinction matters. The quantum part of the paper is not a speculative appendix to the classical theorem. It is the operator-space continuation that becomes exact once the fluctuation object has been linearised and expressed in the correct BKM geometry. What transfers here is conditional visibility geometry in linearised BKM coordinates, not yet the full backbone-whitened detectability geometry of the classical shadow construction.

10.2 Quantum observable envelope

The operator-space transfer theorem of Section 9 implies that once a quantum fluctuation correction is written in weighted Gram form relative to the BKM metric, its visible signal under finite observation is governed by the same envelope hierarchy as in the classical case. The consequences listed below are therefore conditional: they apply only once a physically meaningful quantum fluctuation correction with the required weighted Gram representation has been identified.

Concretely, let Δ_q be a finite-dimensional positive semidefinite quantum fluctuation correction acting on the BKM operator space, and assume that it admits a weighted signal factorisation

$$\Delta_q = B_q B_q^\sharp, \quad (10.1)$$

where $^\sharp$ denotes the adjoint with respect to the BKM metric. Let Π be an M_{BKM} -orthogonal projector of rank d . Then the visible quantum nonequilibrium signal is

$$V_\Pi(\Delta_q) := \text{tr}(\Pi \Delta_q \Pi), \quad (10.2)$$

and by the abstract theory this is the projected signal energy of the Euclideanised operator $\tilde{B}_q := M_{BKM}^{1/2} B_q M_{BKM}^{-1/2}$, with $\tilde{\Pi} := M_{BKM}^{1/2} \Pi M_{BKM}^{-1/2}$ and $V_{\Pi}(\Delta_q) = \|\tilde{\Pi} \tilde{B}_q\|_F^2$.

The immediate consequences are direct quantum analogues of the classical theorems:

1. the exact rank- d visible envelope is given by the Ky Fan top-singular-energy profile of the Euclideanised quantum signal operator,
2. the exact retention fraction is the cumulative singular-energy fraction,
3. the effective observability complexity is the stable rank of the Euclideanised quantum signal operator,
4. the optimal quantum finite observer is the projector onto the dominant singular front,
5. the envelope hierarchy is a spectroscopy of the visible quantum nonequilibrium signal.

Thus, at the level of weighted signal geometry, the classical and linear quantum finite-observation theories are formally identical. What changes is not the projector-level visibility law, but the physical origin and interpretation of the fluctuation object.

This is precisely where the recent BKM and quantum information-geometric literature becomes relevant. BKM geometry is already known to control nonadiabatic entropy production and path-action structure in slow quantum nonequilibrium steady-state transitions [21]. Quantum Fisher decompositions also distinguish incoherent and coherent geometric contributions to thermodynamic and dynamical constraints [22]. The paper adds a complementary statement: once the fluctuation correction is placed in BKM metric form, finite observation has a spectral visibility law.

10.3 Restricted information and quantum observation

The quantum extension is also the natural point of contact with recent work on restricted information and observation-dependent thermodynamic quantities.

Rubino, Brukner, and Manzano show that coarse-grained quantum thermodynamic quantities depend explicitly on observational resolution, even while fluctuation-theorem-type relations survive at each level of access [19]. Pernambuco and Céleri argue that restricted microscopic access should be treated geometrically as a gauge reduction on the space of states, leading to gauge-invariant entropy and fluctuation relations under limited information [20]. These papers differ significantly in mathematical machinery, but they share a common message: thermodynamic content depends on what the observer can resolve.

The present theory gives a new layer beneath that message. In the linear BKM fluctuation regime, restricted observation is not only a question of coarser thermodynamic bookkeeping or gauge-reduced distinguishability. It is a singular-value problem. The observer's access to quantum nonequilibrium structure is determined by the spectrum and singular-mode geometry of the weighted fluctuation signal.

This suggests a clean conceptual synthesis. Restricted information determines which operator-space directions are accessible. The BKM metric determines the physically correct geometry of fluctuation comparison. The weighted signal operator determines how much nonequilibrium structure those accessible directions can actually retain. In the linear BKM regime, the restricted-information question therefore becomes precise: identify a physically meaningful correction Δ_q admitting $\Delta_q = B_q B_q^\sharp$, and the finite-observation visibility theory then follows from the metric-universal Ky Fan framework. In this sense, the theory does not compete with the restricted-information literature. It supplies a missing linear-algebraic piece for that literature in the fluctuation setting.

10.4 Limits of the present quantum theorem

The operator-space transfer should not be overstated. As stressed at the end of Section 9, the paper proves a classical bridge and an abstract weighted visibility law, and then uses the latter to formulate a disciplined linear quantum extension. It does not claim a general nonlinear quantum Donsker-Varadhan theory, nor a complete quantum shadow theorem at the same level of explicitness as in the classical Markov setting. In particular, the full backbone-whitened detectability geometry of the classical shadow construction has not yet been identified and tested in operator-valued form.

Three limitations are especially important.

First, the physical identification of the quantum fluctuation correction remains linearised. The paper says that *if* the correction admits the appropriate weighted Gram representation in finite-dimensional BKM geometry, then the finite-observation machinery transfers. It does not derive that representation in full generality from first principles for arbitrary GKLS dynamics.

Second, the operator-space theory does not yet address noisy or non-projective finite observation channels at the same level of completeness. The projectors of the current paper describe sharp finite-rank observation sectors in the underlying geometry. General measurement channels and faulty observation maps remain a next step.

Third, the current paper does not attempt to settle the nonlinear quantum analogue of the diagonal-qubit falsifier and related beyond-quadratic issues discussed in the recent literature on partial observation and quantum thermodynamics. The paper stays deliberately at the linear fluctuation level where the weighted visibility geometry is exact.

These limitations mark the boundary of the work and indicate the next technical steps.

11 Worked Examples

These benchmarks are selected calibration points for the operator theory. Each is read through the same finite-observation ladder: skew source, backbone-weighted signal, hidden spectral tail under rank constraint, and the whitened spectrum controlling detectability against the detailed-balance backbone.

11.1 A four-state two-cycle network from the entropy-estimation literature

We take the four-state, five-link, two-cycle network used in the transition-based waiting-time literature, with rates given in Appendix E, Table IV of van der Meer, Ertel, and Seifert [37]. In that source problem, the observed link is the transition pair $2 \leftrightarrow 3$, and the forcing parameter F enters only through the rates on that link. At the benchmark value

$$F = \ln 3, \tag{11.1}$$

the rates become

$$Q = \begin{pmatrix} -36.7 & 1 & 35 & 0.7 \\ 8 & -9 & 1 & 0 \\ 0.2 & 6 & -56.2 & 50 \\ 75 & 0 & 2 & -77 \end{pmatrix}. \tag{11.2}$$

This first benchmark is genuinely multicyclic, small enough for transparent computation, and already recognised in the entropy-estimation literature. It therefore provides a compact test of the bridge, envelope, shadow, and detectability results on a standard network.

11.2 Envelope benchmark and the induced observed-link sector

For this generator the stationary law is

$$\pi \approx (0.381628, 0.207197, 0.247191, 0.163983). \quad (11.3)$$

The reduced Fisher space has dimension $m = 3$. The computed weighted signal operator A_{DV} has singular values

$$s_1(A_{DV}) \approx 7.672020, \quad s_2(A_{DV}) \approx 6.298419, \quad s_3(A_{DV}) = 0, \quad (11.4)$$

so the active nonequilibrium geometry is rank two. In reduced dimension three this is structural: the real skew carrier has rank at most two, so there are at most two nonzero singular values. The weighted stable rank is

$$\text{sr}(A_{DV}) \approx 1.673975. \quad (11.5)$$

The Ky Fan visibility levels are therefore

$$K_1^{DV} \approx 14.71497, \quad K_2^{DV} = K_3^{DV} \approx 24.63249, \quad (11.6)$$

which gives retained fractions

$$\Phi_1^{DV} \approx 0.597381, \quad \Phi_2^{DV} = 1. \quad (11.7)$$

Thus a single unconstrained optimal observable direction already retains about 59.7% of the full Gaussian nonequilibrium signal, while two directions recover the entire active sector. The random rank-one baseline is $1/3$.

The bridge relation is also numerically sharp in this example. Direct computation gives

$$\frac{\left\| \Delta_{DV} - \frac{1}{4} \hat{J} H_0^{-1} \hat{J}^T \right\|_F}{\|\Delta_{DV}\|_F} \approx 1.38 \times 10^{-6}, \quad (11.8)$$

which is a finite-precision numerical reconstruction residual rather than a structural mismatch. The benchmark is therefore not merely illustrative at the level of the envelope. It is also quantitatively consistent with the bridge developed in Section 3.

Figure 1 makes the unconstrained envelope structure explicit. To connect the benchmark directly to the source observation pattern, we now instantiate the actually observed link as an orthogonal sector inside the theorem class of the paper.

The waiting-time estimator of [37] is built from transitions along the observed edge $2 \leftrightarrow 3$. That observable is not literally an orthogonal projector on the reduced Fisher space, so we take its canonical Fisher-orthogonal rank-one surrogate

$$b_{23} := D_\pi^{-1/2}(e_2 - e_3) \in T_\pi, \quad q_{23} := \frac{C^T b_{23}}{\|C^T b_{23}\|}, \quad \Pi_{\text{obs}} := q_{23} q_{23}^T. \quad (11.9)$$

Because $\langle b_{23}, e_\pi \rangle = 1 - 1 = 0$, this direction lies naturally in the Fisher tangent space. It is




Figure 1: Observable visibility envelope for the four-state two-cycle benchmark at $F = \ln 3$. The reduced nonequilibrium signal is rank two in a three-dimensional reduced Fisher space. The rank-one retained fraction is $\Phi_1^{DV} \approx 0.597$, while $\Phi_2^{DV} = 1$, so two observable directions recover the entire active sector. Dashed lines indicate the random-observer baselines $1/3$ and $2/3$.

the orthogonal observer induced by the physical access pattern observe only the $2 \leftrightarrow 3$ link, while remaining inside the orthogonal-compression framework of Sections 2 to 6.

For this literature-induced sector we obtain

$$V_{\text{obs}}^{DV} = \text{tr}(\Pi_{\text{obs}} \Delta_{DV}) \approx 4.32702, \quad \Phi_{\text{obs}}^{DV} = \frac{V_{\text{obs}}^{DV}}{\text{tr}(\Delta_{DV})} \approx 0.175663. \quad (11.10)$$

Thus the actual observed edge captures well below the unconstrained optimum $\Phi_1^{DV} \approx 0.597381$, and even below the blind rank-one baseline $1/3$. At the raw-signal level, the source observable is therefore strongly misaligned with the dominant singular front of A_{DV} .

The shadow statistics tell a more refined story. For the same sector,

$$B_{\text{obs}} := \Pi_{\text{obs}} H_0 \Pi_{\text{obs}} |_{\text{Ran } \Pi_{\text{obs}}}, \quad E_{\text{obs}} := \Pi_{\text{obs}} \Delta_{DV} \Pi_{\text{obs}} |_{\text{Ran } \Pi_{\text{obs}}}, \quad (11.11)$$

so in this rank-one case the shadow score is the scalar

$$S_{\text{obs}} = B_{\text{obs}}^{-1/2} E_{\text{obs}} B_{\text{obs}}^{-1/2} \approx 0.252725. \quad (11.12)$$

The reverse Gaussian KL divergence is therefore

$$D_{KL}(\mathcal{N}_{DB,\text{obs}} \parallel \mathcal{N}_{NESS,\text{obs}}) = \frac{1}{2} (S_{\text{obs}} - \log(1 + S_{\text{obs}})) \approx 0.01370, \quad (11.13)$$

which gives

$$\text{TV}(\mathcal{N}_{DB,\text{obs}}, \mathcal{N}_{NESS,\text{obs}}) \leq 0.0828. \quad (11.14)$$

This already separates two questions that are blurred in much of the partial-observation literature. The physically observed edge is poor as a raw proxy for the full nonequilibrium signal, but it is less poor after backbone normalisation. Indeed, with $\text{tr}(T) \approx 0.597640$ and

Table 1: Rank-one comparison at $F = \ln 3$ for the four-state two-cycle benchmark. For raw visibility, the random row uses the Grassmannian rank-one expectation. On the detectability side, the quoted comparison uses the natural whitened rank-one benchmark $\text{tr}(T)/3$. The raw-optimal row maximises Φ^{DV} , the observed row is the literature-induced edge sector Π_{obs} , and the detectability-optimal row maximises $\text{tr}(S)$.

Observer sector	Φ^{DV}	$\text{tr}(S)$	$\tilde{\Phi} = \text{tr}(S)/\text{tr}(T)$	TV bound
Random raw baseline / whitened benchmark	1/3	0.199213	1/3	–
Raw-optimal rank-one	0.597381	0.292204	0.488930	0.0947
Observed 2 $\leftrightarrow$ 3 sector	0.175663	0.252725	0.422872	0.0828
Detectability-optimal rank-one	0.397326	0.298820	0.500000	0.0966

$\Theta_1 \approx 0.298820$, the detectable fraction of the actual observed sector is

$$\tilde{\Phi}_{\text{obs}} := \frac{\text{tr}(S_{\text{obs}})}{\text{tr}(T)} \approx 0.422872, \quad (11.15)$$

to be compared with the detectability optimum $1/2$. At $F = \ln 3$, the observed edge therefore sits below the raw rank-one baseline in visibility, but above the natural rank-one whitened benchmark in backbone-normalised detectability. This is the cleanest indication, on this benchmark, that the detectability landscape is materially flatter than the raw-visibility landscape.

Table 1 is the main point of the worked example. It shows that one can now rank concrete observation schemes geometrically. The gap between the observed 2 $\leftrightarrow$ 3 sector and the raw optimum is an exact alignment deficit, not a vague information-loss slogan. At the same time, the smaller gap in $\tilde{\Phi}$ shows that backbone-normalised detectability can be materially less pessimistic than raw signal retention.

11.3 A short forcing sweep on the actual observed sector

Because the source family is parametrised by a single force F acting on the observed link, it also furnishes a natural one-parameter sweep. Figure 2 plots

$$F \mapsto \Phi_{\text{obs}}^{DV}(F), \quad F \mapsto \text{tr}(S_{\text{obs}}(F)), \quad (11.16)$$

for the same induced observed-link sector $\Pi_{\text{obs}}(F)$.

Across the range $F \in [-2, 4]$, the raw retained fraction rises from about 0.012 to about 0.696, while the shadow trace rises from about 0.019 to about 0.307. The separation is already clear at moderate forcing: at $F = 0$, the observed sector still retains only $\Phi_{\text{obs}}^{DV} \approx 0.081$ of the total raw signal, but its backbone-normalised shadow strength is already $\text{tr}(S_{\text{obs}}) \approx 0.174$. This is the distinction between raw visibility and backbone-normalised detectability formalised in Section 6.

The worked benchmark shows that the present geometry can compare the actual source observable with the visibility-optimal and detectability-optimal observers, and can track how that gap changes across a physically motivated forcing family.




Figure 2: Observed-link forcing sweep for the edge-induced orthogonal sector $\Pi_{\text{obs}}(F)$ associated with the measured $2 \leftrightarrow 3$ link. Blue shows the raw retained fraction $\Phi_{\text{obs}}^{DV}(F)$. Orange shows the shadow trace $\text{tr}(S_{\text{obs}}(F))$. The grey dashed line marks the raw random rank-one baseline $1/3$, and the vertical dotted line marks the benchmark value $F = \ln 3$. The two curves separate visibly, showing that backbone-normalised detectability and raw signal retention are distinct geometric questions even for the same physical observation channel.

11.4 A five-state ribosome proofreading benchmark from the speed-accuracy literature

The four-state benchmark shows that the theorem package can compare actual and optimal observers on a standard partial-observation network. We now test the same machinery on a biologically canonical proofreading model. Banerjee, Kolomeisky, and Igoshin analyse aa-tRNA selection by the *E. coli* ribosome through a five-state kinetic-proofreading network, together with hyperaccurate and error-prone mutants, and argue that the experimentally relevant regime is organised primarily by speed, with energetic proofreading cost as an additional constraint rather than by direct minimisation of error alone [38].

To place that system inside the finite-observation framework of the present paper, we use the closed turnover continuous-time Markov model on the state set

$$(E, ER, ER^*, EW, EW^*), \quad (11.17)$$

where E is the free ribosome, ER and ER^* are the cognate branch before and after hydrolysis, and EW and EW^* are the near-cognate branch before and after hydrolysis. The right-pathway rates and discrimination factors are taken directly from the Banerjee supporting-information tables for the wild type (WT), hyperaccurate mutant (HYP, rpsL141), and error-prone mutant (ERR, rpsD12) ribosomes [38].

As in Banerjee’s translation analysis, we impose the nearly irreversible choices $k_{-2,R} = k_{-3,R} = 10^{-3} \text{ s}^{-1}$ and $f_{-2} = 1$, and derive f_{-3} from the cycle constraint. The generator used here is not the backward first-passage operator from the source paper. It is the ergodic turnover generator in which both catalytic completion and proofreading reset return the system to the free state E .

Before applying finite observation, we validate the reconstruction against the source biology. The stationary fluxes recover the correct speed and error ordering

$$\text{speed: WT} > \text{HYP} > \text{ERR}, \quad \text{error: HYP} < \text{WT} < \text{ERR}, \quad (11.18)$$

and the flux-based proofreading cost agrees with the analytic Banerjee formula to relative error below 2×10^{-3} for all three variants. The benchmark is a source-locked turnover model




Figure 3: Raw weighted-signal spectrum and cumulative Ky Fan retention for the Banerjee ribosome benchmark. For WT, HYP, and ERR, the singular front of A_{DV} is sharply concentrated. Rank one retains roughly two thirds of the full correction, and rank two retains more than 95% in every variant.

that reproduces the intended biological regime before any spectral analysis is performed.

11.5 Raw spectral concentration and interpretation of effective complexity

The reduced Fisher tangent space has dimension $m = 4$ for each ribosome variant. Figure 3 and Table 2 show the singular spectrum of the weighted signal operator A_{DV} and the associated Ky Fan retention hierarchy.

Across all three variants the raw nonequilibrium signal is sharply concentrated. The weighted stable rank lies in the narrow interval

$$1.4130 \leq \text{sr}(A_{DV}) \leq 1.5166, \quad (11.19)$$

while the rank-one retained fractions satisfy

$$0.6594 \leq \Phi_1^{DV} \leq 0.7077. \quad (11.20)$$

More importantly, two observer directions already recover almost the entire Donsker-Varadhan correction:

$$\Phi_2^{DV}(\text{WT}) \approx 0.9867, \quad \Phi_2^{DV}(\text{HYP}) \approx 0.9913, \quad \Phi_2^{DV}(\text{ERR}) \approx 0.9504. \quad (11.21)$$

Thus the active signal is effectively rank two in a four-dimensional reduced Fisher space, despite the fact that the underlying proofreading network is multibranch and is usually discussed through several competing biological performance measures. Within the raw finite-observation geometry, the benchmark exhibits a pronounced appearance of spectral simplicity despite topological and functional richness: topological complexity does not imply spectral complexity.

11.6 Backbone-normalised detectability and biologically legible observer directions

The whitened detectability spectrum sharpens the comparison between variants. Figure 4 shows that WT and HYP have an almost perfectly rank-two detectability structure: in both

Table 2: Finite-observation summary for the Banerjee ribosome benchmark. The last two columns separate raw concentration from backbone-normalised detectability concentration. Raw rank two suffices for all three variants, whereas the ERR mutant requires full rank to exceed the 90% detectability threshold.

Variant	$\text{sr}(A_{DV})$	Φ_1^{DV}	Φ_2^{DV}	$d_{0.9}^{\text{raw}}$	$d_{0.9}^{\text{det}}$	top θ_k
WT	1.4709	0.6799	0.9867	2	2	0.324613
HYP	1.4130	0.7077	0.9913	2	2	0.320700
ERR	1.5166	0.6594	0.9504	2	4	0.327211




Figure 4: Backbone-normalised detectability spectrum for the Banerjee ribosome benchmark. WT and HYP are nearly rank two after whitening, while ERR retains a substantial lower pair. This is the clearest place where raw visibility and detectability separate across biologically meaningful variants.

cases a nearly degenerate leading pair at

$$\theta_1 \approx \theta_2 \approx 0.32 \quad (11.22)$$

dominates, while the residual tail is tiny,

$$\theta_3 = \theta_4 \approx 0.0062 \quad \text{for WT}, \quad \theta_3 = \theta_4 \approx 0.0026 \quad \text{for HYP}. \quad (11.23)$$

Consequently, two detectability directions already capture more than 98% of the backbone-normalised signal for WT and HYP. The ERR mutant shows a more mixed geometry. Its leading pair remains dominant,

$$\theta_1 \approx \theta_2 \approx 0.3272, \quad (11.24)$$

but the lower pair remains substantial,

$$\theta_3 \approx \theta_4 \approx 0.1405, \quad (11.25)$$

so rank two captures only about 70% of detectability and one needs full rank to exceed the 90% threshold. Raw retention is therefore highly concentrated in all three variants, but detectability distinguishes the speed-optimised WT and HYP variants from the more diffuse ERR mutant.

Figure 5 shows that the dominant observer directions are also biologically legible. Up to the usual overall sign convention, the top raw observer loads primarily on the cognate branch $E \leftrightarrow ER \leftrightarrow ER^*$, while the subdominant direction separates cognate from near-cognate structure. The top detectability observer places its weight mainly on the hydrolysis stage, which is consistent with the role of $ER \rightarrow ER^*$ as the main distinguishability channel after backbone normalisation and with hydrolysis as the natural discrimination checkpoint of the




Figure 5: Top raw and top detectability observer loadings for WT, HYP, and ERR. The labels E , ER , ER^* , EW , and EW^* denote the free, cognate pre-hydrolysis, cognate post-hydrolysis, near-cognate pre-hydrolysis, and near-cognate post-hydrolysis states. Overall signs are conventional. The dominant directions remain concentrated on biologically recognisable proofreading coordinates rather than on diffuse mixtures of all states.

proofreading cycle. The observer geometry is therefore not only low rank. It is interpretable in terms of the proofreading mechanism itself. As the matched-null comparison later makes clear, however, the resulting low-rank concentration is not by itself specific evidence for proofreading architecture.

11.7 Robustness under hydrolysis-rate variation

Banerjee’s own speed-accuracy analysis identifies the hydrolysis step as a key kinetic control axis for the WT ribosome [38]. We therefore perform a one-parameter finite-observation sweep over $k_{2,R}$ while keeping the remaining WT parameters fixed. Figure 6 shows that the low-rank picture is robust across this family. The rank-one retained fraction stays above 0.66, the rank-two retained fraction stays above 0.969, the weighted stable rank remains close to 1.5, and the top detectability eigenvalue rises smoothly toward a plateau near 0.33. Across the hydrolysis-rate sweep, the dominant detectability structure remains low-rank and varies smoothly. There is no sharp spectral transition in the biologically relevant band. The spectral simplicity of the WT proofreading network is therefore a structural feature of the benchmark family, not a knife-edge artefact of one fitted parameter set.

Taken together, the Banerjee benchmark upgrades the theorem package from a compact literature network to a biologically canonical proofreading system with experimentally grounded parameters and variant structure. The conclusion is unusually clean. A network that is functionally rich enough to support speed-accuracy-cost trade-off analysis can still be spectrally simple in the finite-observation geometry. In this benchmark, most of the nonequilibrium correction lives in two observer directions, and the remaining distinction between variants becomes visible only after backbone whitening.




Figure 6: WT hydrolysis-rate sweep for the Banerjee ribosome benchmark. The raw rank-one retained fraction, raw rank-two retained fraction, weighted stable rank, and top detectability eigenvalue remain in a narrow low-rank regime across the relevant range of $k_{2,R}$. The dashed vertical line marks the WT value from the source model.

11.8 A matched random null for the Banerjee benchmark

The Banerjee ribosome benchmark is useful, but its structural interpretation must be stated carefully. We therefore compare the WT proofreading network against a matched null ensemble of dense random five-state nonequilibrium CTMCs with the same mean exit-rate scale and with total DV energy restricted to a narrow band around the real benchmark. The null ensemble was designed to answer a specific question: is the strong raw rank-two compression in the ribosome case a special fingerprint of proofreading architecture, or is it largely a dimensional effect of working in a four-dimensional reduced Fisher space?

This random-ensemble perspective is also consonant with recent work showing that, in broad classes of random Markov models, many irreversibility-sensitive observables are controlled primarily by heterogeneity rather than by asymmetry alone except near symmetry [39].

The answer is the latter. As Figure 7 shows, across the primary raw metrics the real Banerjee WT point sits close to the middle of the matched null distribution rather than in an extreme low-rank tail. In particular,

$$\Phi_1^{DV}(\text{WT}) \approx 0.6799, \quad \Phi_2^{DV}(\text{WT}) \approx 0.9867, \quad \text{sr}(A_{DV}) \approx 1.4709, \quad (11.26)$$

so the WT point sits close to the middle of the sampled null ensemble on all three raw metrics. Thus the Banerjee network remains an important demonstration that the finite-observation machinery can be applied to a real biochemical system with biologically legible observer directions, but the raw low-rank compression itself cannot be claimed as a special consequence of proofreading architecture.




Figure 7: Matched random null for the Banerjee ribosome benchmark. The WT proofreading network lies near the centre of the null distribution for the main raw concentration measures once the total DV scale is matched. The benchmark therefore remains useful as a biologically interpretable worked example, but its strong raw rank-two concentration is not by itself evidence for a proofreading-specific structural simplification.

11.9 A force-driven kinesin benchmark from the molecular-motor literature

To test whether the same spectral compression extends beyond proofreading, we next turn to the chemomechanical kinesin model of Liepelt and Lipowsky [40, 41]. This benchmark is mechanistically distinct from the Banerjee ribosome network. Its primary form is a six-state motor with two competing chemomechanical cycles and one purely dissipative cycle, together with explicit load-force control through a mechanical load-distribution factor and force-sensitive chemical prefactors. The supplement also gives a reduced seven-state small-load extension in which a doubly ADP-bound state opens an additional forward cycle. We use the source-locked six-state model at arbitrary load, and the reduced seven-state model only in the zero-load regime where the source supplement explicitly justifies that truncation and supplies the extra rates.

For the six-state model the transition rates are built in the source convention as

$$\omega_{ij} = \kappa_{ij} I_{ij}([ATP], [ADP], [P]) \Phi_{ij}(F), \quad (11.27)$$

with rate constants from Table I of [40], force parameters from Table II and Appendix A of [41], and the backward-cycle rate $\kappa_{54} = (\kappa_{52}/\kappa_{25})^2 \kappa_{21}$ imposed as in the source derivation. The reduced seven-state model retains only the additional edges $1 \leftrightarrow 7$ and $5 \leftrightarrow 7$, again as prescribed in Appendix D of [41].

Before applying finite observation, we validate the reconstruction against the source dynamics: both force-data parametrisations reproduce a stall force near 7.2 pN, the six-state model reproduces the source asymmetry between $[P]$ - and $[ADP]$ -dependence at zero load, and the reduced seven-state model restores the missing high-ADP slowdown while producing the

Figure 1: 6-state raw DV spectrum across force




Figure 8: Six-state kinesin benchmark under force. Left panels: singular values of the weighted signal operator A_{DV} for the [15] and [9] source fits. Middle panels: cumulative raw Ky Fan retention fractions Φ_1^{DV} and Φ_2^{DV} . Right panels: weighted stable rank. The signal is almost rank one at low load and broadens only near stall, where the forward and backward chemomechanical cycles compete most strongly.

predicted crossover to the added F_{DD} cycle at $[ADP] \approx 1083 \mu\text{M}$.

11.10 Raw concentration across force and concentration

The six-state kinesin benchmark is sharply low-rank across the tested regimes. At zero load, all three published source parametrisations are essentially rank one in the raw Donsker–Varadhan correction:

$$\Phi_1^{DV}([15], F=0) \approx 0.9777, \quad \Phi_1^{DV}([9], F=0) \approx 0.9988, \quad \Phi_1^{DV}([12], F=0) \approx 0.9563, \quad (11.28)$$

with weighted stable rank in the narrow interval

$$1.0012 \leq \text{sr}(A_{DV}) \leq 1.0457. \quad (11.29)$$

Thus a single observer direction already captures more than 95% of the nonequilibrium correction in every zero-load case.

Figure 8 shows how this picture evolves under load for the two source force-data parametrisations. Broadening does occur near stall, but it is physically transparent rather than pathological. For the [15] fit, rank-one retention drops from 0.9777 at zero load to about 0.5260 near 7 pN, while the weighted stable rank rises from 1.0228 to 1.9012.

For the [9] fit, the broadening is milder, with $\Phi_1^{DV} \approx 0.6462$ and $\text{sr}(A_{DV}) \approx 1.5474$ at 7 pN. In both cases, however, the low-rank tendency persists away from the stall region. This is the behaviour one would expect if the singular front is controlled by the competition between forward and backward chemomechanical cycles: the spectrum is narrow when one cycle dominates and broadens only when the cycle currents come close to balance.

Figure 3: 6-state zero-load concentration stress test ([12] params)




Figure 9: Zero-load concentration stress test for the six-state kinesin model using the [12] source parametrisation. The $[P]$ -sweep leaves the raw singular spectrum and cumulative retention almost unchanged. The $[ADP]$ -sweep produces a mild increase in concentration while the reference velocity falls strongly, matching the source observation that the six-state model captures the phosphate trend but underestimates the ADP slowdown.

The zero-load concentration stress test in Figure 9 shows a second important feature. The six-state geometry is almost insensitive to $[P]$, for which the source paper already reports good kinetic agreement, while $[ADP]$ produces only a mild spectral shift even in the regime where the six-state velocity law is known to underpredict the slowdown. For the [12] parametrisation at $F = 0$, the raw rank-one retained fraction stays near 0.956 throughout the $[P]$ -sweep and rises only slightly, from about 0.958 to 0.967, over four decades of $[ADP]$. Thus the kinetic failure of the six-state model at high ADP is not a failure of raw spectral concentration.

11.11 Backbone-normalised detectability and the reduced seven-state extension

The whitened detectability spectrum is materially broader than the raw spectrum. Figure 10 shows that the top detectability eigenvalue remains of order 10^{-2} to 10^{-1} , but the detectability mass is spread over more directions than the raw DV correction. In the representative six-state cases of Table 3, one typically needs rank four to exceed the 90% detectability threshold even when the raw signal is already more than 95% concentrated in rank one. This reproduces, in a mechanochemically different system, the same separation between raw visibility and backbone-normalised detectability that was already visible in the ribosome benchmark.

The reduced seven-state model strengthens the cross-model recurrence case further. Its purpose in the source literature is to repair the missing $[ADP]$ -dependence of the six-state motor by adding a doubly ADP-bound state and the corresponding forward cycle $F_{DD}^+ = |12571\rangle$. In the finite-observation geometry, however, the added mechanistic structure does not broaden the signal. Figure 11 shows that the rank-one retained fraction actually increases with $[ADP]$, from about 0.9560 at $1 \mu\text{M}$ to 0.9740 at $5000 \mu\text{M}$, while the weighted stable rank falls from 1.0460 to 1.0267. At the same time, the top detectability eigenvalue decreases from 6.24×10^{-2} to 2.82×10^{-2} , and the minimal detectability rank for 90% retention improves from four to two. The source paper's extra cycle becomes

Figure 2: 6-state detectability spectrum across force




Figure 10: Backbone-normalised detectability for the six-state kinesin benchmark under force. Detectability is substantially broader than the raw weighted signal, and the broadening becomes strongest near stall. The theory therefore distinguishes raw nonequilibrium visibility from statistical distinguishability against the detailed-balance backbone even within the same molecular-motor family.

dynamically important without destroying the low-rank observer geometry.

Figure 12 shows that the leading observer directions remain structured rather than diffuse. In the six-state family, the dominant raw and detectability directions concentrate on only a few numbered states at a time, and in the seven-state ADP regime the new state 7 enters coherently rather than spreading weight uniformly across the enlarged network. This again points to a low-dimensional active nonequilibrium core rather than to generic high-dimensional mixing.

Taken together, the kinesin benchmark materially strengthens the case for cross-model recurrence. The Liepelt–Lipowsky motor is not a proofreading network. It has competing chemomechanical cycles, explicit load-force control, and, in the reduced seven-state extension, a third forward cycle that becomes dominant at high ADP. Yet the finite-observation geometry is at least as concentrated as in the Banerjee ribosome benchmark across most of the tested regimes. The one place where concentration weakens, near stall, is physically transparent: it is exactly where the forward and backward motor cycles come closest to balance. This is therefore not a failure of the low-rank theory but one of its clearest mechanistic confirmations.

11.12 Matched negative controls for the kinesin motor benchmark

Unlike the Banerjee benchmark, the kinesin benchmark behaves differently. Here the positive result survives matched negative controls. We compare three real operating points against graph-matched nonequilibrium control ensembles: anchor A, the six-state motor at zero load; anchor B, the same six-state model near stall; and anchor C, the reduced seven-state model at zero load and high ADP. Two control classes are used. NC1 keeps the graph but

Figure 4: Reduced 7-state model ADP sweep ($F=0$, [12] params)




Figure 11: Reduced seven-state kinesin benchmark at zero load under an $[ADP]$ sweep. The added doubly ADP-bound state and cycle $F_{DD}^+ = |12571\rangle$ improve the source kinetic fit at high ADP, but they do not broaden the finite-observation geometry. Raw rank-one retention increases, the stable rank falls slightly, and detectability becomes easier to compress at large ADP.

Table 3: Representative finite-observation summary for the kinesin benchmark. Raw rank-one concentration is already extreme at low load and remains strong in the reduced seven-state extension. Detectability is broader than raw visibility in the six-state motor but becomes easier to compress in the high-ADP seven-state regime.

Case	n	Φ_1^{DV}	Φ_2^{DV}	$sr(A_{DV})$	$\text{top } \theta_1$	$d_{0.9}^{\text{raw}}$	$d_{0.9}^{\text{det}}$
6s,[15], $F = 0$	6	0.9777	0.9911	1.0228	0.0623	1	4
6s,[15], $F = 7$ pN	6	0.5260	0.8695	1.9012	0.0465	3	4
6s,[9], $F = 0$	6	0.9988	0.9995	1.0012	0.0623	1	4
6s,[12], $F = 0$	6	0.9563	0.9825	1.0457	0.0624	1	4
7s,[12], $[ADP] = 5000 \mu\text{M}$	7	0.9740	0.9966	1.0267	0.0282	1	2

replaces the structured motor rates by random positive rates. NC2 is stricter: it keeps the graph and approximately matches the stationary law while randomising the nonequilibrium edge asymmetry.

At the two main low-rank anchors, the real motor is substantially more concentrated than the controls, as shown in Figure 13. For anchor A the real six-state motor has

$$\Phi_1^{DV} \approx 0.9777, \quad sr(A_{DV}) \approx 1.0228, \quad (11.30)$$

placing it above all sampled NC1 controls and above most sampled values in the stricter π -matched NC2 ensemble. For anchor C the reduced seven-state model has

$$\Phi_1^{DV} \approx 0.9740, \quad sr(A_{DV}) \approx 1.0267, \quad (11.31)$$

placing it above nearly all sampled NC1 controls and above most sampled NC2 controls. Thus the low-rank concentration is not a generic consequence of graph size or edge count alone, and it is only partly explained by the stationary-law profile.

The one exception is the near-stall anchor B, which is now taken at the same $F = 7$ pN operating point reported in Table 3. At this anchor the real motor becomes *less* concentrated




Figure 12: Leading raw and detectability observer loadings for representative six-state and reduced seven-state kinesin operating points in the canonical numbered state basis. Overall signs are conventional. The dominant directions remain concentrated on a small number of states and reorganise coherently as load or ADP changes, rather than diffusing across the full state graph.




Figure 13: Kinesin matched negative controls. The red curves are the real motor benchmarks and the blue bands are graph-matched control ensembles. At the zero-load six-state anchor and the high-ADP seven-state anchor, the real motor retains substantially more DV signal in low observer rank than the controls. Near stall, the situation reverses: the real motor broadens more strongly than the controls, consistent with direct competition between forward and backward chemomechanical cycles.

than the matched random controls, with

$$\Phi_1^{DV} \approx 0.5260, \quad \text{sr}(A_{DV}) \approx 1.9012, \quad (11.32)$$

placing the real point below most sampled random controls. The qualitative conclusion is unchanged but now better calibrated: near stall the real motor is less concentrated than typical controls, though not pathologically so. This is the regime in which the forward and backward chemomechanical cycles compete most strongly. The negative control sharpens the mechanistic reading of the benchmark by identifying the stall region as the place where low-rank concentration weakens.

Taken together, the negative controls sharpen the interpretation of both biological examples. For the Banerjee ribosome, the benchmark is valuable mainly as an experimentally grounded and biologically interpretable worked example, but not as evidence that proofreading architecture is unusually low-rank. For kinesin, by contrast, the control ensembles show that the observed spectral compression is not generic to comparable nonequilibrium graphs. The molecular-motor result therefore provides the stronger empirical support for a nontrivial cross-model recurrence claim.

11.13 A ten-state chemotaxis feedback benchmark from the sensory-adaptation literature

The kinesin benchmark shows that strong finite-observation compression can survive outside proofreading and across explicit motor-force control. We now test a larger sensory-feedback system. Sartori, Granger, Lee, and Horowitz study *E. coli* chemotaxis through a ten-state receptor model with two activity states $a \in \{0, 1\}$ and five methylation levels $m \in \{0, 1, 2, 3, 4\}$, maintained out of equilibrium by SAM-driven methylation [42]. At fixed ligand concentration $[L]$, the state space is

$$(a, m) \in \{0, 1\} \times \{0, 1, 2, 3, 4\}, \quad (11.33)$$

ordered in the implementation as

$$(0, 0), (1, 0), (0, 1), (1, 1), (0, 2), (1, 2), (0, 3), (1, 3), (0, 4), (1, 4). \quad (11.34)$$

The source kinetics are fixed by the Tar receptor parameters

$$K_I = 18.2 \mu\text{M}, \quad K_A = 3000 \mu\text{M}, \quad \Delta_m = 2, \quad m_0 = 1, \quad (11.35)$$

together with the separated timescales

$$\tau_a = 10^{-3} \text{ s}, \quad \tau_m = 10 \text{ s}, \quad (11.36)$$

and chemical driving

$$\Delta\mu = 6 k_B T. \quad (11.37)$$

We build the fixed-ligand generator directly from the source rate formulas for activity flips and nearest-neighbour methylation moves, and evaluate it at

$$[L] \in \{94, 200, 500, 720, 1500, 3000, 5760\} \mu\text{M}, \quad (11.38)$$

which spans the adaptive regime $K_I \ll [L] \ll K_A$, its upper edge, and the strongly saturated outside region [42].

Before any spectral analysis, we impose the same source-locking standard used in the previous biological examples. The implemented generator reproduces the ten-state receptor graph of the source model, including all horizontal activity flips and all vertical nearest-neighbour methylation transitions. Local detailed balance holds on the activity sector by construction, and the SAM-driven methylation asymmetry breaks detailed balance at every tested ligand value. The resulting stationary laws are positive and normalised. They also reproduce a key qualitative feature of the source paper: the methylation marginal remains heavily concentrated at the highest methylation level, with Bhattacharyya overlaps above 0.999 across representative ligand pairs. Thus the benchmark is not an abstract ten-state toy. It is a source-locked receptor model in the specific high-methylation operating regime used by Sartori *et al.* to discuss the limited memory capacity of the feedback architecture.

The raw Donsker–Varadhan correction is sharply concentrated across the entire ligand family, as shown in Figure 14. In the adaptive regime, the one-direction retained fraction is already

$$\Phi_1^{DV} \approx 0.972 \quad \text{at } [L] = 94 \mu\text{M}, \quad (11.39)$$

with

$$\Phi_2^{DV} \approx 0.999, \quad \text{sr}(A_{DV}) \approx 1.029. \quad (11.40)$$

Figure 1: Raw DV spectrum and cumulative retention




Figure 14: Raw weighted-signal spectrum and cumulative Ky Fan retention for the Sartori ten-state chemotaxis benchmark. Across all tested ligand concentrations, the weighted signal operator A_{DV} has a sharply dominant singular front. One observer direction captures about 85% to 97% of the raw DV signal, and two directions capture at least about 99.5%.

Even at the largest tested ligand concentration,

$$[L] = 5760 \mu\text{M}, \quad (11.41)$$

well beyond K_A , the signal remains strongly compressed:

$$\Phi_1^{DV} \approx 0.850, \quad \Phi_2^{DV} \approx 0.995, \quad \text{sr}(A_{DV}) \approx 1.176. \quad (11.42)$$

Thus a ten-state sensory-feedback network with a nine-dimensional reduced Fisher space still behaves, in raw finite-observation geometry, almost as a rank-one system. This is materially stronger than the Banerjee proofreading example, where an analogous low-rank conclusion could still be questioned as a small-state-space effect.

As in the previous worked examples, raw visibility and backbone-normalised detectability do not coincide, as shown in Figure 15. The top detectability eigenvalue rises from about

$$\theta_1 \approx 6.68 \times 10^{-6} \quad (11.43)$$

at $94 \mu\text{M}$ to a peak near

$$\theta_1 \approx 3.86 \times 10^{-5} \quad (11.44)$$

around $3000 \mu\text{M}$, while the detectability stable rank grows from about 2.14 to 3.00 in the present backbone-normalised scaling. In cumulative terms, one raw direction suffices for 90% retention over most of the adaptive regime, whereas the backbone-whitened concentration typically requires two to four directions. The chemotaxis benchmark therefore reinforces, in a mechanistically different setting, the main message of Section 6: the visible signal and the statistically distinguishing signal against the detailed-balance backbone are related, but they are not the same geometric object.

Figure 2: Detectability spectrum and cumulative concentration





Figure 15: Backbone-normalised detectability spectrum for the Sartori chemotaxis benchmark. Raw visibility is almost rank one throughout the ligand family, but detectability is broader: the leading whitened eigenmodes retain most, but not all, of the concentration. This is the same structural separation between raw nonequilibrium signal and backbone-normalised distinguishability already seen in the ribosome and kinesin benchmarks.

Figure 3: Observer loadings in the ten-state basis (blue = inactive $a=0$, red = active $a=1$)



The leading observer directions remain interpretable rather than diffuse. Figure 16 shows that the dominant raw observer is concentrated almost entirely on the lowest methylation level, as an activity contrast between $(0, 0)$ and $(1, 0)$, with only a small secondary correction at $m = 1$. The detectability-optimal direction is broader and more ligand-dependent: in some regimes it is dominated by the inactive low-methylation state $(0, 0)$, in others by the active low-methylation state $(1, 0)$, with smaller balancing weight on $(0, 1)$ or $(1, 1)$. This is the pattern one expects when the raw nonequilibrium correction is nearly rank one but the detailed-balance backbone is anisotropic.

Figure 4: Ligand-regime sweep




Figure 17: Ligand sweep for the Sartori ten-state chemotaxis benchmark. The green band marks the adaptive regime $K_I \ll [L] \ll K_A$. Top left: Φ_1^{DV} and Φ_2^{DV} . Top right: weighted stable rank. Bottom left: top detectability eigenvalue. Bottom right: total DV signal energy. The raw concentration is sharpest inside the adaptive regime and degrades only gradually outside it, while the total DV signal grows monotonically and saturates above K_A .

The continuous ligand sweep in Figure 17 makes the operating-point dependence clear. Inside the adaptive window, the raw signal is most concentrated:

$$\Phi_1^{DV} \approx 0.97, \quad \text{sr}(A_{DV}) \approx 1.03. \quad (11.45)$$

As $[L]$ rises through and beyond K_A , the rank-one retained fraction decays smoothly toward about 0.79, while Φ_2^{DV} remains close to 0.99, the weighted stable rank rises only to about 1.27, and the total DV signal energy grows monotonically before saturating. There is no sharp spectral transition at the edge of the adaptive regime. The geometry changes gradually with ligand level rather than catastrophically.

Figure 5: Methylation-distribution validation (cf. Sartori S3)




Figure 18: Validation against the source methylation-distribution picture. The stationary methylation marginal is heavily concentrated at $m = 4$ for both in-regime and out-of-regime ligand pairs, with Bhattacharyya overlaps above 0.999. This reproduces the source paper’s limited-memory observation and shows that the finite-observation signal concentration is extracted on top of, rather than instead of, the high-overlap methylation statistics.

The methylation validation in Figure 18 is important for interpretation. The source model operates in a regime where the stationary methylation marginals remain extremely similar across ligand values, even when the mean activity changes appreciably. The finite-observation result is therefore not that the methylation memory is broad and low-rank. It is subtler. The raw nonequilibrium correction is sharply concentrated even though the methylation marginals themselves are almost indistinguishable. This explains why the dominant observer localises on a low-methylation activity contrast rather than on a broad redistribution over all methylation levels.

Taken together, the Sartori benchmark strengthens the empirical recurrence case in a more demanding state space than Banerjee. The model is mechanistically distinct from both proofreading and chemomechanical transport, and its reduced Fisher space has dimension nine rather than four or five. Yet the raw finite-observation geometry remains strongly compressed throughout the tested ligand family. At the same time, the benchmark preserves the same structural lesson already seen in the earlier biological examples: detectability is broader than raw visibility, and the leading observer directions remain physically interpretable. The correct conclusion is therefore not that every aspect of chemotactic memory is low-dimensional. It is that, in this source-locked ten-state feedback receptor model, the Donsker–Varadhan correction itself has a sharply concentrated spectral front. What remains open for this benchmark is a matched-null control family comparable to the kinesin negative-control analysis; until that is in hand, the chemotaxis example should be read as a source-locked worked example rather than as evidence that sensory-feedback architecture generically enforces low rank.

12 Related-work connections

12.1 Estimators, hidden events, and coarse observation

A large part of the partial-observation literature studies estimator quality, lower bounds, hidden events, unresolved transitions, and coarse-graining schemes for irreversibility detection [5–11, 13, 36, 43]. Those works ask how much irreversibility can be inferred from restricted

data.

The present paper addresses a complementary question. Alongside asking which estimator is best, one can ask how much nonequilibrium signal the chosen observable sector can retain in principle. In the regime treated here, that question is governed by a weighted spectral envelope, its hidden share is the spectral tail, and observer quality is controlled by alignment with the singular front of the correction. The Gaussian shadow divergences studied here are local observed-Gaussian comparison laws, conceptually adjacent to trajectory-level irreversibility measures such as Roldán–Parrondo, but not the same object.

12.2 Observer dependence, coarse graining, and design

Restricted-information and coarse-grained thermodynamic frameworks emphasise that observation is part of the thermodynamic problem itself [8, 17–20]. Work on current fluctuations, random counting observables, and optimised coarse-graining schemes makes the same point from a design perspective: some observations are much better than others for irreversibility detection [14–16].

The alignment and shadow theorems make that statement concrete in the present setting. A visibility-optimal observer aligns with the singular front of the raw signal. A detectability-optimal observer aligns with the singular front of the backbone-whitened signal. The theory therefore gives a geometric ranking of observers rather than only a qualitative statement that coarse observation can help or hurt.

12.3 Conceptual and geometric context

Several neighbouring literatures sharpen the conceptual space. Frenesy and time-symmetric dynamical activity make clear that nonequilibrium structure is richer than entropy-production rates alone [44]. Conditioned-process theory shows how distinct microscopic dynamics can collapse onto equivalent effective descriptions under large-deviation tilts [45].

Spectral work on nonreversible chains shows that singular-value and gap objects can carry operational content rather than serving as linear-algebraic decoration; in the present context such gap control helps set the scale of the detailed-balance backbone against which the weighted signal is read [46]. The cycle-current perspective of Schnakenberg and Andrieux–Gaspard clarifies the physical source of the skew sector, while reversible-chain information geometry provides a natural detailed-balance backdrop for the Fisher-coordinate formulation used here [28, 30, 31]. We also note a nearby geometric literature around Souriau’s Lie-group thermodynamics and Barbaresco’s Souriau–Fisher information geometry [47–49].

Taken together, these neighbouring literatures locate the results clearly. The paper is complementary to work on bounds, estimators, and coarse-grained thermodynamics. Its specific contribution is a finite-dimensional visibility law for the Donsker–Varadhan correction in the Markov Gaussian regime, together with its observer geometry, shadow criterion, and ordered spectral readout.

13 Scope

13.1 What is proved

Within this setting, the Donsker–Varadhan correction is identified as a positive weighted correction operator generated by the reduced skew channel and factorising through the weighted signal operator and the detailed-balance backbone. Within its stated Gaussian and orthogonal-observer scope, the theory gives the visibility law, the full envelope hierarchy, retention fractions, hidden tail, stable-rank obstruction, alignment formula, optimal and generic orthogonal observers, the shadow formulas, and the backbone-whitened detectability reformulation. In particular, the envelope hierarchy determines the eigenvalue profile of Δ_{DV} , equivalently the singular-energy profile $\frac{1}{4}s_k(A_{DV})^2$, while the shadow hierarchy supplies the distinct spectrum relevant for Gaussian distinguishability.

The abstract weighted-geometry theorem also holds in finite dimension. It is a transfer theorem: once a positive semidefinite correction admits a weighted Gram factorisation relative to a positive-definite metric, the envelope, alignment, visibility hierarchy, and spectral-readout machinery transfer by Euclideanisation. This includes BKM-symmetrised operator space.

It is equally important to state what is not proved here. Projection suppression and shared observed-Gaussian behaviour are asymptotic Gaussian corollaries, not new bridge theorems. The exact separation proved here is a Hessian-level statement, and beyond quadratic order the paper does not rule out mixed dissipative-skew terms in the full rate function. The linear quantum extension is only at the level of weighted geometry transfer. It is therefore a conditional geometric framework rather than a tested model-specific theory, since no concrete quantum correction operator with the required structure is exhibited here. The full nonlinear quantum fluctuation bridge remains open.

13.2 Falsification

At the classical bridge level, the theory would fail if the reduced-coordinate correction were not representable as a weighted Gram signal built from the skew sector and the detailed-balance backbone. Numerically, the bridge claim fails if a direct computation gives $\|\Delta_{DV} - \frac{1}{4}\hat{J}H_0^{-1}\hat{J}^T\|$ above tolerance. The bridge identity is therefore a genuine key step.

At the finite-observation level, the envelope law would fail if direct optimisation over orthogonal rank- d projectors did not agree with the singular-energy formula for K_d^{DV} , or if finite observation in the relevant fluctuation geometry were not correctly modelled by orthogonal compression on the reduced Euclidean Fisher space. In other words, if the physically relevant observable sectors could not be represented in this way, then the visibility law would no longer be the correct model.

At the statistical level, the Gaussian shadow package fails if KL, Hellinger, or total-variation quantities computed directly from the compressed Gaussian covariances do not match the spectral formulas in the shadow operator. At the same time, the equilibrium-shadow theorem is only a Gaussian theorem. It would be falsified as a general statement about observed nonequilibrium if the relevant finite-observation signatures in a given system were dominated by strongly non-Gaussian effects outside the quadratic regime, even when the projected Hessian correction is small.

At the spectroscopy level, the claim fails if the envelope increments do not reproduce the nonzero spectrum of Δ_{DV} , or if the H_0 -whitened Ky Fan values do not agree with the generalised-eigenvalue detectability spectrum.

Finally, the operator-space extension would fail as a physical theory if no meaningful linear fluctuation correction in BKM geometry could be identified for the relevant quantum nonequilibrium setting. The weighted transfer theorem would remain mathematically true, but it would lose physical force.

13.3 What the worked examples support, and what they do not

The worked examples play a calibrating role. They show that the distinction between raw visibility and backbone-normalised detectability remains informative across mechanistically distinct nonequilibrium models. In this empirical arc, the ribosome benchmark provides utility together with restraint, because the matched null shows that low-rank concentration in a small five-state system should not by itself be read as architecture-specific specialness. The kinesin benchmark is the stronger cross-domain recurrence case: it shows robust low-rank concentration at low load together with physically sensible broadening near stall under direct cycle competition. The chemotaxis benchmark is the larger-state-space stress test: it indicates that the same operator theory can remain informative beyond the smallest reduced spaces.

At the same time, the examples do not yet prove a universal low-rank law definitively. They support repeated utility of the two-spectrum framework, and they support cross-domain recurrence of the visibility geometry in the tested regimes, but they do not by themselves settle how much of the observed concentration is controlled by model class, operating regime, state-space dimension, or the number of effective irreversibility-injection channels. That wider question remains open.

Appendix I sharpens this point analytically. A homogeneous driven ring gives a clean counterexample to the naïve idea that few cycles should force low raw stable rank, while a finite-rank skew channel gives an exact bounded-complexity mechanism. Taken together, these results show that the relevant control variable is not bare cycle count by itself but the spectral localisation of the skew sector relative to the detailed-balance backbone.

Read laterally, the examples also sharpen the hidden-engine anatomy proved in the main text. What survives finite observation is organised by the same ladder in every case: skew source, detailed-balance transmission geometry, backbone-weighted visible signal, hidden return or hidden tail under rank constraint, and detectability after whitening.

13.4 What remains open

The first open direction is noisy observation-channel design. Section 6 gives a Gaussian shadow theorem for fixed linear noisy channels, but the optimisation problem over channels remains open. This is naturally an optimal experimental-design problem in the observed covariance geometry. Unlike deterministic subspace observation, the observation matrix enters nonlinearly in both the covariance defect and the backbone normalisation, so the design problem is no longer a direct Ky Fan optimisation. A related structural question is which backbone heterogeneities or projection mismatches force the detectability spectrum to broaden relative to the raw visibility spectrum.

The second open direction is the full nonlinear quantum bridge. The weighted operator-space

extension shows that the visibility geometry is ready once the fluctuation correction is identified, but that identification remains to be derived beyond the linear regime.

The third open direction is extension beyond finite-state Markov models. Recent work on thermodynamically consistent coarse graining across particle and field descriptions suggests that structure can survive substantial change of description [50]. At the same time, the need to coarse grain empirical densities and currents in continuous space makes clear that observation scale must remain part of the structure rather than an external nuisance [51, 52, 56]. A natural question is whether a useful finite-observation complexity invariant, perhaps weighted stable rank or a field-theoretic analogue, survives renormalised coarse graining.

The fourth open direction is the non-Markov and driven large-deviation frontier. Periodic level-2.5 theories and recent non-Markov self-interacting large-deviation results suggest that empirical-measure and current structures persist well beyond the time-homogeneous Markov class treated here [53, 54]. It is therefore natural to ask whether an analogue of the weighted signal operator and its finite-observation envelope exists for memoryful or periodically driven systems.

The fifth open direction is implementation. The envelope hierarchy suggests a direct route to irreversibility spectroscopy: estimate the singular front of the nonequilibrium signal from controlled families of observables or from designed operator-space compressions. A natural next step is to study how robustly the leading singular front of A_{DV} can be estimated from finite data under partial observation. A closely related frontier is active or response-based detection, where recent finite-frequency fluctuation-response bounds indicate that nonequilibrium detectability can also be organised in the frequency domain [55]. In the background sits the broader tension between spectral structure and scalar uncertainty bounds, now familiar from the thermodynamic uncertainty-relation literature [24, 25].

13.5 Conclusion and outlook

The main results of this paper are finite-dimensional and specific to the Donsker–Varadhan correction, but their significance is broader. The central lesson is that finite observation does not merely discard information; it organises nonequilibrium structure into an exact visible geometry. In the present setting that geometry is carried first by the weighted signal operator $A_{DV} = \widehat{J}H_0^{-1/2}$, which governs raw visibility, and second by its backbone-whitened detectability counterpart, which governs Gaussian distinguishability relative to the detailed-balance backbone. This yields an exact theory of nonequilibrium visibility, observer design, detectability, and spectral readout under explicit rank constraints.

That exact core is already enough to support a strong conclusion, but the conclusion must be stated at the correct level. We do not merely have a collection of examples with similar numerics. We have identified a common mathematical form. Across the finite-dimensional class treated here, nonequilibrium fluctuation structure is governed by a positive weighted correction operator, its weighted signal factorisation, its Ky Fan envelope hierarchy, and a second backbone-normalised spectrum for detectability. Within the orthogonal-observer model, the envelope hierarchy is the exact rank-constrained visibility law generated by A_{DV} . It determines the singular-energy profile $\frac{1}{4}s_k(A_{DV})^2$, and within the finite irreducible Markov Gaussian setting treated here this weighted geometry is the organising structure.

The sharper question is where a low-rank regime sits inside that exact geometry. The worked examples show recurrent concentration, but Appendix I makes clear that low-rank concentration is not a theorem of the whole class. A homogeneous driven ring has raw

weighted stable rank

$$\text{sr}(A_{DV}) = \frac{n-2}{2 \cos^2(\pi/n)} \sim \frac{n}{2}, \quad (13.1)$$

so topological simplicity by itself does not force raw spectral simplicity. By contrast, if the reduced skew channel has rank at most $2r$, then both raw visibility and backbone-whitened detectability have rank at most $2r$. This isolates one exact bounded-complexity mechanism. The recurrent concentration seen in the biological benchmarks is consistent with a regime in which the skew channel is effectively localised to a small number of irreversibility-injection directions, although the paper does not prove that benchmark by benchmark. The natural conclusion is therefore not that every finite nonequilibrium system is low-rank, but that the low-rank behaviour seen in the worked examples belongs to a candidate regime inside the exact weighted geometry.

We therefore state the conclusion in its disciplined form: this work identifies the exact weighted geometry of finite observation and separates two optimisation problems that are often conflated, raw visibility and backbone-normalised detectability. It also provides evidence for a candidate low-rank universality regime within that geometry.

What recurs across model classes is not merely the fact of compression. It is the same organising architecture: skew source, detailed-balance transmission geometry, weighted visible signal, hidden tail under rank constraint, and backbone-whitened detectability. The worked examples, together with the analytic boundary note, show that the open problem is now sharply posed. The relevant control variable is not bare cycle count alone. It is how broadly the skew sector is transmitted across the backbone modes.

That point matters for the next stage of the programme. The task is no longer to ask vaguely whether low-rank concentration appears again. The task is to determine which structural features keep the transmitted skew sector spectrally localised, which features delocalise it, and how the basin should be organised. Upgraded matched nulls, controlled scaling families, and explicit breakdown classes are the natural next tests.

The frontier is therefore clear. The finite-dimensional Gaussian orthogonal-observer theory developed here now has an exact structural form. The low-rank regime is strongly suggested in a nontrivial empirical arc, but it is not yet the whole class. If the basin analysis now underway confirms that the same low-rank regime persists under broader structural sweeps, then finite observation will not merely be a practical limitation on what can be seen. It will pick out a genuine organising regime of nonequilibrium structure.

A Reduced-coordinate identities

This appendix records the reduced-coordinate identities underlying the bridge theorem. The aim is to place the key algebra in one location so that the main text remains readable. It is the same reduced conductance-and-skew decomposition used in Section 3, written here in the explicit (x, y) variables of the envelope expansion.

A.1 Reduced tangent coordinates

Let

$$T_\pi = \{u \in \mathbb{R}^n : \langle u, e_\pi \rangle = 0\}, \quad e_\pi = (\sqrt{\pi_1}, \dots, \sqrt{\pi_n})^T, \quad (\text{A.1})$$

and let $C \in \mathbb{R}^{n \times (n-1)}$ be an orthonormal frame of T_π , so that

$$C^T C = I_{n-1}, \quad C C^T = P_\pi. \quad (\text{A.2})$$

Every tangent fluctuation is written uniquely as $u = Cx$ with $x \in \mathbb{R}^{n-1}$.

In these coordinates, the Fisher-conjugated generator reduces to

$$\widehat{K} = C^T D_\pi^{1/2} Q D_\pi^{-1/2} C, \quad \widehat{K} = \widehat{G} + \widehat{J}, \quad (\text{A.3})$$

with

$$\widehat{G} = \frac{\widehat{K} + \widehat{K}^T}{2}, \quad \widehat{J} = \frac{\widehat{K} - \widehat{K}^T}{2}. \quad (\text{A.4})$$

The canonical detailed-balance reference Hessian on the same reduced space is

$$H_0 = -\frac{1}{2} \widehat{G}. \quad (\text{A.5})$$

A.2 Reduced envelope structure

Write the Donsker–Varadhan variational formula as [1–3]

$$I_{DV}(\mu) = \sup_h \left(- \sum_{i,j} \mu_i q_{ij} (e^{h_j - h_i} - 1) \right). \quad (\text{A.6})$$

Choose reduced Fisher coordinates $x \in \mathbb{R}^{n-1}$ and reduced logarithmic coordinates $y \in \mathbb{R}^{n-1}$ by

$$\mu(x) = \pi + D_\pi^{1/2} Cx + O(|x|^2), \quad h(y) = D_\pi^{-1/2} Cy. \quad (\text{A.7})$$

Expanding the variational integrand to second order gives

$$\mathcal{L}(x, y) = y^T \widehat{G} y - x^T \widehat{K} y + O(|(x, y)|^3), \quad \widehat{K} = \widehat{G} + \widehat{J}. \quad (\text{A.8})$$

Because $y^T \widehat{J} y = 0$, the conductance block depends only on $\widehat{G}$. Hence

$$M_Q = M_G = -\widehat{G}. \quad (\text{A.9})$$

Setting $S := -\widehat{G} \succ 0$, one has

$$\mathcal{L}(x, y) = -y^T S y + y^T (S + \widehat{J}) x + O(|(x, y)|^3). \quad (\text{A.10})$$

Completing the square in y yields

$$I_{DV}(\mu(x)) = \frac{1}{4} x^T (S + \widehat{J})^T S^{-1} (S + \widehat{J}) x + O(|x|^3) = \frac{1}{4} x^T (S + \widehat{J} S^{-1} \widehat{J}^T) x + O(|x|^3). \quad (\text{A.11})$$

Therefore

$$H_{DV} = \frac{1}{2} S + \frac{1}{2} \widehat{J} S^{-1} \widehat{J}^T. \quad (\text{A.12})$$

Since $H_0 = \frac{1}{2} S$, one gets $H_0^{-1} = 2S^{-1}$, and hence

$$H_{DV} = H_0 + \frac{1}{4} \widehat{J} H_0^{-1} \widehat{J}^T. \quad (\text{A.13})$$

Equivalently,

$$\Delta_{DV} = H_{DV} - H_0 = \frac{1}{4} \widehat{J} H_0^{-1} \widehat{J}^T. \quad (\text{A.14})$$

A.3 Absence of mixed terms

The quadraticity theorem can be read as a cancellation statement. A generic two-sector reduced envelope would allow mixed terms involving both the dissipative and skew sectors in the Schur complement. The identity $M_Q = M_G$ removes those terms because the conductance arises from the quadratic form in the auxiliary variable, and a skew matrix has zero quadratic form. The result is that the correction depends quadratically on the skew sector and is weighted only by the inverse detailed-balance backbone. This is the algebraic source of the weighted signal form

$$\Delta_{DV} = \frac{1}{4} A_{DV} A_{DV}^T, \quad A_{DV} = \widehat{J} H_0^{-1/2}. \quad (\text{A.15})$$

A.4 Observed block formula

Let Π be a rank- d orthogonal projector whose range $U = \text{Ran } \Pi$ is H_0 -invariant, and choose a basis adapted to the splitting

$$\mathbb{R}^{n-1} = U \oplus U^\perp, \quad U = \text{Ran } \Pi, \quad (\text{A.16})$$

so that H_0 is block diagonal:

$$H_0 = \begin{pmatrix} H_{0,S} & 0 \\ 0 & H_{0,F} \end{pmatrix}, \quad \widehat{J} = \begin{pmatrix} J_{SS} & J_{SF} \\ J_{FS} & J_{FF} \end{pmatrix}. \quad (\text{A.17})$$

Then

$$\Delta_{DV} = \frac{1}{4} \widehat{J} H_0^{-1} \widehat{J}^T = \frac{1}{4} \begin{pmatrix} J_{SS} & J_{SF} \\ J_{FS} & J_{FF} \end{pmatrix} \begin{pmatrix} H_{0,S}^{-1} & 0 \\ 0 & H_{0,F}^{-1} \end{pmatrix} \begin{pmatrix} J_{SS}^T & J_{FS}^T \\ J_{SF}^T & J_{FF}^T \end{pmatrix}. \quad (\text{A.18})$$

Its slow block is therefore

$$(\Delta_{DV})_{SS} = \frac{1}{4} J_{SS} H_{0,S}^{-1} J_{SS}^T + \frac{1}{4} J_{SF} H_{0,F}^{-1} J_{SF}^T. \quad (\text{A.19})$$

This is the converter formula used in Section 3.

B Linear algebra of weighted signal operators

This appendix collects the linear algebra used repeatedly in Sections 4 to 9.

B.1 Visibility interval

Let A be any finite matrix with singular values

$$s_1(A) \geq s_2(A) \geq \cdots \geq s_m(A) \geq 0. \quad (\text{B.1})$$

For a rank- d orthogonal projector Π , define

$$V_\Pi(A) := \|\Pi A\|_F^2. \quad (\text{B.2})$$

Then Ky Fan gives

$$\sup_{\text{rank } \Pi = d} V_\Pi(A) = \sum_{k=1}^d s_k(A)^2, \quad (\text{B.3})$$

and the minimum is

$$\inf_{\text{rank } \Pi = d} V_\Pi(A) = \sum_{k=m-d+1}^m s_k(A)^2, \quad (\text{B.4})$$

with the convention that zero singular values fill out the list if $\text{rank}(A) < m$.

Applied to

$$\Delta_{DV} = \frac{1}{4} A_{DV} A_{DV}^T, \quad (\text{B.5})$$

this gives the full visibility interval for rank- d observation:

$$\frac{1}{4} \sum_{k=m-d+1}^m s_k(A_{DV})^2 \leq V_\Pi^{DV} \leq \frac{1}{4} \sum_{k=1}^d s_k(A_{DV})^2. \quad (\text{B.6})$$

B.2 Optimal hidden tail

The optimal hidden residual after rank- d observation is

$$\frac{1}{4} \sum_{k>d} s_k(A)^2. \quad (\text{B.7})$$

This is the Eckart-Young residual in Frobenius norm, expressed at the signal level. In the present paper it is the irrecoverable hidden nonequilibrium tail under the best rank- d observer.

B.3 Stable-rank obstruction

For any nonzero matrix A ,

$$\text{sr}(A) = \frac{\|A\|_F^2}{\|A\|_{op}^2}. \quad (\text{B.8})$$

Then

$$\sum_{k=1}^d s_k(A)^2 \leq d s_1(A)^2 = d \frac{\|A\|_F^2}{\text{sr}(A)}. \quad (\text{B.9})$$

Dividing by $\|A\|_F^2$ gives the basic stable-rank obstruction

$$\frac{\sum_{k=1}^d s_k(A)^2}{\sum_k s_k(A)^2} \leq \frac{d}{\text{sr}(A)}. \quad (\text{B.10})$$

B.4 Direct-sum envelope law

If

$$A = A_1 \oplus A_2 \quad (\text{B.11})$$

is block diagonal with respect to an orthogonal splitting, then the rank- d envelope satisfies

$$K_d(A_1 \oplus A_2) = \max_{d_1+d_2=d} (K_{d_1}(A_1) + K_{d_2}(A_2)), \quad (\text{B.12})$$

where $K_d(A)$ denotes the Ky Fan top-singular-energy envelope [32]. This expresses optimal budget allocation across independent sectors.

Proof. The singular values of $A_1 \oplus A_2$ are the multiset union of the singular values of A_1 and A_2 . The optimal sum of the top d squared singular values is obtained by distributing the rank budget between the two blocks in the best possible way.

B.5 Comparison of weighted and unweighted envelopes

Let H_0 be positive definite with eigenvalue window

$$\mu_{\min} I \preceq H_0 \preceq \mu_{\max} I. \quad (\text{B.13})$$

Then

$$\mu_{\max}^{-1} I \preceq H_0^{-1} \preceq \mu_{\min}^{-1} I. \quad (\text{B.14})$$

Hence for

$$A_{DV} = \widehat{J} H_0^{-1/2}, \quad \Delta_{DV} = \frac{1}{4} A_{DV} A_{DV}^T, \quad (\text{B.15})$$

one has

$$\frac{1}{4\mu_{\max}} \widehat{J} \widehat{J}^T \preceq \Delta_{DV} \preceq \frac{1}{4\mu_{\min}} \widehat{J} \widehat{J}^T. \quad (\text{B.16})$$

Taking optimal compressed traces gives

$$\frac{1}{4\mu_{\max}} K_d(\widehat{J}) \leq K_d^{DV} \leq \frac{1}{4\mu_{\min}} K_d(\widehat{J}). \quad (\text{B.17})$$

C Random projector calculations

C.1 Expectation of a Haar projector

Let Π be Haar-uniform among rank- d orthogonal projectors on $\mathbb{R}^m$. Then orthogonal invariance implies

$$\mathbb{E}[\Pi] = cI \tag{C.1}$$

for some scalar c . Taking traces gives

$$cm = \mathbb{E}[\text{tr } \Pi] = d, \tag{C.2}$$

hence

$$\mathbb{E}[\Pi] = \frac{d}{m}I. \tag{C.3}$$

C.2 Expected retained signal

For any fixed matrix A ,

$$\mathbb{E}[\|\Pi A\|_F^2] = \mathbb{E}[\text{tr}(A^T \Pi A)] = \text{tr}(A^T \mathbb{E}[\Pi] A) = \frac{d}{m} \|A\|_F^2. \tag{C.4}$$

Applied to the weighted DV signal,

$$\mathbb{E}[V_\Pi^{DV}] = \frac{1}{4} \frac{d}{m} \|A_{DV}\|_F^2. \tag{C.5}$$

C.3 Expected directional overlap

For any fixed unit vector $u \in \mathbb{R}^m$,

$$\mathbb{E}[\|\Pi u\|^2] = u^T \mathbb{E}[\Pi] u = \frac{d}{m}. \tag{C.6}$$

Thus each singular direction of the signal is seen at the same average rate by a generic rank- d observer.

C.4 Remark on concentration

The present paper only needs the mean baseline. If desired, concentration of $\|\Pi u\|^2$ and of $\|\Pi A\|_F^2$ around their means can be obtained from standard concentration on the Grassmannian. We omit those estimates because they are not needed for the main results.

D Gaussian comparison formulas

This appendix records the centred Gaussian formulas used in Section 6.

D.1 KL divergence

For centred Gaussian laws on $\mathbb{R}^d$ with covariance matrices Σ_1, Σ_0 ,

$$D_{KL}(\mathcal{N}(0, \Sigma_1) \parallel \mathcal{N}(0, \Sigma_0)) = \frac{1}{2} \left[\text{tr}(\Sigma_0^{-1} \Sigma_1) - d - \log \det(\Sigma_0^{-1} \Sigma_1) \right]. \quad (\text{D.1})$$

In the shadow setting,

$$\Sigma_{DB, \Pi} = B_{\Pi}^{-1}, \quad \Sigma_{NESS, \Pi} = B_{\Pi}^{-1/2} (I + S_{\Pi})^{-1} B_{\Pi}^{-1/2}. \quad (\text{D.2})$$

Hence

$$\Sigma_{NESS, \Pi}^{-1} \Sigma_{DB, \Pi} = B_{\Pi}^{1/2} (I + S_{\Pi}) B_{\Pi}^{-1/2}, \quad (\text{D.3})$$

whose eigenvalues are $1 + \lambda_i$, where λ_i are the eigenvalues of S_{Π} . This yields

$$D_{KL}(\mathcal{N}_{DB, \Pi} \parallel \mathcal{N}_{NESS, \Pi}) = \frac{1}{2} \sum_i (\lambda_i - \log(1 + \lambda_i)). \quad (\text{D.4})$$

Likewise,

$$\Sigma_{DB, \Pi}^{-1} \Sigma_{NESS, \Pi} = B_{\Pi}^{1/2} (I + S_{\Pi})^{-1} B_{\Pi}^{-1/2}, \quad (\text{D.5})$$

whose eigenvalues are $(1 + \lambda_i)^{-1}$, giving

$$D_{KL}(\mathcal{N}_{NESS, \Pi} \parallel \mathcal{N}_{DB, \Pi}) = \frac{1}{2} \sum_i \left(\log(1 + \lambda_i) - \frac{\lambda_i}{1 + \lambda_i} \right). \quad (\text{D.6})$$

D.2 Hellinger distance

For centred Gaussians,

$$1 - H^2 = \frac{\det(\Sigma_1)^{1/4} \det(\Sigma_0)^{1/4}}{\det\left(\frac{1}{2}(\Sigma_1 + \Sigma_0)\right)^{1/2}}. \quad (\text{D.7})$$

Substituting the detailed-balance and nonequilibrium observable covariances and factoring out $B_{\Pi}^{-1/2}$ gives

$$H^2(\mathcal{N}_{DB, \Pi}, \mathcal{N}_{NESS, \Pi}) = 1 - \frac{\det(I + S_{\Pi})^{1/4}}{\det(I + \frac{1}{2} S_{\Pi})^{1/2}}. \quad (\text{D.8})$$

D.3 Pinsker and Bayes error

Pinsker gives

$$\text{TV}(P, Q) \leq \sqrt{\frac{1}{2} D_{KL}(P \| Q)}. \quad (\text{D.9})$$

For equal priors, the optimal Bayes classification error satisfies

$$P_e^* = \frac{1 - \text{TV}(P, Q)}{2}. \quad (\text{D.10})$$

Combining these with the KL bound

$$D_{KL}(\mathcal{N}_{DB, \Pi} \| \mathcal{N}_{NESS, \Pi}) \leq \frac{1}{4} \|S_{\Pi}\|_F^2 \quad (\text{D.11})$$

yields the total-variation and Bayes-error estimates stated in Section 6.

E Worked details for the literature examples

This appendix records the precise interpretive use of the external literature discussed in Sections 11 and 12.

E.1 Four-state two-cycle benchmark

The worked benchmark of Section 11 uses the four-state, five-link, two-cycle network from Appendix E, Table IV of van der Meer, Ertel, and Seifert [37]. Two roles are separated carefully.

First, the paper uses the generator as a compact multicyclic test of the bridge, envelope, and shadow theorems.

Second, to connect directly to the source observation pattern, the paper defines the edge-induced rank-one sector

$$b_{23} = D_{\pi}^{-1/2}(e_2 - e_3), \quad \Pi_{\text{obs}} = \frac{C^T b_{23} (C^T b_{23})^T}{\|C^T b_{23}\|^2}. \quad (\text{E.1})$$

This is an orthogonal surrogate for the physical access pattern observe only the $2 \leftrightarrow 3$ link. It is not a claim that the waiting-time estimator of [37] itself is literally an orthogonal projector on reduced Fisher space. The point is more disciplined: once one restricts to the theorem class of the present paper, Π_{obs} is the canonical one-dimensional sector induced by the observed link within that class.

The numerical comparison in Section 11 therefore separates three quantities on the same benchmark family: raw retained nonequilibrium signal, backbone-normalised detectability, and the gap from the observed-link sector to the unconstrained optimal rank-one observers. The short forcing sweep is performed for this same induced sector, not for a re-optimised observer.

E.2 Estimator papers

In the finite-resolution and missing-data literature, the central practical question is how well one can estimate entropy production from restricted observations [5, 6, 57]. The present paper does not replace those estimators. It inserts a prior geometric question: how much of the weighted nonequilibrium signal is observable at all in the chosen sector? The retained fraction and the alignment gain are the relevant pre-estimation quantities.

E.3 Lumping and milestoning

The milestoning work of Blom et al. [13] shows that reprocessed coarse observation can improve dissipation estimation. The present theory interprets that phenomenon as better alignment with the singular front of the weighted signal, rather than as a paradox in which a coarser observation outperforms a finer one.

E.4 Unresolved events

Harunari’s unresolved-event framework [7] studies hidden irreversibility at the event-statistics level. The present paper addresses the same obstruction at the Gaussian Hessian level. These should be seen as complementary rather than competing approaches.

E.5 Observation-dependent thermodynamics

Rubino, Brukner, and Manzano [19] and Pernambuco and Céleri [20] show that thermodynamic content depends on access and resolution. The present paper adds the spectral law for how much of the weighted fluctuation correction survives finite-rank observation in the corresponding linearised setting.

E.6 BKM and de Bruijn geometry

Lacerda, Bettmann, and Goold [21], Bettmann and Goold [22], and Wirth [23] provide the metric and entropy-decay background showing that BKM geometry and de Bruijn identities are the correct neighbourhood for the operator-space continuation of the theory.

F Operator-space technicalities

This appendix records the Euclideanisation used in Section 9.

Let M be a positive-definite metric on a finite-dimensional space V . If Π is M -orthogonal, then

$$\tilde{\Pi} := M^{1/2} \Pi M^{-1/2} \tag{F.1}$$

is an ordinary Euclidean orthogonal projector.

Indeed,

$$\tilde{\Pi}^2 = M^{1/2} \Pi^2 M^{-1/2} = \tilde{\Pi}. \quad (\text{F.2})$$

Moreover,

$$\tilde{\Pi}^\dagger = (M^{-1/2})^\dagger \Pi^\dagger (M^{1/2})^\dagger = M^{-1/2} \Pi^\dagger M^{1/2}. \quad (\text{F.3})$$

Since Π is M -self-adjoint,

$$\Pi^\dagger M = M \Pi, \quad \Pi^\dagger = M \Pi M^{-1}, \quad (\text{F.4})$$

so

$$\tilde{\Pi}^\dagger = M^{-1/2} M \Pi M^{-1} M^{1/2} = M^{1/2} \Pi M^{-1/2} = \tilde{\Pi}. \quad (\text{F.5})$$

Thus $\tilde{\Pi}$ is Euclidean orthogonal.

If $\Delta = BB^\sharp$, define

$$\tilde{B} = M^{1/2} B M^{-1/2}. \quad (\text{F.6})$$

Then the Euclideanised correction is

$$M^{1/2} \Delta M^{-1/2} = M^{1/2} B B^\sharp M^{-1/2} = M^{1/2} B M^{-1} B^\dagger M^{1/2} = \tilde{B} \tilde{B}^\dagger. \quad (\text{F.7})$$

Hence the weighted projector problem becomes an ordinary Euclidean Gram problem after conjugation by $M^{1/2}$.

G Interpretation and boundaries

The finite-observation theory proved in the main text can be placed inside a slightly broader interpretive viewpoint, provided that the proof-status boundary is kept sharp.

The basic distinction is between structural invisibility and practical weakness. Structural invisibility is vertical: degrees of freedom are intrinsically unobservable because they lie in fibres that are annihilated by the observation map or by the first positive descended object. Practical weakness is horizontal: after the invisible directions are quotiented out, the descended positive object may still be spectrally weak or anisotropic on the active quotient. In the present paper, the relevant descended object is the Donsker–Varadhan correction Δ_{DV} , and the active quotient is the reduced Euclidean Fisher space fixed in Sections 2 and 3.

A small observed correction can arise for conceptually different reasons. It may arise because the observation genuinely kills the relevant directions, because the quotient signal is low-rank but aligned away from the chosen observer, because the quotient signal is broad and weak, or because the observable sector is simply the wrong thin slice through an otherwise informative quotient geometry. These are not the same mechanism.

The work proves one finite-dimensional branch of this broader viewpoint: once the finite-state Markov Gaussian setting is fixed, the nonequilibrium correction descends to a weighted positive object, its finite-rank visible mass is solved by the envelope hierarchy, and its backbone-normalised distinguishability is solved by the shadow spectrum. That is the scope established here.

The finite-observation picture can fail by structural fibre collapse, by quotient-rank or quotient-anisotropy failure, by mismatch between the chosen thin observer and the dominant quotient directions, or by genuinely nonlinear aliasing beyond the quadratic regime. These mechanisms should be kept distinct, rather than collapsed into generic talk of “information loss.”

H Interpretive synthesis for the worked examples

The theorems of Sections 3 to 7 admit a compact synthesis that is useful when reading the worked benchmarks.

The algebraic source of the nonequilibrium correction is the reduced skew channel $\widehat{J}$. The detailed-balance backbone H_0 does not generate the correction by itself, but it supplies the transmission geometry through which the skew source becomes visible. The natural signal object is therefore the backbone-weighted operator

$$A_{DV} = \widehat{J}H_0^{-1/2}, \quad \Delta_{DV} = \frac{1}{4}A_{DV}A_{DV}^T. \quad (\text{H.1})$$

Finite observation compresses this weighted signal, not the full generator.

When an observed sector is H_0 -adapted, the block formula gives a second useful layer of anatomy:

$$(\Delta_{DV})_{SS} = \frac{1}{4}J_{SS}H_{0,S}^{-1}J_{SS}^T + \frac{1}{4}J_{SF}H_{0,F}^{-1}J_{SF}^T. \quad (\text{H.2})$$

The first term is the intrinsic observed correction. The second is the hidden mixed-return channel from the complementary sector, throttled by the inverse stiffness of the fast detailed-balance block. Thus hidden nonequilibrium is not a vague reservoir. On adapted sectors it has a return mechanism.

At finite observer rank d , the optimally hidden budget is the singular-energy tail of the weighted signal operator,

$$K_{\text{tot}}^{DV} - K_d^{DV} = \frac{1}{4} \sum_{k>d} s_k (A_{DV})^2, \quad (\text{H.3})$$

so the envelope hierarchy is literally a spectroscopy of what finite observation cannot recover at the chosen rank. After whitening by the backbone, the relevant distinguishability object is the shadow or detectability spectrum. This is a second spectrum. It agrees with raw visibility only in special cases and should not be identified with it.

Taken together, the worked examples can be read through one recurring ladder: skew source, weighted signal, hidden return or hidden tail, and whitened detectability. This paper develops that ladder only for the finite-observation setting treated in the main text. Broader extensions are left outside the present scope.

I An analytic note on the scope of the low-rank universality claim

This appendix records two exact analytic complements to the worked examples. The first gives a clean boundary family showing that low raw stable rank is not forced by topological simplicity alone. The second gives an exact bounded-complexity mechanism. Together they clarify what the universality claim of the main text does and does not assert.

I.1 Homogeneous driven ring: broad raw visibility and bounded detectability

Consider the nearest-neighbour n -state ring with uniform stationary law and homogeneous clockwise versus anticlockwise bias. On the mean-zero reduced Euclidean Fisher space, the

detailed-balance backbone and reduced skew channel are, up to positive constants that cancel from stable-rank ratios,

$$H_0 \propto L_n := 2I - S - S^{-1}, \quad \widehat{J} \propto K_n := S - S^{-1}, \quad (\text{I.1})$$

where S is the cyclic shift. Both operators are circulant, so they diagonalise in the Fourier basis. For the nonzero Fourier mode

$$q_m = \frac{2\pi m}{n}, \quad m = 1, \dots, n-1, \quad (\text{I.2})$$

one has

$$\lambda_{H_0}(m) \propto 1 - \cos q_m, \quad |\lambda_{\widehat{J}}(m)| \propto |\sin q_m|. \quad (\text{I.3})$$

Proposition I.1 (Homogeneous driven ring: raw stable rank grows linearly). *For the homogeneous driven ring,*

$$A_{DV} = \widehat{J}H_0^{-1/2} \quad (\text{I.4})$$

has squared singular values proportional to

$$\sigma_m(A_{DV})^2 \propto \frac{\sin^2 q_m}{1 - \cos q_m} = 1 + \cos q_m = 2 \cos^2\left(\frac{\pi m}{n}\right), \quad m = 1, \dots, n-1. \quad (\text{I.5})$$

Consequently,

$$\text{sr}(A_{DV}) = \frac{n-2}{2 \cos^2(\pi/n)} \sim \frac{n}{2}. \quad (\text{I.6})$$

In particular, a one-cycle family can have raw weighted stable rank growing linearly with state-space dimension.

Proof. The Fourier diagonalisation gives the displayed singular-value formula immediately. Since

$$\sum_{m=1}^{n-1} \cos^2\left(\frac{\pi m}{n}\right) = \frac{n-2}{2} \quad (\text{I.7})$$

and the maximal value is attained at $m = 1$ and equals $\cos^2(\pi/n)$, the stated stable-rank identity follows.

Proposition I.2 (Homogeneous driven ring: detectability complexity stays bounded). *For the same family,*

$$C_{DV} = H_0^{-1/2} \widehat{J} H_0^{-1/2} \quad (\text{I.8})$$

has singular values proportional to

$$\sigma_m(C_{DV}) \propto \frac{|\sin q_m|}{1 - \cos q_m} = \left| \cot\left(\frac{\pi m}{n}\right) \right|, \quad m = 1, \dots, n-1. \quad (\text{I.9})$$

Hence

$$\text{sr}(C_{DV}) = \frac{(n-1)(n-2)}{3 \cot^2(\pi/n)} \longrightarrow \frac{\pi^2}{3} \quad (n \rightarrow \infty). \quad (\text{I.10})$$

Thus the same homogeneous ring that is broad in raw visibility remains bounded in backbone-whitened detectability complexity.

Proof. The Fourier formula above gives the singular values. Using the classical identity

$$\sum_{m=1}^{n-1} \cot^2\left(\frac{\pi m}{n}\right) = \frac{(n-1)(n-2)}{3}, \quad (\text{I.11})$$

and the fact that the largest value is $\cot^2(\pi/n)$, we obtain the displayed formula. The limit follows from $\cot(\pi/n) \sim n/\pi$.

Propositions I.1 and I.2 show that the two spectra of the paper can behave very differently even in the simplest analytically transparent family. A homogeneous one-cycle system can be broad in raw visible signal while remaining bounded in detectability complexity. The reason is that inverse-backbone weighting amplifies the lowest-wavenumber backbone modes and suppresses the high-wavenumber tail strongly enough that the detectability stable rank stays bounded even as the raw-visibility stable rank grows linearly. The relevant control variable is therefore not bare cycle count by itself.

I.2 Finite-rank skew channels as an exact bounded-complexity mechanism

The next observation isolates a clean mechanism that does force bounded complexity.

Lemma I.3 (*Finite-rank skew channels force bounded complexity*). *If the reduced skew channel satisfies*

$$\text{rank}(\hat{J}) \leq 2r, \quad (\text{I.12})$$

then

$$\text{rank}(A_{DV}) = \text{rank}(C_{DV}) \leq 2r, \quad (\text{I.13})$$

and therefore

$$\text{sr}(A_{DV}) \leq 2r, \quad \text{sr}(C_{DV}) \leq 2r. \quad (\text{I.14})$$

Proof. Since H_0 is positive definite on the reduced space, both $H_0^{-1/2}$ and left multiplication by $H_0^{-1/2}$ are invertible. Hence

$$\text{rank}(A_{DV}) = \text{rank}(\hat{J}H_0^{-1/2}) = \text{rank}(\hat{J}), \quad \text{rank}(C_{DV}) = \text{rank}(H_0^{-1/2}\hat{J}H_0^{-1/2}) = \text{rank}(\hat{J}). \quad (\text{I.15})$$

The stable-rank bounds follow from $\text{sr}(X) \leq \text{rank}(X)$ for every finite matrix X .

A useful special case is when the skew channel is generated by a fixed number of elementary antisymmetric couplings,

$$\hat{J} = \sum_{a=1}^r \gamma_a (u_a v_a^T - v_a u_a^T), \quad (\text{I.16})$$

with orthonormal or merely linearly independent pair directions u_a, v_a . Each summand has rank at most 2, so the hypothesis of Lemma I.3 holds automatically. This gives an exact mechanism by which localised irreversibility injection keeps both visibility and detectability spectrally compressible, uniformly in ambient state-space dimension.

The conclusion for the universality discussion is precise. The exact weighted geometry of finite observation is common across the full class treated in the main text, but low-rank concentration belongs to a narrower regime inside that geometry. The analytic boundary

family above shows that low rank is not forced by topological simplicity alone. The finite-rank skew lemma shows that one sufficient bounded-complexity mechanism is spectral localisation of the skew sector. This is the structural axis that the next empirical appendix programme should probe.

Acknowledgements

No external funding was received. All work was produced independently without affiliation. Author’s email is contact@nomogenetics.com.

References

- [1] M. D. Donsker and S. R. S. Varadhan, Asymptotic evaluation of certain Markov process expectations for large time. I, *Communications on Pure and Applied Mathematics* **28** (1975), 1–47.
- [2] M. D. Donsker and S. R. S. Varadhan, Asymptotic evaluation of certain Markov process expectations for large time. II, *Communications on Pure and Applied Mathematics* **28** (1975), 279–301.
- [3] H. Touchette, The large deviation approach to statistical mechanics, *Physics Reports* **478** (2009), 1–69.
- [4] J. L. Lebowitz and H. Spohn, A Gallavotti–Cohen-type symmetry in the large deviation functional for stochastic dynamics, *Journal of Statistical Physics* **95** (1999), 333–365.
- [5] M. Baiesi, Y. Nishiyama, and G. Falasco, Effective estimation of entropy production with lacking data, preprint, 2023.
- [6] L. Fritz, M. Ertel, and U. Seifert, Entropy estimation for partially accessible Markov networks based on imperfect observations: role of finite resolution and finite statistics, preprint, 2024.
- [7] P. E. Harunari, Uncovering nonequilibrium from unresolved events, preprint, 2024.
- [8] R. van der Meer and K. Saito, Thermodynamic bounds and error correction for faulty coarse graining, preprint, 2025.
- [9] B. Ertel and U. Seifert, An estimator of entropy production for partially accessible Markov networks based on the observation of blurred transitions, *Physical Review E* **111** (2025), 044106.
- [10] J. Ehrich, Tightest bound on hidden entropy production from partially observed dynamics, *Journal of Statistical Mechanics: Theory and Experiment* (2021), 083214.
- [11] J. Degünther, J. van der Meer, and U. Seifert, Fluctuating entropy production on the coarse-grained level: inference and localization of irreversibility, *Physical Review Research* **6** (2024), 023175.
- [12] M. Esposito, Stochastic thermodynamics under coarse graining, *Physical Review E* **85** (2012), 041125.
- [13] K. Blom, H. Song, S. Vouga, R. Godec, and D. E. Makarov, Milestoning estimators of dissipation in systems observed at a coarse resolution, preprint, 2025.
- [14] R. Bao, N. Ohga, and S. Ito, Measuring irreversibility by counting: a random coarse-graining framework, preprint, 2025.
- [15] T. R. Gingrich, G. M. Rotskoff, and J. M. Horowitz, Inferring dissipation from current fluctuations, *Journal of Physics A: Mathematical and Theoretical* **50** (2017), 184004.

- [16] Q. Yu, M. P. Leighton, and C. W. Lynn, Coarse-graining dynamics to maximize irreversibility, preprint, 2025.
- [17] G. Teza, A. L. Stella, and T. GrandPre, Coarse-graining via lumping: exact calculations and fundamental limitations, preprint, 2025.
- [18] G. Teza and A. L. Stella, Exact coarse graining preserves entropy production out of equilibrium, *Physical Review Letters* **125** (2020), 110601.
- [19] G. Rubino, Č. Brukner, and G. Manzano, Coarse-grained quantum thermodynamics: observation-dependent quantities, observation-independent laws, preprint, 2025.
- [20] J. Pernambuco and L. C. Céleri, Geometry of restricted information: the case of quantum thermodynamics, preprint, 2026.
- [21] T. Lacerda, N. Bettmann, and J. Goold, Information geometry of transitions between quantum nonequilibrium steady states, preprint, 2025.
- [22] N. Bettmann and J. Goold, Information geometry approach to quantum stochastic thermodynamics, preprint, 2025.
- [23] F. Wirth, Exponential relative entropy decay along quantum Markov semigroups, preprint, 2025.
- [24] A. C. Barato and U. Seifert, Thermodynamic uncertainty relation for biomolecular processes, *Phys. Rev. Lett.* **114** (2015), 158101.
- [25] J. M. Horowitz and T. R. Gingrich, Thermodynamic uncertainty relations constrain non-equilibrium fluctuations, *Nat. Phys.* **16** (2020), 15–20.
- [26] J. Gu, Counterexamples to the conjectured ordering between the waiting-time bound and the thermodynamic uncertainty bound on entropy production, arXiv:2601.04039 (2026).
- [27] S. Ito, Stochastic thermodynamic interpretation of information geometry, *Phys. Rev. Lett.* **121** (2018), 030605.
- [28] G. Wolfer and S. Watanabe, Information geometry of reversible Markov chains, *Information Geometry* **4** (2021), 393–433.
- [29] C. Maes, K. Netočný, and B. Wynants, Steady state statistics of driven diffusions, *Physica A* **387** (2008), 2675–2689.
- [30] J. Schnakenberg, Network theory of microscopic and macroscopic behavior of master equation systems, *Rev. Mod. Phys.* **48** (1976), 571–585.
- [31] D. Andrieux and P. Gaspard, Fluctuation theorem for currents and Schnakenberg network theory, *J. Stat. Phys.* **127** (2007), 107–131.
- [32] K. Fan, Maximum properties and inequalities for the eigenvalues of completely continuous operators, *Proceedings of the National Academy of Sciences* **37** (1951), 760–766.
- [33] D. Petz, Monotone metrics on matrix spaces, *Linear Algebra and its Applications* **244** (1996), 81–96.
- [34] R. Bhatia, *Matrix Analysis*, Graduate Texts in Mathematics Vol. 169, Springer, New York, 1997.
- [35] M. Rudelson and R. Vershynin, Sampling from large matrices: An approach through geometric functional analysis, *J. ACM* **54** (2007), Article 21.
- [36] É. Roldán and J. M. R. Parrondo, Entropy production and Kullback–Leibler divergence between stationary trajectories of discrete systems, *Physical Review E* **85** (2012), 031129.

- [37] J. van der Meer, B. Ertel, and U. Seifert, Thermodynamic inference in partially accessible Markov networks: A unifying perspective from transition-based waiting time distributions, *arXiv:2203.12020*, 2022.
- [38] K. Banerjee, A. B. Kolomeisky, and O. A. Igoshin, Elucidating interplay of speed and accuracy in biological error correction, *Proceedings of the National Academy of Sciences of the United States of America* **114** (2017), 5183–5188.
- [39] F. Mosam and E. De Giuli, Heterogeneity dominates irreversibility in random Markov models, *arXiv:2602.04905* (2026).
- [40] S. Liepelt and R. Lipowsky, Kinesin’s network of chemomechanical motor cycles, *Physical Review Letters* **98** (2007), 258102.
- [41] S. Liepelt and R. Lipowsky, EPAPS supplementary material for “Kinesin’s network of chemomechanical motor cycles,” supplementary document to *Physical Review Letters* **98** (2007), 258102.
- [42] P. Sartori, L. Granger, C. F. Lee, and J. M. Horowitz, Thermodynamic costs of information processing in sensory adaptation, *PLoS Computational Biology* **10** (2014), e1003974.
- [43] E. Meyberg, J. Degünther, and U. Seifert, Entropy production from waiting-time distributions for overdamped Langevin dynamics, *Journal of Physics A: Mathematical and Theoretical* **57** (2024), 25LT01.
- [44] C. Maes, Frenesy: time-symmetric dynamical activity in nonequilibria, *Physics Reports* **850** (2020), 1–33.
- [45] R. Chetrite and H. Touchette, Nonequilibrium Markov processes conditioned on large deviations, *Annales Henri Poincaré* **16** (2015), 2005–2057.
- [46] S. Chatterjee, Spectral gap of nonreversible Markov chains, *The Annals of Applied Probability* **35** (2025), 2644–2677.
- [47] J.-M. Souriau, *Structure of Dynamical Systems: A Symplectic View of Physics*, Birkhäuser, Boston, 1997.
- [48] F. Barbaresco, Geometric theory of heat from Souriau Lie groups thermodynamics and Koszul Hessian geometry: applications in information geometry for exponential families, *Entropy* **18** (2016), 386.
- [49] F. Barbaresco, Higher order geometric theory of information and heat based on poly-symplectic geometry of Souriau Lie groups thermodynamics and their contextures, *Entropy* **20** (2018), 840.
- [50] S. Mohite and H. Rieger, Thermodynamically consistent coarse-graining: from interacting particles to fields via second quantization, preprint, 2026.
- [51] C. Dieball and A. Godec, Mathematical, thermodynamical, and experimental necessity for coarse graining empirical densities and currents in continuous space, *Physical Review Letters* **129** (2022), 140601.
- [52] G. Falasco and M. Esposito, Macroscopic stochastic thermodynamics, *Reviews of Modern Physics* **97** (2025), 015002.
- [53] L. Bertini, R. Chetrite, A. Faggionato, and D. Gabrielli, Level 2.5 large deviations for continuous-time Markov chains with time periodic rates, *Annales Henri Poincaré* **19** (2018), 3197–3238.
- [54] F. Coghi, A. Budhiraja, and J. P. Garrahan, Level 2.5 large deviations and uncertainty relations for non-Markov self-interacting dynamics, preprint, 2026.
- [55] J. Zheng and Z. Lu, Thermodynamic and kinetic bounds for finite-frequency fluctuation-response, preprint, 2026.

- [56] C. Dieball and A. Godec, Perspective: Time irreversibility in systems observed at coarse resolution, *The Journal of Chemical Physics* **162** (2025), 090901.
- [57] U. Seifert, Stochastic thermodynamics, fluctuation theorems and molecular machines, *Reports on Progress in Physics* **75** (2012), 126001.