

Quotient Observation

Hidden Locality, Schur Mediation, and Visible Precision Geometry

J. R. Dunkley

21st March 2026

Abstract

We formulate quotient observation as a nonlinear map on the cone of finite-dimensional positive definite precision operators and show that it is exactly the minimum-energy visible precision induced by the upstairs Gaussian law. This yields well-posedness, positive homogeneity, monotonicity, operator concavity, and exact composition along nested observation towers. For positive upstairs corrections, the finite-observation law from our earlier work is the global tangent ceiling for the visible quotient correction. Compression and quotient differ by a positive semidefinite Schur gap, and on H_0 -adapted sectors the first genuinely quotient-specific term is a sign-definite quartic hidden-mediation subtraction. We then develop the latent geometry of quotient observation: finite hidden sectors preserve tangent support exactly, so hidden load attenuates inside the tangent ceiling but cannot create a new visible direction; beneath a fixed tangent ceiling, the visible laws are classified by a unique positive hidden-load operator, which also fixes the minimal hidden latent dimension and the contraction volume. This identifies the class closed by marginal elimination as the quotient Gaussian precision class itself, with the finite-state Donsker–Varadhan Hessian surviving only as tangent ceiling and microscopic skew-square cone inside that larger attenuation class. We further show that fixed-observed-dimension skew-square closure fails generically by a paired-spectrum obstruction. For hidden-local precision graphs we prove exponential suppression of long-range visible couplings by hidden graph distance and conditioning, identify corridor and flat-bundle extremal architectures, and verify the predicted signatures in synthetic tests. On the Bell square, common visible Gaussian gluing is a law-level compatibility problem with two independent coordinates: remote-setting variance consistency and Bell-locality of the arcsine lift. Bell-locality alone is therefore not sufficient for common gluing of the full visible pair laws. An empirical appendix broadens the programme to real-time homodyne threshold monitoring of $g^{(2)}(0)$, and a second appendix records a transformed semi-coarse-graining benchmark from partially observed entropy-production inference in which the visible law collapses exactly to a finite parity-lifted semi-Markov observer with a two-kernel timing module and an exact finite CTMC completion. Full nonlinear finite-state Markov Donsker–Varadhan re-embedding and completion-independent renewal hidden-load theory remain open.

Contents

1	Introduction	3
2	Quotient precision on the positive cone	4
3	Hierarchical quotient observation	6
4	Compression versus quotient	8
5	The Donsker–Varadhan tangent bridge	8
6	Latent contraction geometry and the visible cone	9
7	Hidden mediation and hidden locality	14
8	Bell frontier: common visible gluing and the classical quotient boundary	17
9	Synthetic falsification battery	20
10	Scope, guardrails, and the open problems	22
11	Discussion	24
12	Blind digital benchmark under partial observation	26
A	Bell-frontier details	28
B	Scalar summaries of the Schur gap	36
C	Quotient adjacency and visible fill-in diagnostics	36
D	Real-time threshold monitoring of $g^{(2)}(0)$	37
E	Renewal-visible transformed semi-coarse-graining benchmark	40
F	Blind benchmark diagnostics	43
G	Methodological bridges and limits of the positive-survivor programme	45

1 Introduction

Given that you only see part of a system, what can you recover exactly, what can you only bound, and what can you never reconstruct uniquely?

Our finite-observation theory isolates a quadratic nonequilibrium visibility law for finite-state Donsker–Varadhan geometry and shows that spectral complexity can be much smaller than topological complexity at the level of observed quadratic fluctuations [1].

That question sits inside the broader Donsker–Varadhan and nonequilibrium large-deviation programme [2–7], as well as the network and stochastic-thermodynamic view of nonequilibrium observables [8–12]. The natural next question is whether that quadratic law is the tangent shadow of a deeper observation theory whose primitive operation is marginal elimination rather than orthogonal compression.

That distinction is structural, not cosmetic. Compression keeps a chosen observed subspace and discards the rest. Quotient observation pushes forward the covariance of the full Gaussian fluctuation law and then re-inverts on the visible space. These operations coincide only in reducing situations. Away from that setting, quotient observation introduces a Schur correction, and that correction is the precise mathematical source of hidden mediation [13–15].

Here we work first on the cone of finite-dimensional positive definite precision operators. The finite-state Markov Donsker–Varadhan setting enters as a structured subclass, not as the whole nonlinear theorem domain. Within that subclass the finite-observation visibility law reappears as the quadratic tangent law on suitable H_0 -adapted quotient sectors. Beyond that order, quotient observation generically leaves the same observed-dimension skew-square cone, and the full nonlinear finite-state Markov re-embedding problem remains open. The object closed under marginal elimination is therefore not the Donsker–Varadhan Hessian itself, but the quotient precision, equivalently the visible minimum-energy or shorted operator associated with the upstairs Gaussian precision law [16–18].

The visible law is not an arbitrary projection of the hidden law. It is a quotient precision determined by a minimum-energy principle, and nested quotient observation composes along a tower of surjections. The finite-observation law is the tangent ceiling and microscopic support cone. Quotienting then attenuates inside that visible tangent support, rather than creating new visible correction directions. In ordinary language, hidden complexity cannot rescue a direction that is already blind at tangent level.

Beneath a fixed tangent ceiling, the finite-hidden visible class is parametrised by a positive hidden-load operator on the tangent support. The central geometric object is therefore the support-preserving visible cone beneath a fixed tangent ceiling; the later sections analyse its tangent, its intrinsic coordinates, its locality shadows, and its compatibility frontiers. Hidden locality enters twice: through Green-kernel mediation in physical coordinates and through positive contraction in latent mode space [19–22]. The present paper then pushes that geometry to one sharp foundations-facing application: on the Bell square, common visible Gaussian gluing is a law-level compatibility problem for full pair laws, not merely a correlator test, and its frontier has two independent coordinates.

The paper is organised as follows. Section 2 formulates quotient precision on the positive cone and proves its variational geometry, calculus, and tangent upper-envelope theorem. Section 3 proves tower composition and derives the hierarchical tangent chain rule. Section 4 separates compression from quotient. Section 5 builds the tangent bridge from finite observation in the Donsker–Varadhan sector.

Section 6 develops the latent-contraction geometry, the support-preserving visible cone, the

intrinsic hidden-load parametrisation, the minimal hidden-dimension law, the determinant invariant, and the same-dimension closure criterion. Section 7 proves hidden-component mediation, intrinsic graph-local decay, and corridor asymptotics. Section 8 identifies the Bell-square gluing frontier, proves that it splits into variance consistency and arcsine-CHSH locality, shows that it is strictly sharper than Bell-locality alone, and recovers the CHSH guardrail for single bounded local readout layers.

Section 9 presents the synthetic falsification battery. Section 10 states the boundary of what is proved and what remains open, including the next compatibility frontier in time. After the discussion, Section 12 reports a blind digital benchmark under partial observation whose role is operational rather than foundational: the theory is frozen first, the hidden world is fixed first, and the visible sector is then used to test whether the theorem-level quantities recovered before reveal agree with the concealed truth.

Appendices collect Bell-frontier details, including the first temporal clique case, scalar gap diagnostics, quotient-adjacency diagnostics, an empirical broadening branch on real-time threshold monitoring of $g^{(2)}(0)$, a transformed semi-coarse-graining benchmark from partially observed entropy-production inference, additional blind-benchmark diagnostics, and a final bridge-theorem bank recording three hidden-to-visible reductions outside the main theorem flow.

2 Quotient precision on the positive cone

2.1 Setting

Let V and Y be finite-dimensional real inner-product spaces, and let $C : V \rightarrow Y$ be a surjective linear map. Equivalently, in coordinates, C has full row rank. Let $\text{SPD}(V)$ denote the cone of symmetric positive definite operators on V .

Definition 2.1 (Quotient observation map). For $H \in \text{SPD}(V)$ we define the quotient precision on Y by

$$F_C(H) := (CH^{-1}C^T)^{-1}. \quad (2.1)$$

This is the visible precision obtained by pushing forward the covariance H^{-1} and re-inverting on the visible space. In finite dimensions this is the Schur-complement or shorted-operator construction seen from the visible sector [13, 14, 16, 17].

Theorem 2.2 (Variational characterisation). Let $C : V \rightarrow Y$ have full row rank and let $H \in \text{SPD}(V)$. Then for every $y \in Y$,

$$y^T F_C(H) y = \min\{x^T H x : Cx = y\}. \quad (2.2)$$

In particular, $F_C(H)$ is the minimum-energy visible precision induced by the upstairs quadratic form, in the precise sense of shorted quadratic forms and parallel addition [16, 17, 23].

Proof. Fix $y \in Y$ and minimise $x^T H x$ subject to $Cx = y$. Introduce a multiplier $\lambda \in Y$ and consider

$$\mathcal{L}(x, \lambda) = x^T H x + 2\lambda^T (Cx - y).$$

Stationarity gives $Hx + C^T \lambda = 0$, hence $x = -H^{-1}C^T \lambda$. Imposing the constraint gives

$$-CH^{-1}C^T \lambda = y, \quad \lambda = -(CH^{-1}C^T)^{-1}y.$$

Therefore the minimiser is

$$x_* = H^{-1}C^T(CH^{-1}C^T)^{-1}y,$$

and its energy is

$$x_*^T H x_* = y^T (CH^{-1}C^T)^{-1}y = y^T F_C(H)y.$$

Since $H \succ 0$, the constrained minimum is unique.

Theorem 2.3 (Well-posedness, homogeneity, monotonicity, and concavity). *Let $C : V \rightarrow Y$ have full row rank. For every $H \in \text{SPD}(V)$, the operator $F_C(H)$ is well defined and belongs to $\text{SPD}(Y)$. Moreover,*

$$F_C(\alpha H) = \alpha F_C(H), \quad \alpha > 0, \quad (2.3)$$

and if $H_1 \succeq H_0$ in Loewner order then

$$F_C(H_1) \succeq F_C(H_0). \quad (2.4)$$

Finally, F_C is operator concave on $\text{SPD}(V)$:

$$F_C(\lambda H_1 + (1 - \lambda)H_0) \succeq \lambda F_C(H_1) + (1 - \lambda)F_C(H_0), \quad 0 \leq \lambda \leq 1. \quad (2.5)$$

The order-theoretic and concavity behaviour here sits in the standard Loewner-operator framework for positive matrices and operator means [24–27].

Proof. Well-posedness and homogeneity are immediate from (2.1). Monotonicity follows because $H_1 \succeq H_0$ implies $H_1^{-1} \preceq H_0^{-1}$, conjugation by C preserves order, and inversion reverses it.

For concavity, fix $y \in Y$. By Theorem 2.2,

$$y^T F_C(H)y = \inf_{Cx=y} x^T H x.$$

For fixed x , the map $H \mapsto x^T H x$ is affine. An infimum of affine functionals is concave. Therefore for each y ,

$$y^T F_C(\lambda H_1 + (1 - \lambda)H_0)y \geq \lambda y^T F_C(H_1)y + (1 - \lambda)y^T F_C(H_0)y.$$

Since this holds for every y , the operator inequality (2.5) follows.

2.2 Fréchet calculus and the tangent ceiling

Theorem 2.4 (Fréchet derivative of the quotient map). *The map $F_C : \text{SPD}(V) \rightarrow \text{SPD}(Y)$ is smooth. Its first Fréchet derivative at $H \in \text{SPD}(V)$ in the direction $\Delta = \Delta^T$ is*

$$DF_C(H)[\Delta] = F_C(H)(CH^{-1}\Delta H^{-1}C^T)F_C(H). \quad (2.6)$$

Proof. Let $G(H) := H^{-1}$, $L(M) := CMC^T$, and $J(N) := N^{-1}$. Then $F_C = J \circ L \circ G$. One has

$$DG(H)[\Delta] = -H^{-1}\Delta H^{-1}, \quad DJ(N)[E] = -N^{-1}EN^{-1}.$$

Therefore

$$DF_C(H)[\Delta] = -F_C(H)C(DG(H)[\Delta])C^T F_C(H),$$

which gives (2.6) after substituting the formula for $DG(H)[\Delta]$.

Corollary 2.5 (Global tangent upper envelope). *Let $H_0 \in \text{SPD}(V)$ and let $\Delta \succeq 0$. Then*

$$0 \preceq F_C(H_0 + \Delta) - F_C(H_0) \preceq DF_C(H_0)[\Delta]. \quad (2.7)$$

Thus the tangent law provides an upper envelope for the visible quotient correction under positive upstairs perturbations.

Proof. The lower bound follows from monotonicity. For the upper bound, define

$$\phi(t) := F_C(H_0 + t\Delta), \quad 0 \leq t \leq 1.$$

By operator concavity,

$$\phi(1) - \phi(0) \preceq D\phi(0)[1] = DF_C(H_0)[\Delta].$$

This is the matrix form of the elementary fact that a concave curve lies below its tangent line.

Remark. Corollary 2.5 is the precise sense in which the finite-observation tangent law is the visible ceiling. Exact quotienting can only attenuate below it. This is the first monotonicity signal of the later support-rigidity theorem: hidden load lowers or redistributes visible strength beneath the tangent ceiling, but it does not generate a genuinely new visible direction.

Remark (Shorted-operator viewpoint). The variational characterisation in Theorem 2.2 shows that quotient precision is the visible minimum-energy, or shorted, quadratic form induced by the upstairs Gaussian law. This is the class closed under marginal elimination [16, 23, 28, 29]. The Donsker–Varadhan skew-square class enters later as a structured microscopic subclass whose visible tangent law sits inside that larger quotient class.

3 Hierarchical quotient observation

Consider a tower of surjective maps

$$X_0 \xrightarrow{C_1} X_1 \xrightarrow{C_2} \dots \xrightarrow{C_L} X_L. \quad (3.1)$$

Let $H_\ell \in \text{SPD}(X_\ell)$ and define recursively

$$H_{\ell+1} = F_{C_{\ell+1}}(H_\ell). \quad (3.2)$$

Fix a reference backbone ladder $H_{0,\ell}$ satisfying the same recursion. This makes the section the linear-algebraic analogue of iterative coarse-graining and renormalisation ladders [30–32].

Proposition 3.1 (Compositionality of quotient observation). *Let*

$$C_1 : X_0 \rightarrow X_1, \quad C_2 : X_1 \rightarrow X_2$$

be surjective linear maps. Then for every $H \in \text{SPD}(X_0)$,

$$F_{C_2}(F_{C_1}(H)) = F_{C_2 C_1}(H). \quad (3.3)$$

More generally, for a tower of surjections C_1, \dots, C_L ,

$$F_{C_L} \circ \dots \circ F_{C_1} = F_{C_L \dots C_1}. \quad (3.4)$$

Proof. By definition,

$$F_{C_2}(F_{C_1}(H)) = \left(C_2(F_{C_1}(H))^{-1} C_2^T \right)^{-1}.$$

Since

$$(F_{C_1}(H))^{-1} = C_1 H^{-1} C_1^T,$$

we obtain

$$F_{C_2}(F_{C_1}(H)) = (C_2 C_1 H^{-1} (C_2 C_1)^T)^{-1} = F_{C_2 C_1}(H).$$

The tower statement follows by iteration.

Remark. Proposition 3.1 is the associativity law of nested quotient observation. The hierarchical tangent chain rule below is its first-order linearisation around a reference backbone ladder.

Theorem 3.2 (Hierarchical tangent chain rule). *Suppose*

$$H_\ell = H_{0,\ell} + \Delta_\ell,$$

with Δ_ℓ small in operator norm. Then

$$\Delta_{\ell+1} = T_{\ell+1}(\Delta_\ell) + R_{\ell+1}(\Delta_\ell), \quad (3.5)$$

where

$$T_{\ell+1}(\Delta) := H_{0,\ell+1} (C_{\ell+1} H_{0,\ell}^{-1} \Delta H_{0,\ell}^{-1} C_{\ell+1}^T) H_{0,\ell+1} \quad (3.6)$$

and

$$\|R_{\ell+1}(\Delta_\ell)\| \leq c_{\ell+1} \|\Delta_\ell\|^2$$

for $\|\Delta_\ell\|$ sufficiently small. Consequently,

$$\Delta_L = T_L \circ T_{L-1} \circ \dots \circ T_1(\Delta_0) + \text{higher-order terms}. \quad (3.7)$$

Proof. Apply Theorem 2.4 at each rung to the map $F_{C_{\ell+1}}$ around $H_{0,\ell}$. Taylor expansion with remainder on the positive cone gives (3.5). Iterating the one-step formula yields (3.7).

Remark. Equation (3.7) is the linear transport law for infinitesimal visible corrections along the quotient hierarchy of Proposition 3.1. The hierarchy itself composes without error. The remainder enters only when one truncates to tangent order.

4 Compression versus quotient

4.1 Exact distinction

Fix a decomposition $V = U \oplus U^\perp$ and write

$$H = \begin{pmatrix} H_{SS} & H_{SF} \\ H_{FS} & H_{FF} \end{pmatrix}. \quad (4.1)$$

Compression keeps the visible block, while quotient elimination computes the Schur complement [14, 15, 33].

Proposition 4.1 (Exact compression-quotient comparison). *In the block setting (4.1),*

$$H_U^{\text{comp}} = H_{SS}, \quad H_U^{\text{quot}} = H_{SS} - H_{SF}H_{FF}^{-1}H_{FS}. \quad (4.2)$$

Hence the Schur gap

$$G_U := H_U^{\text{comp}} - H_U^{\text{quot}} = H_{SF}H_{FF}^{-1}H_{FS} \quad (4.3)$$

is positive semidefinite.

Proof. The quotient formula is the standard Schur complement identity for a block positive definite matrix. The expression (4.3) is of the form $AB^{-1}A^T$ with $B \in \text{SPD}(U^\perp)$, hence $G_U \succeq 0$.

Scalar summaries. For concise one-number summaries of the Schur gap, including trace, operator-norm, and log-determinant diagnostics, see Appendix B.

5 The Donsker–Varadhan tangent bridge

We now connect quotient geometry to the finite-observation theory. We use from [1] the quadratic Donsker–Varadhan correction formula, building on the foundational Donsker–Varadhan theory and the Ky Fan spectral envelope viewpoint [2–4, 34, 35],

$$H_{DV} = H_0 + \Delta_{DV}, \quad \Delta_{DV} = \frac{1}{4} \widehat{J} H_0^{-1} \widehat{J}^T \succeq 0, \quad (5.1)$$

and the corresponding block formulas for observed compression sectors.

Introduce an explicit perturbation parameter by

$$\widehat{J} = \varepsilon J_1, \quad H(\varepsilon) = H_0 + \varepsilon^2 \Delta_2 + O(\varepsilon^4). \quad (5.2)$$

Definition 5.1 (H_0 -adapted quotient sector). *A decomposition $V = U \oplus U^\perp$ is called H_0 -adapted if*

$$H_0 = \begin{pmatrix} H_{0,S} & 0 \\ 0 & H_{0,F} \end{pmatrix} \quad (5.3)$$

in the associated block coordinates.

Theorem 5.2 (Signed quartic bridge). Assume the Donsker–Varadhan expansion (5.1) with bookkeeping (5.2), and let $V = U \oplus U^\perp$ be H_0 -adapted. Then:

1. compression recovers the finite-observation visible quadratic law on U ,
2. quotient and compression agree through order ε^2 ,
3. the first genuinely new quotient effect is quartic and sign-definite:

$$H_U^{\text{comp}} - H_U^{\text{quot}} = \varepsilon^4 \Delta_{SF}^{(2)} H_{0,F}^{-1} \Delta_{FS}^{(2)} + O(\varepsilon^6) \succeq 0. \quad (5.4)$$

Proof. Because the splitting is H_0 -adapted, the mixed backbone blocks vanish. Hence

$$H_{SF}(\varepsilon) = \varepsilon^2 \Delta_{SF}^{(2)} + O(\varepsilon^4), \quad H_{FF}(\varepsilon) = H_{0,F} + O(\varepsilon^2).$$

Compression gives the visible block $H_{SS} = H_{0,S} + \varepsilon^2 \Delta_{SS}^{(2)} + O(\varepsilon^4)$, which is the finite-observation visible quadratic law imported from [1]. Proposition 4.1 gives

$$H_U^{\text{comp}} - H_U^{\text{quot}} = H_{SF} H_{FF}^{-1} H_{FS}.$$

Substituting the block expansions yields (5.4). In particular the gap is $O(\varepsilon^4)$, so quotient and compression agree through quadratic order.

Corollary 5.3 (Tangent closure of visible Donsker–Varadhan geometry). In the H_0 -adapted setting,

$$H_U^{\text{quot}} = H_{0,S} + (\Delta_{DV})_{SS} + O(\|\widehat{J}\|^4). \quad (5.5)$$

Thus the finite-observation visible law is the quadratic tangent law of quotient observation in the Donsker–Varadhan sector.

Proof. Combine Theorem 5.2 with (5.1).

6 Latent contraction geometry and the visible cone

6.1 Exact latent-contraction formula

Work in an H_0 -adapted quotient chart and whiten by the backbone so that the upstairs precision takes the form

$$\widetilde{H} = I + LL^T, \quad L = \begin{pmatrix} B \\ C \end{pmatrix}, \quad (6.1)$$

where B contains the visible rows and C the hidden rows. At the level of Gaussian graphical models this is the hidden-variable precision geometry rather than a heuristic latent-variable ansatz [22, 36, 37]. The first theorem of this section makes the tangent ceiling exact and visible: quotienting keeps the same support ceiling as compression, but contracts beneath it by a positive hidden load.

Theorem 6.1 (Exact latent-contraction formula). *In the setting (6.1), the visible quotient correction is*

$$X = B(I + C^T C)^{-1} B^T. \quad (6.2)$$

The tangent ceiling is

$$T := BB^T, \quad (6.3)$$

and the gap is

$$T - X = B C^T (I + C C^T)^{-1} C B^T \succeq 0. \quad (6.4)$$

Equivalently,

$$X = BRB^T, \quad R := (I + C^T C)^{-1}, \quad 0 \prec R \leq I. \quad (6.5)$$

Proof. In block form,

$$\tilde{H} = \begin{pmatrix} I + BB^T & BC^T \\ CB^T & I + CC^T \end{pmatrix}.$$

The visible quotient precision is its Schur complement:

$$I + BB^T - BC^T (I + CC^T)^{-1} CB^T.$$

Using the identity

$$I - C^T (I + CC^T)^{-1} C = (I + C^T C)^{-1},$$

we obtain

$$I + B(I + C^T C)^{-1} B^T.$$

Subtracting the whitened backbone I gives (6.2). The formulae (6.4) and (6.5) follow immediately.

Corollary 6.2 (Exact support theorem). *For every finite hidden sector,*

$$\text{Ran}(X) = \text{Ran}(T), \quad \text{rank}(X) = \text{rank}(T). \quad (6.6)$$

Thus finite hidden mediation preserves the visible correction support. In particular the hidden sector attenuates inside the existing visible support rather than creating new visible directions, which is also the natural completion geometry for latent Gaussian models [22, 38].

Proof. Since $R \succ 0$ on the latent space, BRB^T and BB^T have the same range and rank.

Remark (Blind directions remain blind). Corollary 6.2 is the ordinary-language impossibility theorem behind the quotient geometry. If a visible direction receives zero signal from the tangent ceiling T , then no finite hidden architecture can make that direction informative later, because the actual visible correction X has exactly the same support. Hidden complexity can attenuate, split, or redistribute strength inside the tangent support, but it cannot rescue a blind direction.

6.2 The support-preserving visible cone and hidden-load parametrisation

Proposition 6.3 (Support-preserving interval theorem). Let $T = BB^T$ and let $S := \text{Ran}(T)$. A matrix X arises from a finite hidden sector if and only if

$$X = T^{1/2}\Pi T^{1/2}, \quad 0 \prec \Pi \leq I \text{ on } S. \quad (6.7)$$

Taking closure adds the boundary $0 \leq \Pi \leq I$ on S . Equivalently, the closure of the visible image is the support-preserving Loewner interval beneath T .

Proof. From Theorem 6.1,

$$X = BRB^T, \quad 0 \prec R \leq I.$$

By polar decomposition, $B = T^{1/2}U$ for a partial isometry U with initial space the latent support of B and final space S . Therefore

$$X = T^{1/2}(URU^T)T^{1/2},$$

so (6.7) holds with $\Pi := URU^T$. Since $R \succ 0$ and $R \leq I$, one has $0 \prec \Pi \leq I$ on S .

Conversely, let $0 \prec \Pi \leq I$ on S . Choose the latent space to be S , set $B := T^{1/2}$, and define $R := \Pi$. Since $R \succ 0$ and $R \leq I$, the matrix $G := R^{-1} - I$ is positive semidefinite. Pick C with $C^T C = G$. Then $(I + C^T C)^{-1} = R$, so $X = BRB^T$. The boundary case follows by limit.

Theorem 6.4 (Intrinsic hidden-load parametrisation). Let $T \succeq 0$ and let $S := \text{Ran}(T)$. A visible law X in the support-preserving interval beneath T is equivalently an operator of the form

$$X = T^{1/2}\Pi T^{1/2}, \quad 0 \prec \Pi \leq I \text{ on } S.$$

For such X , define the intrinsic hidden-load operator on S by

$$\Lambda := (T|_S)^{1/2}(X|_S)^{-1}(T|_S)^{1/2} - I_S. \quad (6.8)$$

Then $\Lambda \succeq 0$ on S , and

$$X = T^{1/2}(I_S + \Lambda)^{-1}T^{1/2}. \quad (6.9)$$

Conversely, every $\Lambda \succeq 0$ on S determines a visible law by (6.9). Hence the support-preserving visible class beneath T is in bijection with the positive cone on S .

Proof. Write

$$\Pi := (T|_S)^{-1/2}(X|_S)(T|_S)^{-1/2}.$$

Then $0 \prec \Pi \leq I_S$ on S . Therefore

$$\Lambda = \Pi^{-1} - I_S \succeq 0,$$

and

$$\Pi = (I_S + \Lambda)^{-1}.$$

Substituting this into $X = T^{1/2}\Pi T^{1/2}$ gives (6.9).

Conversely, if $\Lambda \succeq 0$ on S , then $\Pi := (I_S + \Lambda)^{-1}$ satisfies

$$0 \prec \Pi \leq I_S,$$

so Proposition 6.3 yields a visible realisation

$$X = T^{1/2}\Pi T^{1/2} = T^{1/2}(I_S + \Lambda)^{-1}T^{1/2}.$$

This proves the bijection.

Corollary 6.5 (Hidden-load monotonicity). *Let*

$$X_i = T^{1/2}(I_S + \Lambda_i)^{-1}T^{1/2}, \quad \Lambda_i \succeq 0 \text{ on } S,$$

for $i = 0, 1$. Then

$$\Lambda_1 \succeq \Lambda_0 \iff X_1 \preceq X_0. \quad (6.10)$$

Thus increasing hidden load can only attenuate the visible correction.

Proof. Restrict to S , where $T|_S \succ 0$. Congruence by $(T|_S)^{-1/2}$ gives

$$X_1 \preceq X_0 \iff (I_S + \Lambda_1)^{-1} \preceq (I_S + \Lambda_0)^{-1}.$$

Since inversion reverses Loewner order on positive definite operators,

$$(I_S + \Lambda_1)^{-1} \preceq (I_S + \Lambda_0)^{-1} \iff I_S + \Lambda_1 \succeq I_S + \Lambda_0,$$

which is (6.10).

Theorem 6.6 (Minimal hidden latent dimension). *Let $X = T^{1/2}\Pi T^{1/2}$ with $0 \prec \Pi \leq I$ on $S = \text{Ran}(T)$. Then the minimal hidden latent dimension needed to realise X is*

$$d_{\min} = \text{rank}(I - \Pi) = \text{rank}(T - X). \quad (6.11)$$

Proof. By Theorem 6.4, write

$$X = T^{1/2}(I_S + \Lambda)^{-1}T^{1/2}, \quad \Lambda \succeq 0.$$

Any realisation requires a Gram factorisation $\Lambda = C^T C$, so the hidden latent dimension must be at least $\text{rank}(\Lambda)$. Choosing a Gram factorisation of Λ shows that $\text{rank}(\Lambda)$ hidden dimensions suffice. Since

$$T - X = T^{1/2}\left(I_S - (I_S + \Lambda)^{-1}\right)T^{1/2},$$

and $I_S - (I_S + \Lambda)^{-1} = \Lambda(I_S + \Lambda)^{-1}$ has the same rank as Λ , we obtain

$$\text{rank}(\Lambda) = \text{rank}(T - X) = \text{rank}(I - \Pi).$$

This proves the claim.

Corollary 6.7 (Rank and volume of hidden load). *With Λ as in Theorem 6.4,*

$$\text{rank}(\Lambda) = \text{rank}(T - X), \quad (6.12)$$

and

$$\log \text{pdet}(T) - \log \text{pdet}(X) = \log \det(I_S + \Lambda). \quad (6.13)$$

In particular, the minimal hidden latent dimension is $\text{rank}(\Lambda) = \text{rank}(T - X)$.

Proof. From (6.9),

$$T - X = T^{1/2} \left(I_S - (I_S + \Lambda)^{-1} \right) T^{1/2} = T^{1/2} \Lambda (I_S + \Lambda)^{-1} T^{1/2}.$$

Since $(I_S + \Lambda)^{-1}$ is invertible on S , this implies

$$\text{rank}(T - X) = \text{rank}(\Lambda).$$

Also, on S ,

$$X|_S = (T|_S)^{1/2} (I_S + \Lambda)^{-1} (T|_S)^{1/2},$$

hence

$$\det(X|_S) = \det(T|_S) \det(I_S + \Lambda)^{-1}.$$

Taking logarithms gives (6.13).

Remark (Operator-interval picture). Beneath a fixed tangent ceiling T , finite hidden sectors realise exactly the support-preserving Loewner interval

$$X = T^{1/2} \Pi T^{1/2}, \quad 0 \prec \Pi \leq I \text{ on } S = \text{Ran}(T).$$

Here Π records the visible attenuation on the tangent support, $\text{rank}(I - \Pi) = \text{rank}(T - X)$ is the least hidden latent dimension, and

$$\log \det(I_S + \Lambda) = \log \text{pdet}(T) - \log \text{pdet}(X)$$

measures the latent contraction volume. The same-observed-dimension skew-square class analysed next is therefore only a special paired-spectrum slice inside this larger interval [39].

6.3 Same-dimension closure criterion

Proposition 6.8 (Paired-spectrum criterion). *Let $H_{0,\text{vis}} \succ 0$ and let $\Delta_{\text{vis}} \succeq 0$ be a visible correction. Set*

$$W := H_{0,\text{vis}}^{-1/2} \Delta_{\text{vis}} H_{0,\text{vis}}^{-1/2}. \quad (6.14)$$

There exists a visible skew matrix J_{vis} of the same observed dimension such that

$$\Delta_{\text{vis}} = \frac{1}{4} J_{\text{vis}} H_{0,\text{vis}}^{-1} J_{\text{vis}}^T$$

if and only if the positive spectrum of W is pairwise degenerate. In particular, odd visible correction rank is impossible for same-dimension skew-square closure.

This is the same real skew-symmetric pairing obstruction that underlies Williamson-type normal forms and the canonical even-multiplicity structure of skew squares [40, 41].

Proof. If such J_{vis} exists, define

$$K := H_{0,\text{vis}}^{-1/2} J_{\text{vis}} H_{0,\text{vis}}^{-1/2}.$$

Then K is real skew-symmetric and

$$W = \frac{1}{4}KK^T.$$

The nonzero singular values of a real skew matrix come in equal pairs, hence the positive eigenvalues of W do as well.

Conversely, if the positive spectrum of W is pairwise degenerate, choose an orthogonal basis in which

$$W = Q \operatorname{diag}(\mu_1, \mu_1, \mu_2, \mu_2, \dots, 0, \dots, 0) Q^T, \quad \mu_i > 0.$$

Let

$$K := 2Q \operatorname{diag}(\sqrt{\mu_1}J_2, \sqrt{\mu_2}J_2, \dots, 0) Q^T,$$

where $J_2 = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}$. Then K is real skew-symmetric and $W = \frac{1}{4}KK^T$. Setting

$$J_{\text{vis}} := H_{0,\text{vis}}^{1/2} K H_{0,\text{vis}}^{1/2}$$

gives the required representation.

Remark (Nongenericity of same-dimension skew-square closure). Proposition 6.8 identifies the precise same-dimension obstruction. On any fixed visible support stratum where pairwise spectral degeneracy is a nontrivial condition, that obstruction cuts out a proper algebraic subset of the positive cone. Same-observed-dimension skew-square closure is therefore structurally nongeneric inside the quotient-visible cone beneath a fixed tangent ceiling. This is the clean algebraic reason the quotient-visible class is genuinely larger than the fixed-observed-dimension Donsker–Varadhan skew-square class. The synthetic failure fractions in Section 9 should be read as the numerical shadow of that geometry, not as an empirical surprise.

Remark (What is and is not closed). Theorems 6.1 and 6.4, together with Proposition 6.8, show the correct closure picture. Exact quotienting preserves Gaussian precision geometry and the support-preserving attenuation class beneath a fixed tangent ceiling. It does not generically preserve the same observed-dimension Donsker–Varadhan skew-square class. The microscopic Donsker–Varadhan Hessian should therefore be regarded as the tangent ceiling and microscopic cone inside the larger quotient-visible class, not as the class closed by Schur elimination itself.

7 Hidden mediation and hidden locality

7.1 Componentwise hidden mediation

Let $V = V_{\text{vis}} \oplus V_{\text{hid}}$ and let

$$K_X = \begin{pmatrix} K_{VV} & K_{VH} \\ K_{HV} & K_{HH} \end{pmatrix} \in \text{SPD}(V) \tag{7.1}$$

be a local precision operator. This is the standard Gaussian graphical-model setting in which precision sparsity encodes conditional independence and Schur elimination creates visible

fill-in [36, 37, 42]. The point of the next two theorems is that this fill-in is not arbitrary. It is split by hidden connected component and then quantitatively suppressed by hidden graph distance and conditioning.

Theorem 7.1 (Hidden-component mediation). *The visible effective precision after hidden elimination is*

$$K_{\text{eff}} = K_{VV} - K_{VH}K_{HH}^{-1}K_{HV}. \quad (7.2)$$

If the hidden graph decomposes into connected components $H = \sqcup_{\alpha} H_{\alpha}$ respected by K_{HH} , then

$$K_{VH}K_{HH}^{-1}K_{HV} = \sum_{\alpha} K_{VH_{\alpha}}K_{H_{\alpha}H_{\alpha}}^{-1}K_{H_{\alpha}V}. \quad (7.3)$$

Consequently a visible pair can acquire a new coupling only if both visible coordinates connect to the same hidden component.

The decomposition in (7.3) is the Green-kernel or walk-sum mediation picture for hidden elimination [20, 21].

Proof. Equation (7.2) is the Schur complement formula. If the hidden graph splits into connected components respected by K_{HH} , then after permuting indices one has

$$K_{HH} = \text{diag}(K_{H_1H_1}, \dots, K_{H_mH_m}),$$

so the inverse is block diagonal with the same decomposition. Substituting this into the Schur term yields (7.3). The last claim is immediate from the support structure of the sum.

7.2 Intrinsic graph-local decay

This extends the classical decay-of-inverse theory for banded and sparse matrices to the present graph-local hidden precision setting [19, 43, 44].

Theorem 7.2 (Chebyshev graph-distance bound). *Let $D \succ 0$ be a hidden precision matrix supported on a hidden graph G_H , and suppose*

$$\text{spec}(D) \subset [\alpha, \beta], \quad 0 < \alpha < \beta < \infty.$$

Let S, T be hidden vertex sets with graph distance at least δ . Then

$$\|P_S D^{-1} P_T\|_{\text{op}} \leq \frac{2}{\alpha} \frac{\rho^{\delta}}{1 + \rho^{2\delta}}, \quad \rho := \frac{\sqrt{\kappa} - 1}{\sqrt{\kappa} + 1}, \quad \kappa := \beta/\alpha. \quad (7.4)$$

Proof. Let p be any polynomial of degree at most $\delta - 1$. Since D^j propagates only along hidden walks of length at most j , one has

$$P_S p(D) P_T = 0.$$

Hence

$$P_S D^{-1} P_T = P_S (D^{-1} - p(D)) P_T,$$

so

$$\|P_S D^{-1} P_T\|_{\text{op}} \leq \|D^{-1} - p(D)\|_{\text{op}} \leq \max_{x \in [\alpha, \beta]} \left| \frac{1}{x} - p(x) \right|.$$

Minimising the right-hand side over degree- $\delta - 1$ polynomials is the classical minimax approximation problem for $1/x$ on $[\alpha, \beta]$ [45, 46]. The shifted Chebyshev solution gives

$$\min_{\deg p \leq \delta-1} \max_{x \in [\alpha, \beta]} \left| \frac{1}{x} - p(x) \right| = \frac{2}{\alpha} \frac{\rho^\delta}{1 + \rho^{2\delta}},$$

with the stated value of ρ .

Corollary 7.3 (Visible coupling decay). *Let b_u, b_v be visible-to-hidden coupling vectors supported on hidden sets at graph distance at least δ . Then*

$$|b_u^T D^{-1} b_v| \leq \|b_u\|_2 \|b_v\|_2 \frac{2}{\alpha} \frac{\rho^\delta}{1 + \rho^{2\delta}}. \quad (7.5)$$

Proof. Apply Theorem 7.2 and Cauchy–Schwarz.

Theorem 7.4 (Corridor asymptotics and bundle extremisers). *Let $P_{\delta+1}$ be the path on $\delta + 1$ hidden vertices and define the constant-coefficient corridor precision*

$$D_\delta^{\text{corr}} := cI - dA_{P_{\delta+1}}, \quad c := \frac{\alpha + \beta}{2}, \quad d := \frac{\beta - \alpha}{4}. \quad (7.6)$$

Then $\text{spec}(D_\delta^{\text{corr}}) \subset [\alpha, \beta]$, the endpoint distance is δ , and the endpoint transfer obeys

$$(D_\delta^{\text{corr}})_{1, \delta+1}^{-1} = C_{\text{corr}}(\alpha, \beta) \rho^\delta (1 + o(1)), \quad \delta \rightarrow \infty, \quad (7.7)$$

for an explicit positive constant $C_{\text{corr}}(\alpha, \beta)$. Consequently the one-dimensional corridor attains the universal long-range decay rate ρ^δ asymptotically.

Moreover, if one takes k disjoint identical corridors and splits the visible access budget equally across them, the resulting off-mask visible block has flat singular spectrum and asymptotically realises the many-channel threshold-count extremiser permitted by the rank, nuclear, and Frobenius budgets.

Proof. The spectral inclusion for (7.6) follows from the path adjacency spectrum. The inverse of a constant-coefficient Jacobi chain is explicit and gives the stated asymptotic rate (7.7). Since Theorem 7.2 rules out any slower universal decay rate, the corridor is asymptotically optimal at the exponential level.

For the many-channel statement, disjoint identical corridors produce a block proportional to the identity in the corridor basis. Hence the nonzero singular values are equal. This is the flat-spectrum shape that maximises threshold count under fixed rank, nuclear, and Frobenius budgets.

Remark. Theorem 7.4 is an asymptotic graph-extremiser statement. What is proved is the optimal universal rate and its corridor realisation class. Exact finite- δ minimax attainment of the best constant remains open. The synthetic battery below supports this rate-level statement and at the same time suggests that looped topologies can improve finite-distance prefactors under the tested spectral-window normalisation without changing the universal decay exponent.

Adjacency-loss diagnostics. For an intrinsic quotient-adjacency layer based on fibre distance, locality-defect operators, and spectral diagnostics of visible fill-in, see Appendix C.

8 Bell frontier: common visible gluing and the classical quotient boundary

8.1 Bell-square gluing frontier

The graph-local results above explain how strictly local hidden precision geometry can cast dense visible shadows after elimination. That is a theorem about long-range visible fill-in in the precision law. The Bell-facing question is different. Given four setting-wise visible pair laws, when do they glue to one common visible Gaussian law, and what part of that frontier is visible to Bell correlators alone? The answer is sharper than a CHSH-only reading: the Bell-square gluing frontier has two independent coordinates, one correlator-visible and one purely variance-level.

Definition 8.1 (Common Schur gluing of context families). *Let*

$$Q_{xy} \in \text{SPD}(2), \quad x, y \in \{0, 1\},$$

be four visible pair precisions, and let

$$P_{xy} : \mathbb{R}^4 \rightarrow \mathbb{R}^2$$

select the coordinate pair (A_x, B_y) from (A_0, A_1, B_0, B_1) . We say that the family (Q_{xy}) is commonly Schur-glued if there exists $Q_ \in \text{SPD}(4)$ such that*

$$Q_{xy} = F_{P_{xy}}(Q_*) \quad \text{for all } x, y \in \{0, 1\}. \quad (8.1)$$

Theorem 8.2 (Bell-square collapse theorem). *Let*

$$Q_{xy} = \begin{pmatrix} a_{xy} & b_{xy} \\ b_{xy} & d_{xy} \end{pmatrix} \in \text{SPD}(2), \quad \Sigma_{xy} := Q_{xy}^{-1} = \begin{pmatrix} v_{xy}^A & c_{xy} \\ c_{xy} & v_{xy}^B \end{pmatrix},$$

and define

$$\rho_{xy} := \frac{c_{xy}}{\sqrt{v_{xy}^A v_{xy}^B}} = \frac{-b_{xy}}{\sqrt{a_{xy} d_{xy}}}, \quad \phi_{xy} := \arcsin(\rho_{xy}). \quad (8.2)$$

Then the following are equivalent.

1. There exists a centred Gaussian law on (A_0, A_1, B_0, B_1) whose (A_x, B_y) marginals are exactly the pair laws with covariances Σ_{xy} . In the positive definite case this is equivalent to common Schur gluing in the sense of Definition 8.1.
2. The family satisfies remote-setting variance consistency,

$$v_{x0}^A = v_{x1}^A \quad (x = 0, 1), \quad v_{0y}^B = v_{1y}^B \quad (y = 0, 1), \quad (8.3)$$

and the four arcsine CHSH inequalities

$$|\phi_{00} + \phi_{01} + \phi_{10} - \phi_{11}| \leq \pi, \quad (8.4)$$

$$|\phi_{00} + \phi_{01} - \phi_{10} + \phi_{11}| \leq \pi, \quad (8.5)$$

$$|\phi_{00} - \phi_{01} + \phi_{10} + \phi_{11}| \leq \pi, \quad (8.6)$$

$$|-\phi_{00} + \phi_{01} + \phi_{10} + \phi_{11}| \leq \pi. \quad (8.7)$$

Equivalently, with

$$E_{xy} := \frac{2}{\pi} \phi_{xy} = \frac{2}{\pi} \arcsin\left(\frac{-b_{xy}}{\sqrt{a_{xy} d_{xy}}}\right), \quad (8.8)$$

common gluing is equivalent to variance consistency together with Bell-locality of the lifted correlator table $E = (E_{xy})$.

Proof. Appendix A proves the Gaussian gluing criterion and then specialises the Bell square using standard background facts from the correlator literature. The new point for the present paper is that the visible pair-precision frontier is not exhausted by Bell correlators: it also contains the diagonal variance-gluing constraints (8.3).

Proposition 8.3 (Strict refinement over Bell-locality). *Bell-locality of the lifted correlator table E is necessary but not sufficient for common visible Gaussian gluing of the full pair laws.*

Proof. Consider

$$Q_{00} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \quad Q_{01} = \begin{pmatrix} 1/4 & 0 \\ 0 & 1 \end{pmatrix}, \quad Q_{10} = Q_{11} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}.$$

Then every off-diagonal entry vanishes, so (8.8) gives $E_{xy} = 0$ for all x, y , hence the lifted

correlator table is Bell-local. But

$$Q_{00}^{-1} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \quad Q_{01}^{-1} = \begin{pmatrix} 4 & 0 \\ 0 & 1 \end{pmatrix},$$

so the variance of A_0 depends on Bob's setting. The variance-consistency conditions (8.3) therefore fail, and no common visible Gaussian gluing exists.

Remark (Two Bell-square obstructions). For a Bell-square family, define

$$\varepsilon_{\text{var}} := \max\{|v_{00}^A - v_{01}^A|, |v_{10}^A - v_{11}^A|, |v_{00}^B - v_{10}^B|, |v_{01}^B - v_{11}^B|\},$$

and let $\varepsilon_{\text{CHSH}}$ be the excess of the largest left-hand side of (8.4) to (8.7) above π , truncated below at zero. Then common gluing on the Bell square is equivalent to the simultaneous vanishing of both obstructions,

$$\varepsilon_{\text{var}} = 0, \quad \varepsilon_{\text{CHSH}} = 0.$$

This packages the Bell-square frontier as a two-coordinate compatibility problem: one coordinate is invisible to Bell correlators, while the other is the usual arcsine CHSH obstruction. In fact every Bell-local 2×2 correlator table already sits inside a rank-two mediated Gaussian sign-shadow, as recorded in Appendix A.

Corollary 8.4 (Bell cube normal form). *Let S be the 4×4 matrix whose rows are the four CHSH sign vectors*

$$(1, 1, 1, -1), \quad (1, 1, -1, 1), \quad (1, -1, 1, 1), \quad (-1, 1, 1, 1),$$

and define lifted Bell coordinates

$$y := \frac{1}{\pi} S\phi, \quad \phi := (\phi_{00}, \phi_{01}, \phi_{10}, \phi_{11})^\top.$$

Then the Bell-compatible region in arcsine coordinates is exactly the cube

$$|y_i| \leq 1, \quad i = 1, \dots, 4.$$

Equivalently, the four CHSH inequalities are the coordinate inequalities of one global hypercubic normal form. In particular, Euclidean nearest-point repair in y is coordinatewise clipping, and the Bell support function is

$$h_{\mathcal{B}}(u) = \frac{\pi}{4} \|Su\|_1.$$

Appendix A records the derivation and the exact repair formulas.

Corollary 8.5 (Single-readout CHSH locality). *Suppose a Bell-context family (Q_{xy}) is commonly Schur-glued, and let*

$$Z = (Z_{A_0}, Z_{A_1}, Z_{B_0}, Z_{B_1}) \sim N(0, Q_*^{-1})$$

for some Q_* satisfying (8.1). Let

$$A_x = f_x(Z_{A_x}), \quad B_y = g_y(Z_{B_y}), \quad |f_x| \leq 1, \quad |g_y| \leq 1.$$

Then the correlators

$$\tilde{E}_{xy} := \mathbb{E}[A_x B_y]$$

obey the CHSH bound

$$|\tilde{E}_{00} + \tilde{E}_{01} + \tilde{E}_{10} - \tilde{E}_{11}| \leq 2, \tag{8.9}$$

and likewise for the other three sign choices with one minus sign.

Proof. All four outputs are measurable functions of the single common random vector Z . Hence for every realisation one has

$$|A_0(B_0 + B_1) + A_1(B_0 - B_1)| \leq |B_0 + B_1| + |B_0 - B_1| \leq 2,$$

because $|A_x| \leq 1$ and $|B_y| \leq 1$. Taking expectations gives (8.9), and the other sign choices follow by relabelling.

Remark. The Bell-facing lesson is now sharper than a mere guardrail. Common visible Gaussian gluing on the Bell square is a law-level compatibility problem: Bell correlators see only the arcsine shadow of that frontier, while the full visible law also contains the diagonal variance-gluing constraints. Proposition 8.3 shows that the law frontier is strictly sharper than Bell-locality. Corollary 8.5 then recovers the original operational guardrail for any single bounded local readout layer. The same quotient-gluing viewpoint also suggests a temporal frontier: one replaces Bell contexts by an observation graph of times and asks for common gluing of the visible pair laws along that graph. We do not develop that temporal theory here, but it is the natural next compatibility problem because it keeps the same law-level question while replacing the Bell square by a time-indexed marginal pattern.

9 Synthetic falsification battery

This section applies direct falsification pressure to the sharpest theorem-level claims in synthetic hidden-local systems. We tested four families: an H_0 -adapted bridge ensemble built from block-diagonal backbones and skew-generated Donsker–Varadhan perturbations; a generic quotient Gaussian ensemble designed to stress same-dimension skew-square closure; hidden-local chain and sparse-graph families designed to probe graph-distance decay; and a topology-and-bundle family used to compare one-channel and many-channel transfer under local hidden mediation.

The goal was not curve fitting. It was to test the statements that are most exposed: quartic bridge scaling, support preservation, the paired-spectrum closure obstruction, hidden-local controlled-tail decay, and the corridor-versus-bundle extremiser picture. Table 1 summarises the aggregate diagnostics.

Diagnostic	Tested regime	Quantitative outcome	Interpretation
Quartic bridge slope	Adapted bridge sweeps with visible and hidden dimensions 2 to 4	Grand mean slope 3.9957; nine dimension-pair means from 3.9936 to 3.9977	Consistent with the quartic hidden-mediation bridge and stable across the tested dimension range.
Worst bridge gap eigenvalue	Same adapted bridge sweeps	Worst reported bridge-gap eigenvalue -1.17×10^{-15}	Numerically nonnegative up to floating-point tolerance, as required by the sign-definite quartic bridge law.
Support preservation	Same adapted bridge sweeps	Rank-preservation rate 100%; maximum reported principal-angle sine 5.58×10^{-8}	Finite hidden sectors preserve tangent support to numerical precision, matching Corollary 6.2.
Same-dimension closure	Generic quotient ensembles with visible dimensions 3 to 8	Failure fraction 1.0 throughout; pairable fraction 0.0 throughout	Generic same-dimension skew-square closure fails completely in the tested ensemble, matching Proposition 6.8.
Hidden-local decay bound	Random local hidden graphs of size 24, 36, and 48	Bound-compliance rate 100% across all sampled size-distance cells	Visible long-range transfer behaves as a controlled hidden-local tail, consistent with Theorems 7.2 and 7.1.
Topology and bundle tests	Paths, cycles, ladders, trees, and disjoint corridor bundles	Corridor geometry realises the sharp decay class; disjoint bundles give flat singular spectra; under the tested finite-distance normalisation the best one-channel performer was the cycle	Supports the rate-level corridor picture and the flat-spectrum many-channel extremiser law while refining the finite-distance prefactor statement.

Table 1: Synthetic falsification highlights.

Taken together, the synthetic battery does not just illustrate the theory. It exerts direct

falsification pressure on the hardest claims. Across the nine tested visible-hidden dimension pairs the quartic bridge remains pinned near slope 4, the worst bridge-gap eigenvalue stays at floating-point scale, and support preservation holds to numerical precision. Across visible dimensions 3 to 8, generic same-dimension skew-square closure fails in every sampled case, with zero pairable instances. Across the sampled hidden-local graph families every tested entry obeys the decay bound. The substantive refinement forced by the data is that the one-channel topology statement should be phrased at the level of universal rate rather than universal finite-distance constant: corridor geometry cleanly realises the sharp decay class, while looped topologies can improve finite-distance prefactors under the tested spectral-window normalisation.

10 Scope, guardrails, and the open problems

The results above are within finite-dimensional Gaussian precision geometry. They prove, by Theorem 2.3, Corollary 2.5, and Proposition 3.1, that quotient observation defines a variational, monotone, homogeneous, and operator-concave map on the positive cone and composes along towers of surjections.

Corollary 5.3 shows that the finite-observation tangent law is an upper envelope for visible quotient corrections under positive upstairs perturbations, and Proposition 4.1 shows that compression and quotient differ by a positive semidefinite Schur gap. Corollary 6.2, Theorem 6.4, Corollaries 6.5 and 6.7, and Theorem 6.6 show that hidden mediation preserves tangent support while selecting a support-preserving point beneath the tangent ceiling, so hidden complexity attenuates inside the visible ceiling but cannot create a new visible direction. The same results show that the finite-hidden visible laws beneath a fixed tangent ceiling are parametrised by a positive hidden-load operator with monotone attenuation, minimal hidden dimension, and latent contraction volume encoded intrinsically. Proposition 6.8 shows that same-dimension skew-square closure is controlled by a paired-spectrum criterion, giving the clean algebraic reason the quotient-visible cone is generically larger than the fixed-observed-dimension Donsker–Varadhan skew-square class. Theorems 7.2, 7.1, and 7.4 show that graph-local hidden precision forces exponential decay of long-range visible couplings, with corridor architectures providing the asymptotically extremal rate class.

There is also a natural symmetric information-geometric reading of the quotient map. For a surjective observation map $C : U \rightarrow Y$ and an upstairs positive form $H \in \text{SPD}(U)$, the visible quadratic form induced by hidden elimination is uniquely determined by the upstairs geometry and is given invariantly by

$$F_C(H) = (CH^{-1}C^\top)^{-1}.$$

In block coordinates

$$H = \begin{pmatrix} A & B \\ B^\top & D \end{pmatrix}, \quad D \succ 0,$$

this is the Schur complement $A - BD^{-1}B^\top$. We do not develop that observed-information interpretation further here, but it indicates a natural bridge to later work on Fisher and Fisher–Kähler reduction.

The Bell-facing application is equally specific. Section 8 identifies the Bell-square gluing frontier for visible pair laws, proves that common visible Gaussian gluing is equivalent to remote-setting variance consistency together with Bell-locality of the arcsine lift, shows that Bell-locality alone is not sufficient for common gluing of the full pair laws, and then packages the same frontier globally as a cube in lifted Bell coordinates in Corollary 8.4. Corollary 8.5

then recovers the operational statement that any single bounded local readout remains CHSH-local under common Schur gluing. The graph-local results therefore explain visible long-range couplings after elimination, but not Bell violation by themselves. Appendix A records the supporting Bell-frontier derivations, including the Bell cube normal form and repair formulas, Appendix B records scalar summaries of the Schur gap, Appendix C records quotient-adjacency diagnostics, Appendix D records the empirical broadening branch on real-time threshold monitoring of $g^{(2)}(0)$, and Appendix F records additional blind-benchmark diagnostics.

Remark. At present the following six statements should be kept separate.

1. Gaussian quotient closure and tower composition are proved.
2. The closed object is the quotient precision, equivalently the visible minimum-energy or shorted operator induced by the upstairs Gaussian law.
3. Beneath a fixed tangent ceiling, the finite-hidden visible class is completely parametrised by a positive hidden-load operator on the tangent support, and support rigidity proves that blind tangent directions remain blind after finite hidden elimination.
4. On the Bell square, common visible Gaussian gluing is equivalent to variance consistency together with Bell-locality of the arcsine lift, so the law-level frontier has two coordinates rather than one, and any jointly Schur-glued context family obeys CHSH under any single bounded local readout layer.
5. Exact tangent closure of visible Donsker–Varadhan geometry in adapted sectors is proved to quadratic order, while generic same-observed-dimension skew-square closure fails by the paired-spectrum obstruction.
6. Full nonlinear finite-state Markov Donsker–Varadhan re-embedding after quotient observation remains open.

Blurring these six levels would overstate the present state of the theory.

A natural next frontier is temporal compatibility. There the same Gaussian gluing question is indexed not by Bell settings but by an observation graph of times. Appendix A records the three-time triangle theorem as the first chordal temporal case: common Gaussian gluing is equivalent to variance consistency together with positivity of the 3×3 correlation completion, and the arcsine lift gives the Leggett–Garg tetrahedron for the sign-shadow variables. We do not develop a general temporal theory beyond that first clique case here, and the present paper should not be read as claiming a general Leggett–Garg or macrorealism theorem. The proved content of the present paper therefore stops at the Bell-square law frontier, the first temporal clique, and the general quotient-geometry results developed above.

What is not proved is equally important. On the Markov side this sits close to the classical lumpability and coarse-graining problem, but our closure theorem is strictly Gaussian and does not yet solve nonlinear Markov re-embedding [48–50]. In particular, the class closed by marginal elimination is the quotient Gaussian precision class itself. Fixed-observed-dimension Donsker–Varadhan closure is generically the wrong target: the paired-spectrum obstruction shows that it is structurally exceptional inside the quotient-visible cone. The Donsker–Varadhan Hessian instead survives as the tangent ceiling and microscopic skew-square cone inside that larger class.

A second frontier, forced by transformed semi-coarse-graining in partially observed entropy-production inference [?], is renewal-level quotient geometry. There the visible law is not a static state quotient of the original continuous-time Markov chain, but a retrospective path transducer that groups completed hidden excursions. Appendix E shows that in the four-state

benchmark of [?] this transformed observer still collapses exactly to a finite parity-lifted semi-Markov visible law with a two-kernel timing module and admits an explicit exact finite CTMC completion. What remains open is the completion-independent part of the story: whether the renewal visible law itself fixes a canonical tangent ceiling and hence a canonical hidden-load operator, rather than only a completion class.

Definition 10.1 (Full nonlinear re-embedding problem). *Given a finite-state Markov Donsker–Varadhan Hessian upstairs, determine whether quotient observation remains inside the same finite-state Markov Donsker–Varadhan class at every rung.*

11 Discussion

The present paper reveals a two-layer relation between our previous work and the quotient theory. Quadratic order measures hidden-return visibility. This is already present in the finite-observation law and explains why hidden nonequilibrium structure can register strongly on a visible sector before genuinely nonlinear quotient effects appear. Beyond that, quotienting acts as a latent contraction. It preserves the tangent support and attenuates inside it, as in Corollary 6.2. The tangent ceiling is therefore a real ceiling, not a heuristic one. The first genuinely new quotient effect is not a new visible direction but a sign-definite quartic hidden-mediation subtraction. The object preserved by marginal elimination is not the Donsker–Varadhan Hessian as such, but the quotient precision itself. The Donsker–Varadhan skew-square law survives as the tangent microscopic cone and tangent ceiling inside that larger visible attenuation class, in the precise sense of Corollary 5.3.

At the symmetric level, that quotient law also has a simple observed-information reading. Once an upstairs positive form and an observation map are fixed, the visible quadratic form induced by hidden elimination is forced, not chosen, and is given by the same quotient operator $F_C(H) = (CH^{-1}C^T)^{-1}$ that appears throughout the paper. In this sense the quotient precision is naturally read as the visible symmetric information geometry of the observation map. We leave the broader relation to Fisher and Fisher–Kähler reduction to later work.

This clarifies the status of visible long-range couplings. The visible law is not a free-standing dense effective rule appearing from nowhere. It is a Green-kernel shadow of a hidden local precision geometry after elimination [20, 21, 51]. The graph-local theorems show that such shadows are controlled tails, with the universal long-range rate set by conditioning. The extremiser picture is correspondingly rigid at the level of rate and singular-spectrum shape. Corridor geometry cleanly realises the sharp universal rate class, while the synthetic battery indicates that looped topologies can improve finite-distance prefactors in the tested family without changing the rate law. The strongest many-channel threshold pattern remains asymptotically a disjoint bundle of nearly identical corridors with flat singular spectrum.

The Bell-facing lesson is correspondingly precise. Quotient geometry explains how a strictly local hidden precision can cast a visibly dense long-range shadow, but Section 8 shows that the classical frontier is a law-level gluing problem, not merely a correlator problem. On the Bell square, common visible Gaussian gluing is equivalent to remote-setting variance consistency together with Bell-locality of the arcsine lift. The Bell frontier therefore has two coordinates, not one: a variance-level obstruction invisible to correlators, and the usual arcsine-CHSH obstruction. Proposition 8.3 shows that this law frontier is strictly sharper than Bell-locality itself, and Corollary 8.4 shows that in natural lifted coordinates the full Bell-compatible region is globally hypercubic, with coordinatewise repair. Corollary 8.5 then shows that any single bounded local readout remains CHSH-local under common Schur gluing. The relevant obstruction is therefore not visible long-range coupling by itself, but failure of common gluing across contexts, or departure from a single commutative local

readout semantics.

This also points to the next foundations frontier. Once the Bell square is recognised as a compatibility problem for full visible pair laws, the same question can be asked on a graph of observation times. At that level the natural object is again the gluing frontier for the full pair laws, not merely a correlator inequality. Appendix A records the graph-indexed gluing formulation, the automatic chordal and tree gluing theorem, and the three-time triangle theorem as the first temporal clique, where the lifted compatibility region is tetrahedral. We have chosen not to fold a full temporal graph theory into the main theorem flow because the strategic value here is to keep the current contribution tight: the paper proves the quotient geometry itself, one sharp Bell-facing application, and the first temporal clique case. But the compatibility programme suggested by the present work is broader than the Bell square, and time remains the cleanest next direction.

A closely related but genuinely new extension frontier appears when the observer is a transformed path channel rather than a static state quotient. Appendix E treats the four-state transformed semi-coarse-graining benchmark of [?]. There the visible process groups completed hidden excursions into symbols H_n , so the observer is renewal-level rather than state-level. Even so, the benchmark does not explode into an uncontrolled memory sector: the transformed sequence law collapses exactly to a six-symbol parity lift, its waiting laws collapse exactly to a two-kernel timing module, and an explicit exact finite CTMC completion exists. On one such exact completion the transformed observer opens a four-dimensional active visible tangent support with a rank-two hidden-load spectrum, whereas the plain full-CG observer is only two-dimensional and is already a one-hard-direction quotient observer. What is still open is the canonical part of that renewal story: these transformed spectra are exact for a chosen completion, but completion-independence of the tangent ceiling and hidden load is not yet proved.

The visible cone beneath a fixed tangent law is also now transparent. Finite hidden sectors preserve support and realise the support-preserving interior beneath the tangent ceiling by Corollary 6.2. More sharply, Theorem 6.4, Corollaries 6.5 and 6.7, and Theorem 6.6 show that the entire class is parametrised by a unique positive hidden-load operator on the tangent support. In that intrinsic dictionary, the support fixes where visible corrections can live, the attenuation operator fixes how much survives, $\text{rank}(I - \Pi) = \text{rank}(T - X)$ gives the least hidden latent dimension, and $\log \det(I_S + \Lambda)$ measures the latent contraction volume. Taking limits adds the boundary. This explains both why the finite paper remains central and why same-dimension skew-square closure generically fails. The finite paper supplies the tangent ceiling and microscopic support cone. The quotient paper supplies the attenuated law underneath it. In this sense the right renormalisation-level object is the quotient precision geometry, not the original microscopic Donsker–Varadhan parametrisation at fixed visible dimension. Conceptually this places the visible theory closer to an information-geometric coarse description than to a fixed-dimensional microscopic parametrisation [52–54].

The empirical branch collected in Appendix D shows that this strategy already begins to work on a real threshold optical data set. Starting from raw homodyne quadrature records and the supplied $g^{(2)}(0)$ processing chain, the reproduced traces remain numerically faithful to the released derived files, and the resulting current-by-resolution summaries reveal a nontrivial threshold band rather than a smooth one-parameter crossover. In particular, the mean $g^{(2)}(0)$ map shows the expected descent from thermal-like values below threshold toward near-coherent values above threshold, the window-to-window variance of $g^{(2)}(0)$ peaks sharply in the threshold region, and the lag-1 autocorrelation develops a strong scale dependence across the same band. The screened currents therefore supply three complementary empirical readers of the transition: level, fluctuation, and memory, with 70.8 mA emerging as the clearest transition exemplar in the reduced sweep. We therefore treat

the threshold data as empirical pressure on the quotient observables themselves rather than as a claim of full hidden-law reconstruction. Table 1 together with Appendix D shows that the quartic bridge, the closure obstruction, the hidden-local decay law, and the multiscale observational programme all survive direct numerical and empirical pressure. Threshold optical monitoring is therefore a natural next testbed for the quotient geometry proved here, independent of the still-open full nonlinear Markov re-embedding problem.

For readers interested in the broader hidden-to-visible methodology rather than the present theorem domain alone, Appendix G records three bridge theorems and states the current synthesis in deliberately conservative form. In particular, Bridge A now records the inverse split for the hidden-load class: conditional on the tangent ceiling the intrinsic hidden-load operator is uniquely recoverable, while without an independently fixed ceiling the same visible law admits a full cone of ceiling-load explanations.

12 Blind digital benchmark under partial observation

This was a blind test. The hidden world was fixed first, only a small visible part was shown, predictions were locked in, and only then was the full answer revealed. The method correctly said things about what it *could not yet see*. This turns the mathematics into an instrument, using what is visible to make testable claims.

The benchmark was designed to test the operational content of the present paper without weakening the theorem boundary. Each digital world coupled three exact layers. First, a finite-state nonequilibrium backbone fixed the reduced detailed-balance geometry H_0 , the skew object \widehat{J} , and the finite Donsker–Varadhan signal $\Delta_{DV} = \frac{1}{4}\widehat{J}H_0^{-1}\widehat{J}^\top$ on the observer space. Second, an adapted visible sector U fixed the tangent ceiling $T = P_U\Delta_{DV}P_U|_U$, and the actual quotient-visible correction was generated inside the proved attenuation class

$$X = T^{1/2}(I + \Lambda)^{-1}T^{1/2}, \quad \Lambda \succeq 0.$$

Third, Bell-context families were attached in three exact types: commonly glueable families, Bell-local but variance-inconsistent families, and variance-consistent families with arcsine-CHSH violation. Hidden truth and visible challenge artefacts were separated on disk, the analysis stage read only the visible manifest, predictions were hash-locked before reveal, and reveal then checked the lock before emitting the concealed truth. In the exact track the visible payload consisted only of a stated initial observer basis together with the projected objects $Q_0^\top H_0 Q_0$ and $Q_0^\top \Delta_{DV} Q_0$; in the sampled track it consisted only of detailed-balance and nonequilibrium samples on that same visible sector.

The exact-mode hero world already shows the main effect. Figure 1 compares the locked finite prediction with the revealed truth for the retained-signal curve K_d . The corresponding relative errors were 2.61×10^{-2} for the K -curve and 1.67×10^{-3} for the detectability curve Θ_d . On the quotient side, Figure 2 shows the visible spectrum sitting beneath the tangent ceiling on the same adapted sector exactly as the theory predicts. In the scored hero world the minimum hidden dimension and the contraction volume were both recovered with zero score error, and the identity residuals for the quotient formulas remained at machine scale.

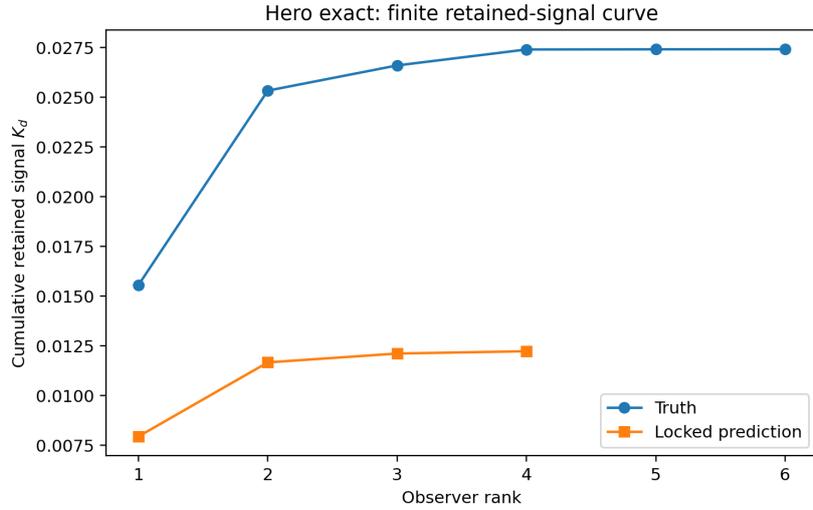


Figure 1: Exact-mode hero world. Locked finite prediction versus revealed truth for the cumulative retained-signal curve K_d .

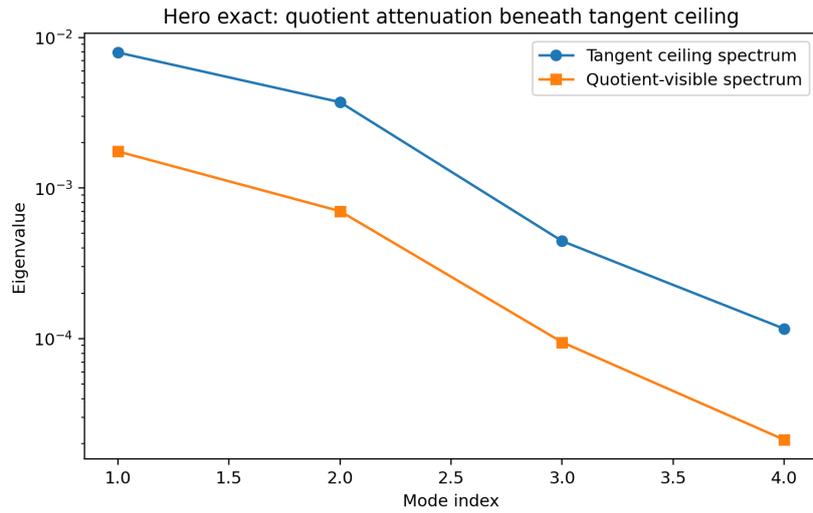


Figure 2: Exact-mode hero world. Tangent-ceiling spectrum and quotient-visible spectrum on the same adapted visible sector. The quotient-visible law sits beneath the tangent ceiling, as required by the support-preserving attenuation geometry.

The stronger test is the frozen 256-world exact ensemble. Across that ensemble the mean relative error was 1.461×10^{-1} for the finite K -curve and 2.541×10^{-2} for the detectability curve Θ_d . The quotient scores remained exact on the recorded ensemble, with d_{\min} error and contraction-volume error equal to zero throughout the machine-readable score files. Bell verdict agreement was 1.0 across all canonical and seeded families. Locality compatibility matched in 92.58% of worlds, while motif agreement was 45.70%; we therefore treat locality only as a class-level diagnostic reader, not as an exact hidden-graph reconstruction claim. Figure 3 summarises these exact-ensemble agreement rates.

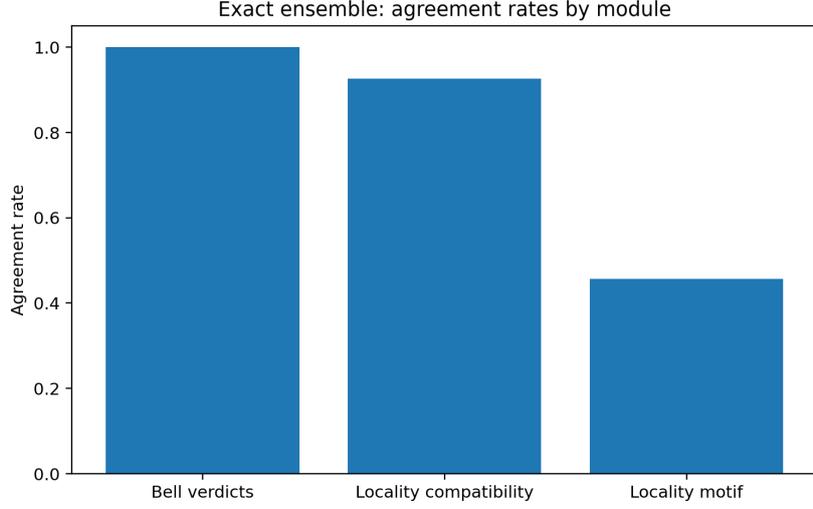


Figure 3: Exact 256-world ensemble. Bell verdict agreement is exact on the recorded ensemble; locality compatibility is strong; motif classification is intentionally weaker and remains class-level.

Two further remarks are important. First, the sampled track is a robustness layer rather than the primary claim. On the sampled hero world the finite module remained stable, with K -curve relative error 8.46×10^{-2} and Θ -curve relative error 3.56×10^{-2} , while quotient scores again remained exact; however, the glueable Bell families were missed in the sampled Bell layer, so the sampled benchmark should be read as support for interface robustness rather than as the cleanest demonstration of the law-level frontier. Second, adaptive reveal is already informative but should still be described cautiously. On a 64-world exact panel, the oracle theorem-target reveal improved the mean second-step raw revealed signal over a random baseline by 6.39×10^{-2} and the corresponding detectability score by 1.64×10^{-2} . We therefore treat adaptive reveal as a diagnostic of what the observation ladder can target next, not yet as a fully deployed blind acquisition policy. Additional distributions and sampled-mode benchmark diagnostics are collected in Appendix F.

Taken together, the blind benchmark does not enlarge the theorem domain of the paper, but it does make the existing theory operational. It shows that the finite tangent law, the quotient-visible attenuation geometry, and the Bell-square law frontier are not merely formal descriptions of hidden elimination. In the theorem-led track they can be recovered from a restricted visible sector under a real lock-then-reveal protocol, with the hidden world checked only after the predictions have been committed.

A Bell-frontier details

A.1 Gaussian gluing on the Bell square

We first isolate the Gaussian completion problem behind Theorem 8.2. Write

$$\Sigma_{xy} = \begin{pmatrix} \sigma_{A_x}^2 & c_{xy} \\ c_{xy} & \sigma_{B_y}^2 \end{pmatrix}, \quad K_{xy} := \frac{c_{xy}}{\sigma_{A_x} \sigma_{B_y}},$$

where the one-site variances are assumed to satisfy the remote-setting consistency conditions (8.3). Let $\mathcal{E}_{2,2}$ denote the bipartite elliptope, namely the set of all 2×2 matrices of the form

$$L_{xy} = u_x \cdot v_y \quad \text{with } u_x, v_y \text{ unit vectors in a real Hilbert space.}$$

Proposition A.1 . *Under the variance-consistency conditions (8.3), the following are equivalent.*

1. *There exists a centred Gaussian law on (A_0, A_1, B_0, B_1) whose (A_x, B_y) marginals are exactly Σ_{xy} .*
2. *The normalised cross block $K = (K_{xy})$ belongs to $\mathcal{E}_{2,2}$.*

If the resulting full covariance is positive definite, this is equivalent to common Schur gluing of the visible pair precisions.

Proof. Assume first that a common centred Gaussian law exists. After standardising each coordinate by its one-site standard deviation, one obtains a correlation matrix on (A_0, A_1, B_0, B_1) with diagonal entries equal to 1 and cross block equal to K . Every correlation matrix is positive semidefinite and therefore a Gram matrix of unit vectors. Hence $K \in \mathcal{E}_{2,2}$.

Conversely, suppose $K_{xy} = u_x \cdot v_y$ for unit vectors u_x, v_y . Form the block Gram matrix

$$R := \begin{pmatrix} (u_x \cdot u_{x'})_{x,x'} & (u_x \cdot v_y)_{x,y} \\ (v_y \cdot u_x)_{y,x} & (v_y \cdot v_{y'})_{y,y'} \end{pmatrix} \succeq 0.$$

Now scale R by the diagonal matrix

$$D := \text{diag}(\sigma_{A_0}, \sigma_{A_1}, \sigma_{B_0}, \sigma_{B_1}).$$

Then $\Sigma_* := DRD \succeq 0$ is a covariance matrix whose (A_x, B_y) marginal block is exactly Σ_{xy} . If $\Sigma_* \succ 0$, inversion yields a visible precision $Q_* = \Sigma_*^{-1}$, and each pair precision is the corresponding Schur complement block. This is exactly common Schur gluing.

Proof of Theorem 8.2. By Proposition A.1, common Gaussian gluing is equivalent to variance consistency together with membership of the normalised cross block K in $\mathcal{E}_{2,2}$. For pair precisions,

$$K_{xy} = \frac{c_{xy}}{\sqrt{v_{xy}^A v_{xy}^B}} = \frac{-b_{xy}}{\sqrt{a_{xy} d_{xy}}} = \rho_{xy},$$

which is (8.2).

Now define

$$E_{xy} := \frac{2}{\pi} \arcsin(K_{xy}).$$

Standard Bell-square background facts identify the frontier in three equivalent ways: Fine's theorem characterises Bell-locality of E by the CHSH inequalities; the Gaussian sign-correlation formula shows that E is the correlator table of the signs of a Gaussian family with correlation matrix having cross block K ; and the Tsirelson–Landau–Masanes characterisation identifies the same Bell-square correlator set with the bipartite elliptope $\mathcal{E}_{2,2}$ after the sine or arcsine lift. Hence $K \in \mathcal{E}_{2,2}$ if and only if the lifted correlator table E is Bell-local, equivalently if and only if the four arcsine CHSH inequalities (8.4) to (8.7) hold. Combining this with Proposition A.1 proves the theorem.

A.2 Strict refinement over Bell-locality

Proposition 8.3 already gives the minimal counterexample. The point is structural: Bell inequalities constrain only the binary-output correlator shadow, while common Gaussian gluing constrains the full pair law. The diagonal sector is therefore visible to the quotient frontier even when the Bell shadow is trivial.

A convenient quantitative summary is to separate the two obstructions. Define

$$\varepsilon_{\text{var}} := \max\{|v_{00}^A - v_{01}^A|, |v_{10}^A - v_{11}^A|, |v_{00}^B - v_{10}^B|, |v_{01}^B - v_{11}^B|\}, \quad (\text{A.1})$$

and let $\varepsilon_{\text{CHSH}}$ be the excess of the largest left-hand side of (8.4) to (8.7) above π , truncated below at zero. Then common gluing on the Bell square is equivalent to the simultaneous vanishing of both obstructions:

$$\varepsilon_{\text{var}} = 0, \quad \varepsilon_{\text{CHSH}} = 0.$$

This packages the Bell-square frontier as a two-coordinate compatibility problem: one coordinate is invisible to Bell correlators, while the other is the usual arcsine correlator obstruction.

A.3 Low-rank mediated Gaussian sign shadows

The Bell-local shadow can itself be realised by a very small mediated Gaussian model.

Corollary A.2. *Let $E = (E_{xy})$ be a Bell-local 2×2 correlator table, and set*

$$K_{xy} := \sin\left(\frac{\pi}{2} E_{xy}\right).$$

Then there exist unit vectors $u_0, u_1, v_0, v_1 \in \mathbb{R}^2$ and a standard Gaussian $Z \sim N(0, I_2)$ such that

$$U_x := u_x \cdot Z, \quad V_y := v_y \cdot Z, \quad \mathbb{E}[\text{sgn}(U_x) \text{sgn}(V_y)] = E_{xy} \quad \text{for all } x, y.$$

Thus every Bell-local Bell-square correlator table already sits inside a rank-two Gaussian sign-shadow.

Proof. By Theorem 8.2, Bell-locality of E is equivalent to $K \in \mathcal{E}_{2,2}$. In the Bell square one may choose a planar realisation of the corresponding unit vectors, so let $u_x, v_y \in \mathbb{R}^2$ satisfy $u_x \cdot v_y = K_{xy}$. For $Z \sim N(0, I_2)$, the Gaussian variables $U_x = u_x \cdot Z$ and $V_y = v_y \cdot Z$ have correlations K_{xy} . The Gaussian sign-correlation formula gives

$$\mathbb{E}[\text{sgn}(U_x) \text{sgn}(V_y)] = \frac{2}{\pi} \arcsin(K_{xy}) = E_{xy},$$

as required.

A.4 Bell cube normal form and exact repair

The Bell square admits a global normal form in lifted coordinates. This globalises the one-face repair formulas and makes the Bell witness geometry completely explicit.

Theorem A.3 . *Let S be the 4×4 matrix whose rows are the four CHSH sign vectors*

$$s_1 = (1, 1, 1, -1), \quad s_2 = (1, 1, -1, 1), \quad s_3 = (1, -1, 1, 1), \quad s_4 = (-1, 1, 1, 1).$$

Then

$$S^\top S = 4I_4.$$

For

$$y := \frac{1}{\pi} S \phi, \quad \phi := (\phi_{00}, \phi_{01}, \phi_{10}, \phi_{11})^\top,$$

the Bell-compatible region in arcsine coordinates becomes exactly

$$|y_i| \leq 1, \quad i = 1, \dots, 4.$$

Equivalently, the Bell square is globally the cube $[-1, 1]^4$ in the lifted coordinates y . Moreover,

$$\phi = \frac{\pi}{4} S^\top y, \quad \|\delta\phi\|^2 = \frac{\pi^2}{4} \|\delta y\|^2.$$

Hence the exact Euclidean nearest-point map is coordinatewise clipping,

$$y_i^* = \text{clip}(y_i, -1, 1), \quad \phi^* = \frac{\pi}{4} S^\top y^*,$$

and the exact squared Euclidean distance to the Bell frontier is

$$\text{dist}(\phi, \mathcal{B})^2 = \frac{\pi^2}{4} \sum_{i=1}^4 (|y_i| - 1)_+^2.$$

If only one lifted CHSH coordinate violates the cube, with defect $e > 0$, then the exact repair is

$$\delta\phi^* = -\frac{e\pi}{4} s_i, \quad \|\delta\phi^*\|_\infty = \frac{e\pi}{4}.$$

Proof. The four inequalities (8.4) to (8.7) are precisely

$$|s_i \cdot \phi| \leq \pi, \quad i = 1, \dots, 4.$$

Since the rows of S are the vectors s_i , this is equivalent to $|y_i| \leq 1$ for all i , so the compatible region is the cube $[-1, 1]^4$ in the lifted coordinates. Direct computation gives $S^\top S = 4I_4$, hence $S^{-1} = \frac{1}{4} S^\top$ and therefore

$$\phi = \frac{\pi}{4} S^\top y.$$

The Euclidean metric transforms accordingly,

$$\|\delta\phi\|^2 = \frac{\pi^2}{16} \delta y^\top S S^\top \delta y = \frac{\pi^2}{4} \|\delta y\|^2,$$

because $S S^\top = 4I_4$. Euclidean nearest-point repair in y is therefore just coordinatewise

clipping to the cube, and transforming back gives the stated projection formula and distance formula. If only one coordinate exceeds the cube by e , then clipping changes only that coordinate by e , so

$$\delta\phi^* = \frac{\pi}{4} S^\top(-ee_i) = -\frac{e\pi}{4} s_i,$$

which also gives $\|\delta\phi^*\|_\infty = e\pi/4$.

Corollary A.4 . *For any covector $u \in \mathbb{R}^4$, the Bell support function in arcsine coordinates is*

$$h_{\mathcal{B}}(u) := \sup_{\phi \in \mathcal{B}} u \cdot \phi = \frac{\pi}{4} \|Su\|_1.$$

Proof. Using $\phi = (\pi/4)S^\top y$ and $|y_i| \leq 1$, one has

$$h_{\mathcal{B}}(u) = \sup_{|y_i| \leq 1} \frac{\pi}{4} (Su) \cdot y = \frac{\pi}{4} \sum_{i=1}^4 |(Su)_i| = \frac{\pi}{4} \|Su\|_1.$$

A.5 Observation graphs and the first temporal clique

The Bell square is the first nonchordal compatibility graph treated in the main text. The same gluing question may be asked on any finite observation graph, with Bell and temporal contexts appearing as two graph-indexed instances of one Gaussian pair-law completion problem. The next proposition isolates the law-level object.

Proposition A.5 (Graph-indexed Gaussian gluing). *Let $G = (V, E)$ be a finite graph. For each edge $\{i, j\} \in E$, let*

$$\Sigma_{ij} = \begin{pmatrix} v_i^{(ij)} & c_{ij} \\ c_{ij} & v_j^{(ij)} \end{pmatrix} \in \text{SPD}(2), \quad \rho_{ij} := \frac{c_{ij}}{\sqrt{v_i^{(ij)} v_j^{(ij)}}}.$$

Then there exists a centred Gaussian law on $(X_i)_{i \in V}$ whose (i, j) marginal is exactly Σ_{ij} for every edge if and only if both of the following hold:

(i) *vertex variance consistency, namely there exist numbers $v_i > 0$ with*

$$v_i^{(ij)} = v_i \quad \text{for every vertex } i \text{ and every incident edge } \{i, j\};$$

(ii) *correlation completion, namely there exists a correlation matrix $R \succeq 0$ on V such that*

$$R_{ii} = 1, \quad R_{ij} = \rho_{ij} \quad \text{for all } \{i, j\} \in E.$$

Proof. If a global centred Gaussian law exists, its one-site variances are independent of which edge marginal is inspected, so (i) holds. Dividing the global covariance by the diagonal standard-deviation matrix gives a correlation matrix R with the stated edge restrictions, so (ii) holds.

Conversely, assume (i) and (ii). Let $D = \text{diag}(v_i)_{i \in V}$ and set $\Sigma := D^{1/2} R D^{1/2}$. Then $\Sigma \succeq 0$

and every edge marginal is

$$\Sigma_{\{i,j\}} = \begin{pmatrix} v_i & \sqrt{v_i v_j} R_{ij} \\ \sqrt{v_i v_j} R_{ij} & v_j \end{pmatrix} = \begin{pmatrix} v_i^{(ij)} & c_{ij} \\ c_{ij} & v_j^{(ij)} \end{pmatrix} = \Sigma_{ij},$$

because $R_{ij} = \rho_{ij} = c_{ij} / \sqrt{v_i^{(ij)} v_j^{(ij)}}$ and the visible variances agree by (i).

The contextual object is therefore the edge-correlation completion set

$$\mathcal{E}(G) := \{\rho \in (-1, 1)^E : \rho \text{ extends to a correlation matrix on } V\}.$$

Bell and temporal inequalities sit below $\mathcal{E}(G)$ as readout shadows. At law level the primary obstruction is gluing failure for the full pair laws.

The next theorem identifies the easy graph class.

Theorem A.6 (Chordal clique gluing). *Let G be a chordal graph with maximal cliques C_1, \dots, C_m . For each clique C_α , let ν_{C_α} be a centred Gaussian law with positive definite clique covariance. Assume that whenever two neighbouring cliques in a clique tree meet in a separator $S = C_\alpha \cap C_\beta$, the induced S -marginals of ν_{C_α} and ν_{C_β} agree. Then there exists a centred Gaussian law on all of V whose C_α -marginal is exactly ν_{C_α} for every α .*

Proof. The key separator-gluing step is explicit. If centred Gaussian laws on index sets U and W agree on the overlap $S = U \cap W$, then

$$p(x_{U \cup W}) = \frac{p_U(x_U) p_W(x_W)}{p_S(x_S)}$$

defines a centred Gaussian law on $U \cup W$ with the prescribed U - and W -marginals. The quotient is well-defined because the overlap marginal agrees, and direct integration recovers the stated marginals.

Now choose a clique tree for G and argue by induction on the number of cliques. For one clique there is nothing to prove. For more than one clique, remove a leaf clique C_ℓ with separator $S = C_\ell \cap C'$. By induction the remaining clique family admits a common Gaussian law. Gluing that law to ν_{C_ℓ} across S by the separator formula above produces the desired global Gaussian law.

In particular, if G is a tree, then the maximal cliques are just edges. Hence edgewise positive pair laws together with vertex variance consistency are sufficient for common Gaussian gluing. Trees therefore have no further law-level obstruction. The first genuinely nontrivial temporal obstruction appears when one passes from a tree to the first temporal clique.

The first fully observed temporal case is the triangle K_3 . Because it is chordal, the correlation-completion condition reduces to positivity of the single 3×3 correlation matrix.

Theorem A.7 (Exact three-time triangle theorem). Let $\Sigma_{01}, \Sigma_{12}, \Sigma_{02} \in \text{SPD}(2)$ be centred Gaussian pair laws with entries

$$\Sigma_{ij} = \begin{pmatrix} v_i^{(ij)} & c_{ij} \\ c_{ij} & v_j^{(ij)} \end{pmatrix}, \quad \rho_{ij} := \frac{c_{ij}}{\sqrt{v_i^{(ij)} v_j^{(ij)}}}.$$

There exists a centred Gaussian law on (X_0, X_1, X_2) with exactly these three pair marginals if and only if

$$v_0^{(01)} = v_0^{(02)}, \quad v_1^{(01)} = v_1^{(12)}, \quad v_2^{(02)} = v_2^{(12)},$$

and

$$R = \begin{pmatrix} 1 & \rho_{01} & \rho_{02} \\ \rho_{01} & 1 & \rho_{12} \\ \rho_{02} & \rho_{12} & 1 \end{pmatrix} \succeq 0.$$

Equivalently,

$$1 + 2\rho_{01}\rho_{02}\rho_{12} - \rho_{01}^2 - \rho_{02}^2 - \rho_{12}^2 \geq 0. \quad (\text{A.2})$$

Setting

$$\phi_{ij} := \arcsin(\rho_{ij}), \quad y_{ij} := \frac{2}{\pi}\phi_{ij},$$

the lifted compatibility region is exactly the tetrahedron cut out by

$$1 + y_{01} + y_{02} + y_{12} \geq 0, \quad (\text{A.3})$$

$$1 + y_{01} - y_{02} - y_{12} \geq 0, \quad (\text{A.4})$$

$$1 - y_{01} + y_{02} - y_{12} \geq 0, \quad (\text{A.5})$$

$$1 - y_{01} - y_{02} + y_{12} \geq 0. \quad (\text{A.6})$$

Proof. The gluing criterion is Proposition A.5 specialised to $G = K_3$. For the determinant form, positivity of R is equivalent to nonnegativity of its determinant because the 1×1 and 2×2 principal minors are already positive. A direct computation gives (A.2).

For the lifted form, write $\rho_{ij} = \sin \phi_{ij}$ with $\phi_{ij} \in [-\pi/2, \pi/2]$. Then

$$\det R = 4 \prod_{\sigma \in \{(+,+,+), (+,-,-), (-,+,-), (-,-,+)\}} \sin\left(\frac{\pi/2 + \sigma_1\phi_{01} + \sigma_2\phi_{02} + \sigma_3\phi_{12}}{2}\right).$$

Because each factor lies in $[0, \pi]$ precisely when the corresponding linear form is nonnegative, and because $\phi_{ij} = (\pi/2)y_{ij}$, the determinant condition is equivalent to (A.3)–(A.6).

Corollary A.8 (Temporal sign-shadow tetrahedron). Let (X_0, X_1, X_2) be any centred Gaussian triple satisfying Theorem A.7, and define $S_i := \text{sgn}(X_i)$. Then

$$\mathbb{E}[S_i S_j] = \frac{2}{\pi} \arcsin(\rho_{ij}) = y_{ij},$$

so the lifted tetrahedron (A.3)–(A.6) is exactly the classical three-variable correlation polytope for the sign readout. In particular, the first temporal lifted shadow is the Leggett–Garg tetrahedron attached to the three-time clique.

Proof. For a centred Gaussian pair with correlation ρ_{ij} , Sheppard's formula gives

$$\mathbb{E}[\text{sgn}(X_i) \text{sgn}(X_j)] = \frac{2}{\pi} \arcsin(\rho_{ij}).$$

Substituting into Theorem A.7 identifies the lifted inequalities with the classical three-variable correlation tetrahedron.

The temporal tree-versus-clique divide is therefore now visible. Trees glue automatically after repeated-marginal consistency, but they need not determine a unique global law.

Proposition A.9 (Path completion interval and nonuniqueness). *Fix unit variances on the path $0 - 1 - 2$ and edge correlations $\rho_{01}, \rho_{12} \in (-1, 1)$. A Gaussian completion exists for every*

$$\rho_{02} \in \left[\rho_{01}\rho_{12} - \sqrt{(1 - \rho_{01}^2)(1 - \rho_{12}^2)}, \rho_{01}\rho_{12} + \sqrt{(1 - \rho_{01}^2)(1 - \rho_{12}^2)} \right].$$

This interval is nondegenerate whenever $|\rho_{01}| < 1$ and $|\rho_{12}| < 1$. The distinguished tree-Markov construction selects the centre point $\rho_{02} = \rho_{01}\rho_{12}$.

Proof. A completion with correlation matrix

$$R = \begin{pmatrix} 1 & \rho_{01} & \rho_{02} \\ \rho_{01} & 1 & \rho_{12} \\ \rho_{02} & \rho_{12} & 1 \end{pmatrix}$$

exists if and only if $R \succeq 0$. Since the 2×2 principal minors are positive, this is equivalent to

$$1 + 2\rho_{01}\rho_{12}\rho_{02} - \rho_{01}^2 - \rho_{12}^2 - \rho_{02}^2 \geq 0.$$

This is a quadratic inequality in ρ_{02} with roots

$$\rho_{01}\rho_{12} \pm \sqrt{(1 - \rho_{01}^2)(1 - \rho_{12}^2)},$$

so the admissible interval is exactly the one stated. Its midpoint is $\rho_{01}\rho_{12}$.

The same graph-indexed viewpoint also makes the strict Bell refinement structural rather than accidental.

Proposition A.10 (Fibre crossing at fixed correlation shadow). *Let G have a vertex of degree at least 2, and let $\rho \in \mathcal{E}(G)$ lie in the interior of the correlation-completion set. Then there exist both gluable and non-gluable full edge-law families having exactly the same edge-correlation data ρ .*

Proof. Choose any centred Gaussian completion with unit variances, so each edge block is

$$\Sigma_{ij} = \begin{pmatrix} 1 & \rho_{ij} \\ \rho_{ij} & 1 \end{pmatrix}$$

and the family is gluable. Now pick a vertex i used in at least two edges and choose one incident edge $\{i, j\}$. Replace only the variance of i on that edge by a number $\lambda \neq 1$, and replace the covariance on that edge by $\rho_{ij}\sqrt{\lambda}$. The normalised correlation on that edge

remains ρ_{ij} , and all other edge correlations are untouched, so the edge-correlation shadow is unchanged. But the variance of vertex i is now inconsistent across its incident edges, so Proposition A.5 forbids common gluing.

Thus Bell correlators do not exhaust the law-level frontier even before one passes to binary readout inequalities. They are shadows of a fuller pair-law compatibility problem, and the Bell square treated in the main text is the first nonchordal instance of that broader contextual branch.

B Scalar summaries of the Schur gap

B.1 Scalar mediation invariants

Definition B.1 (Hidden-mediation observables). For the Schur gap $G_U \succeq 0$ define

$$\Gamma_{\text{tr}} := \text{Tr}(G_U), \quad \Gamma_{\text{op}} := \|G_U\|, \quad \Gamma_{\text{det}} := \log \det H_U^{\text{comp}} - \log \det H_U^{\text{quot}}. \quad (\text{B.1})$$

Proposition B.2 . The quantities in (B.1) are all nonnegative.

Proof. Since $G_U \succeq 0$, its trace and operator norm are nonnegative. Also $H_U^{\text{comp}} \succeq H_U^{\text{quot}} \succ 0$ by Proposition 4.1, so $\det H_U^{\text{comp}} \geq \det H_U^{\text{quot}}$, hence $\Gamma_{\text{det}} \geq 0$.

C Quotient adjacency and visible fill-in diagnostics

C.1 Fibre-induced visible adjacency

To turn adjacency loss into an intrinsic quotient invariant, the visible mask must be induced by quotient geometry rather than chosen externally.

Definition C.1 (Fibre distance and quotient adjacency). Let $\pi : X \rightarrow Y$ be a quotient of a metric graph or metric state space (X, d_X) . Define the lifted fibre distance by

$$d_\pi(y, y') := \inf\{d_X(u, v) : u \in \pi^{-1}(y), v \in \pi^{-1}(y')\}. \quad (\text{C.1})$$

If upstairs locality has interaction range r , define visible adjacency by

$$y \sim_\pi y' \iff d_\pi(y, y') \leq r. \quad (\text{C.2})$$

Let M_π denote the mask that keeps only matrix entries consistent with (C.2).

Definition C.2 (Locality-defect operator). Let $H_{\text{vis}} \in \text{SPD}(Y)$ be a visible precision and let the visible Hilbert structure on Y be fixed once and for all. The quotient locality-defect operator is

$$N_{\pi} := H_{\text{vis}} - M_{\pi}(H_{\text{vis}}). \quad (\text{C.3})$$

Its singular values, stable rank, and retained Frobenius mass are called the quotient adjacency-loss diagnostics. These are natural numerical readers of Schur-complement fill-in and spectral complexity in sparse positive systems [42, 47].

Proposition C.3 (Intrinsic diagnostics). Relative to a fixed visible Hilbert structure and interaction range r , the singular spectrum of N_{π} and the scalar mediation observables of Section 4 are intrinsic to the quotient construction.

Proof. Once d_{π} and r are fixed, the mask M_{π} is determined by the quotient itself. The operator N_{π} is then determined by H_{vis} and the quotient geometry. The scalar observables depend only on the Schur gap, which is itself determined by the quotient-compression comparison.

D Real-time threshold monitoring of $g^{(2)}(0)$

This section replaces analogies with a direct threshold-data stress test. We used the released quadrature data behind the real-time homodyne monitoring experiment of Lüders, Thewes, and Aßmann [55]. Each matched raw file contains a quadrature array X and is paired to derived arrays ($g2vec, ada, times$) at the published resolution 1000. Using the supplied processing formulas, we reproduced the released derived traces from the raw quadratures to floating-point tolerance across the available current sweep. We then ran a reduced multiresolution screen on currents 68.0, 69.0, 70.0, 70.8, 71.6, 72.5, 78.0, and 90.0 mA at resolutions 250, 1000, and 4000.

The aim of this appendix is deliberately narrow. We ask whether quotient-style observables reveal a structured threshold band in a real optical data set. We therefore retain only the outputs that are already numerically stable in the fast pass: the mean of $g^{(2)}(0)$, the window-to-window variance of $g^{(2)}(0)$ at the nondegenerate resolutions 250 and 1000, and the lag-1 autocorrelation of the $g^{(2)}(0)$ time series. These give three complementary readers of the threshold band: level, fluctuation, and memory. We record the photon-number proxy N and the mean quadrature variance as consistency readers, but we do not use exploratory HMM fits in the present paper.

D.1 Multiresolution threshold summary

The fast pass isolates three complementary empirical readers of the threshold band: the level of $g^{(2)}(0)$, the window-to-window fluctuation strength of $g^{(2)}(0)$, and the lag-1 memory of the same windowed series. Figure 4 collects the level and memory summaries, while Figure 5 and Table 2 record the fluctuation diagnostic.

The mean $g^{(2)}(0)$ heatmap shows a clear crossover rather than a smooth wash. At resolution 1000 the mean values are approximately 2.021 at 68.0 mA, 2.019 at 70.0 mA, 1.972 at 70.8 mA, 1.875 at 71.6 mA, 1.485 at 72.5 mA, 1.074 at 78.0 mA, and 1.063 at 90.0 mA. The lag-1 autocorrelation heatmap shows that the same band is structured dynamically as well:

below threshold the windowed $g^{(2)}(0)$ series are only weakly correlated, while around and above threshold the memory structure becomes strong and resolution dependent. In this reduced sweep, 70.8 mA is the clearest transition exemplar because it combines an elevated mean, the strongest fluctuation peak, and a nontrivial lag-1 pattern distinct from both the lower-current and clearly above-threshold cases.

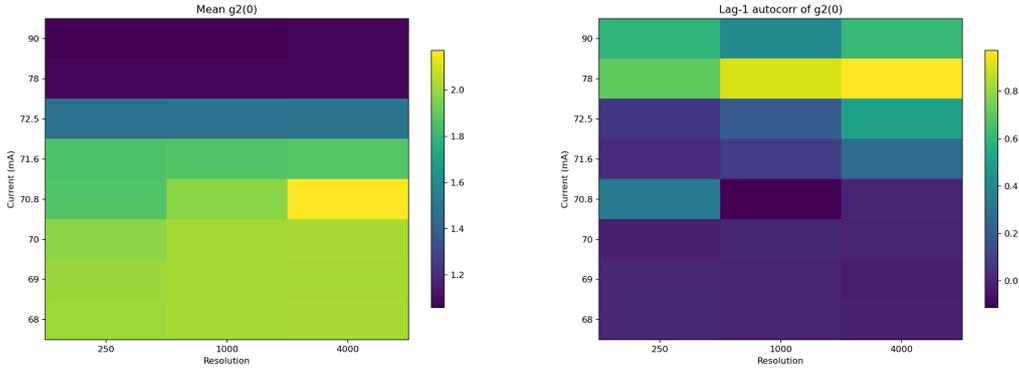


Figure 4: Threshold-monitoring branch. Left: mean $g^{(2)}(0)$ across current and resolution. Right: lag-1 autocorrelation of the corresponding windowed $g^{(2)}(0)$ series. The transition is visible both in level and in temporal structure.

The fluctuation reader proved equally informative. Figure 5 shows the window-to-window variance of the same $g^{(2)}(0)$ series, and Table 2 records the threshold-band values explicitly. The variance peaks sharply in the transition region. The largest value occurs at 70.8 mA, where the variance is 0.0715 at resolution 250 and 0.0393 at resolution 1000. The neighbouring points at 70.0 and 71.6 mA are already smaller, and by 72.5 mA the values have fallen to 0.00758 and 0.00159. Far above threshold, at 78.0 mA, the corresponding variances are only 6.73×10^{-4} and 1.40×10^{-4} . The apparent collapse of the 4000-resolution variance reader in the fast pass is not interpreted structurally: at that coarsest resolution the number of windows is too small for a stable variance estimate, so only the nondegenerate 250 and 1000 columns are used in the interpretation.

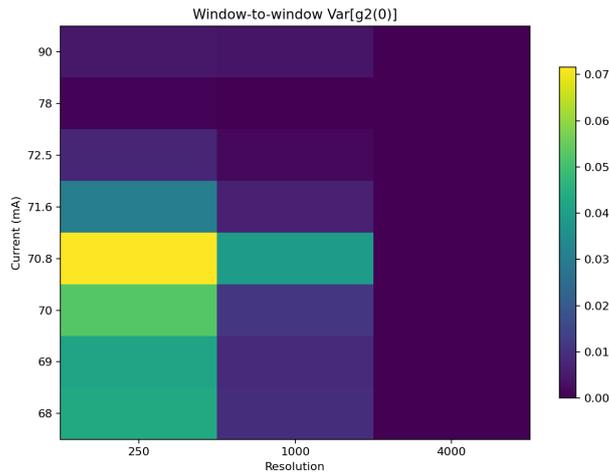


Figure 5: Window-to-window variance of $g^{(2)}(0)$ across current and resolution. The fluctuation peak localises sharply in the threshold band, with 70.8 mA giving the strongest value in the reduced sweep. The 4000-resolution column is sample-limited and is not used structurally in the interpretation.

Current (mA)	Mean $g^{(2)}(0)$ at 1000	Var[$g^{(2)}(0)$] at 250	Var[$g^{(2)}(0)$] at 1000	Interpretation
68.0	2.0211	0.04344	0.00934	Thermal-like mean with substantial fast fluctuations.
70.0	2.0197	0.05259	0.01143	Threshold band with enhanced variability.
70.8	1.9716	0.07154	0.03934	Strongest fluctuation point in the screened set.
71.6	1.8751	0.03049	0.00651	Still super-Poissonian, but already contracting.
72.5	1.4846	0.00758	0.00159	Clear drop in fluctuation strength across the crossover.
78.0	1.0738	0.00067	0.00014	Near-coherent mean with weak residual fluctuations.
90.0	1.0632	0.00462	0.00394	Well above threshold, with low mean but nontrivial coarse memory.

Table 2: Fast-pass threshold summary for the released homodyne data. The threshold band is singled out not only by the level of $g^{(2)}(0)$ but also by the magnitude of its window-to-window fluctuations.

The additional consistency readers are aligned with this picture. The photon-number proxy N and the mean quadrature variance track one another closely, as expected from the underlying homodyne construction, but they are not numerically identical. Their role here is auxiliary: they show that the threshold-band structure is not a one-column artefact of the $g^{(2)}(0)$ estimator.

D.2 Interpretation

We do not claim a full hidden-law reconstruction from the threshold data, and the present appendix does not rely on exploratory HMM fits whose first-pass outputs were not yet stable enough for inclusion. What it does show is already useful. Starting from released raw quadratures and a reproduced published processing chain, the quotient-facing observables pick out a structured transition band in a real optical system. The threshold region is singled out not only by elevated mean $g^{(2)}(0)$, but also by sharply enhanced $g^{(2)}(0)$ variability and by nontrivial temporal structure across observation scales. The reduced sweep therefore already supports the narrower and defensible claim used in the main text: threshold optical monitoring exhibits a nontrivial observational ladder on which quotient-style readers are informative, with 70.8 mA serving as the clearest transition exemplar in the present data slice. Had the current-by-resolution maps been flat, had the fluctuation peak failed to localise near threshold, or had the memory reader shown no scale dependence, the empirical motivation for a multiscale quotient reading would have weakened sharply. Instead, the data sharpen the empirical story without forcing a stronger hidden-law claim than the present fast pass can

justify.

E Renewal-visible transformed semi-coarse-graining benchmark

This appendix records a concrete visible law that lies just beyond the static quotient theorem domain of the main text. The source model is the four-state benchmark introduced in the transformed semi-coarse-graining study of entropy-production estimation under partial observation by Kapustin, Ghosal, and Bisker [?]. The benchmark is useful here because the transformed observer is not a static quotient of the original continuous-time Markov chain, yet its visible law is still exactly tractable.

E.1 Four-state benchmark and transformed observer

Let the original continuous-time Markov chain have states $\{1, 2, 3, 4\}$ with hidden block $H = \{3, 4\}$ and rates

$$1 \rightarrow 2 : 2e^x, \quad 1 \rightarrow 4 : 1, \quad 2 \rightarrow 1 : 3e^{-x}, \quad 2 \rightarrow 3 : 2, \quad 2 \rightarrow 4 : 35, \quad (\text{E.1})$$

$$3 \rightarrow 2 : 50, \quad 3 \rightarrow 4 : 0.7, \quad 4 \rightarrow 1 : 8, \quad 4 \rightarrow 2 : 0.2, \quad 4 \rightarrow 3 : 75. \quad (\text{E.2})$$

Following [?], a transformed semi-coarse-grained trajectory groups each completed run of consecutive visits to the hidden macrostate H into a symbol H_n , where n is the number of hidden visits in that run and the waiting time of H_n is the sum of the waiting times in those n visits. To keep the previous visible state explicit, write

$$A_n := (1, H_n), \quad B_n := (2, H_n).$$

Define

$$r_3 := \frac{0.7}{50.7}, \quad r_4 := \frac{75}{83.2}, \quad r := r_3 r_4. \quad (\text{E.3})$$

Numerically,

$$r \approx 0.0124459490, \quad \Pr(n \leq 2 \mid \text{enter hidden}) = 1 - r \approx 0.987554051, \quad \Pr(n > 4 \mid \text{enter hidden}) = r^2 \approx 1.5$$

Thus the transformed observer is already sharply concentrated on very short hidden excursions.

Proposition E.1 (Exact parity collapse of the transformed sequence law). For the benchmark (E.1) to (E.3), the countable augmented transformed sequence law on

$$\{1, 2\} \cup \{A_n, B_n : n \geq 1\}$$

collapses exactly, at sequence level, to the six-symbol first-order Markov chain

$$\mathcal{Y}_{\text{par}} = \{1, 2, A_o, A_e, B_o, B_e\}, \quad (\text{E.4})$$

where parity records whether the completed hidden run has odd or even length. More precisely, conditioned on entry from visible state 1,

$$P_o^{(1)} = \frac{1 - r_4}{1 - r}, \quad P_e^{(1)} = \frac{r_4(1 - r_3)}{1 - r}, \quad (\text{E.5})$$

and conditioned on entry from visible state 2,

$$P_o^{(2)} = \frac{\frac{2}{37}(1 - r_3) + \frac{35}{37}(1 - r_4)}{1 - r}, \quad P_e^{(2)} = \frac{\frac{2}{37}r_3(1 - r_4) + \frac{35}{37}r_4(1 - r_3)}{1 - r}. \quad (\text{E.6})$$

The next visible state depends on the completed excursion only through the previous visible state and this odd or even parity label.

Proof. Inside the hidden block the chain alternates between states 3 and 4 until escape to the visible sector. For entry through state 4 the completed excursion exits from hidden state 4 exactly when the number of hidden visits is odd and from hidden state 3 exactly when it is even. The probabilities (E.5) follow by summing the corresponding geometric series. For entry from visible state 2, the first hidden state is a fixed mixture of 3 and 4, and summing the odd and even geometric series gives (E.6). Since the visible return law depends only on the hidden exit microstate, it depends only on previous visible state and parity.

The corresponding visible return probabilities are explicit. In particular,

$$A_o \rightarrow 1 \text{ with probability } \frac{8}{8.2}, \quad A_o \rightarrow 2 \text{ with probability } \frac{0.2}{8.2}, \quad A_e \rightarrow 2 \text{ with probability } 1, \quad (\text{E.7})$$

and for the B -symbols,

$$\Pr(1 | B_o) \approx 0.6207014153, \quad \Pr(2 | B_o) \approx 0.3792985847, \quad (\text{E.8})$$

$$\Pr(1 | B_e) \approx 8.5325870 \times 10^{-5}, \quad \Pr(2 | B_e) \approx 0.9999146741. \quad (\text{E.9})$$

Thus the transformed observer exposes the hidden $3 \leftrightarrow 4$ mode as an almost visible parity channel.

Theorem E.2 (Two-kernel timing collapse). *Let*

$$L_i(s) := \frac{\lambda_i}{\lambda_i + s}, \quad \lambda_3 = 50.7, \quad \lambda_4 = 83.2.$$

Then the waiting-time transforms of the parity symbols lie exactly in the two-dimensional span generated by

$$G_{4,o}(s) = \frac{(1-r)L_4(s)}{1-rL_3(s)L_4(s)}, \quad G_{4,e}(s) = \frac{(1-r)L_3(s)L_4(s)}{1-rL_3(s)L_4(s)}. \quad (\text{E.10})$$

More precisely,

$$G_{3,o}(s) = \frac{\lambda_3}{\lambda_4} G_{4,o}(s) + \left(1 - \frac{\lambda_3}{\lambda_4}\right) G_{4,e}(s) = \frac{39}{64} G_{4,o}(s) + \frac{25}{64} G_{4,e}(s), \quad (\text{E.11})$$

$$G_{B_o}(s) = w G_{4,o}(s) + (1-w) G_{4,e}(s), \quad w = \frac{15093}{17593}, \quad (\text{E.12})$$

while

$$G_{B_e}(s) = G_{4,e}(s). \quad (\text{E.13})$$

All four transforms therefore share the same quadratic denominator

$$D(s) = (\lambda_3 + s)(\lambda_4 + s) - r\lambda_3\lambda_4 = (\lambda_3 + s)(\lambda_4 + s) - 52.5. \quad (\text{E.14})$$

Proof. The odd and even transforms for entry through state 4 are summed geometric series in the product $L_3(s)L_4(s)$, which gives (E.10). Entry through state 3 either exits immediately through 3 or first moves to 4, which yields (E.11). The B -symbol transforms are then the conditional mixtures induced by the two possible hidden entry microstates, giving (E.12) and (E.13). The common denominator is immediate from (E.10).

Theorem E.2 shows that the hidden timing sector of the transformed observer is not an uncontrolled memory tower. It is exactly two-dimensional at the visible-law level.

Proposition E.3 (Exact finite completion). *The parity-lifted transformed observer admits an explicit exact finite CTMC completion. One such completion has visible states 1, 2 together with four hidden phase blocks: a two-state block for A_o , a two-state block for A_e , a four-state block for B_o , and a four-state block for B_e . The two basic hidden excursion generators are*

$$S^{(4)} = \begin{pmatrix} -50.7 & 50.7 \\ 175/169 & -83.2 \end{pmatrix}, \quad (\text{E.15})$$

$$S^{(3)} = \begin{pmatrix} -50.7 & 525/832 \\ 83.2 & -83.2 \end{pmatrix}, \quad (\text{E.16})$$

and the B -blocks are direct sums of these two generators with the conditional odd or even entry weights. Hence an exact finite completion exists with total state count 14.

Proof. The parity transforms in Theorem E.2 are rational phase-type laws of order at most two. The A -symbols require one conditioned two-state excursion block each. The B -symbols are conditional mixtures over the two possible hidden entry microstates, so each is realised exactly by the corresponding direct-sum block. Adding the visible states 1, 2 gives $2 + 2 + 2 + 4 + 4 = 14$ states.

Remark (One exact completion and the renewal frontier). On the explicit 14-state completion from Proposition E.3, the transformed observer has active visible tangent rank 4 and hidden-load rank 2 at the visible stalling point. Numerically,

$$\lambda(T_{\text{tr}}) \approx (0, 0.81855, 0.88615, 9.26522, 10.93096),$$

$$\lambda(X_{\text{tr}}) \approx (0, 0.69542, 0.85140, 8.94220, 10.77309),$$

$$\lambda(\Lambda_{\text{tr}}) \approx (0, 0, 0.07834, 0.19437).$$

By contrast, the original full-CG observer on the 4-state chain is only two-dimensional and is already a one-hard-direction quotient observer, with hidden-load spectrum approximately $(0, 0.274678)$ at stall. These transformed spectra are exact for the chosen completion, but the completion-independent part of the renewal story remains open. What is missing at present is a theorem that fixes a canonical tangent ceiling and canonical hidden-load operator directly from the renewal visible law rather than from a chosen finite completion. This is the natural extension frontier beyond the static quotient geometry proved in the main text.

F Blind benchmark diagnostics

The main text reports only the benchmark outputs needed to establish the operational point. The additional diagnostics here show the spread of ensemble finite errors, the adaptive-reveal comparison, and the exact-versus-sampled Bell contrast. They should be read as benchmark diagnostics, not as extensions of the theorem domain.

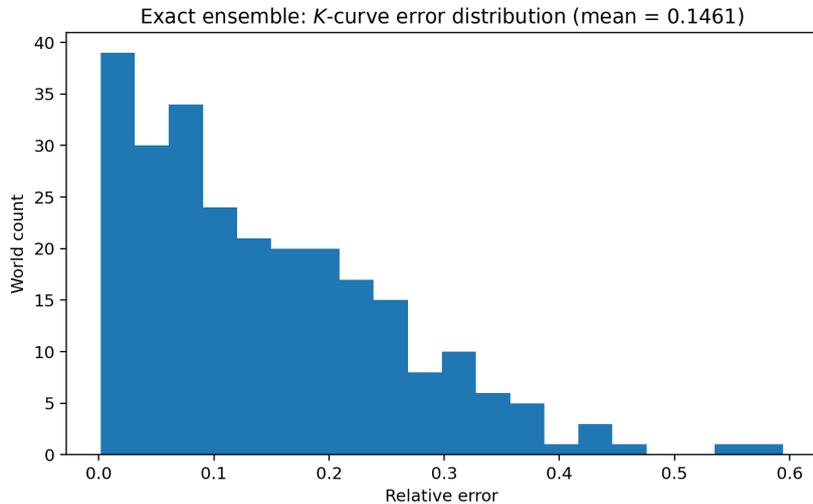


Figure 6: Exact 256-world ensemble. Distribution of relative errors for the finite retained-signal curve K_d .

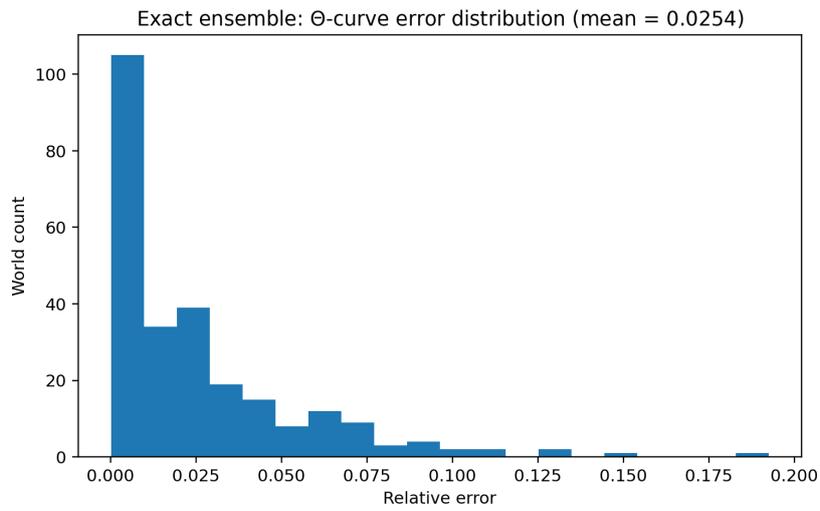


Figure 7: Exact 256-world ensemble. Distribution of relative errors for the detectability curve Θ_d .

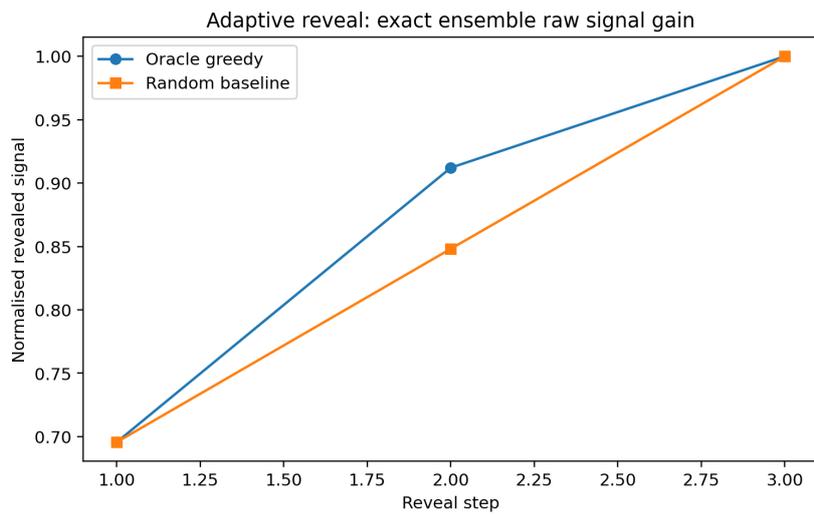


Figure 8: Adaptive reveal on the 64-world exact diagnostic panel. The oracle theorem-target reveal improves the mean second-step raw revealed signal over the random baseline.

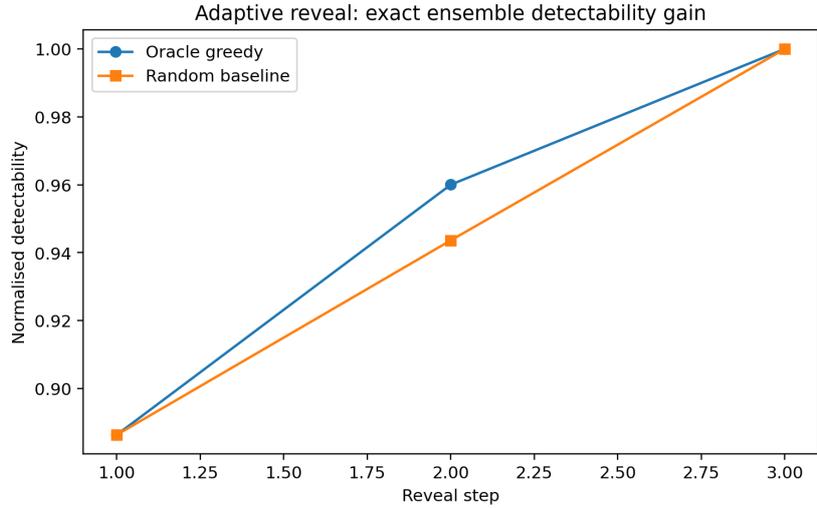


Figure 9: Adaptive reveal on the same 64-world exact diagnostic panel, now measured in the normalised detectability score. The gain is positive but smaller than in the raw revealed signal, which is why the present paper reports adaptive reveal only as a diagnostic frontier.

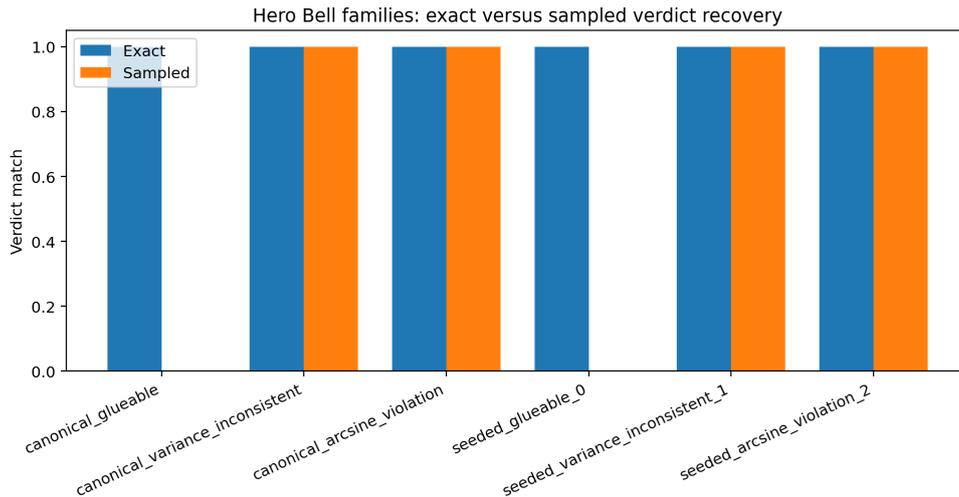


Figure 10: Hero Bell families in exact and sampled modes. Exact-mode verdict recovery is clean on the recorded hero world, whereas the sampled mode misses the glueable families and is therefore treated as a supporting robustness layer rather than as the primary Bell-facing claim.

G Methodological bridges and limits of the positive-survivor programme

This final appendix records three bridge theorems that sit outside the main theorem flow of the paper but follow the same hidden-to-visible methodology. They are included as a theorem bank, not as inputs to the proofs above. Their role is to show that the present quotient construction belongs to a broader reduction programme while keeping proof-status discipline explicit. What is *not* claimed here is a universal abstract theorem covering all hidden-to-visible reductions in one stroke. The statements below exhibit two species only: a second-order visible-medium species and a passive-transfer species. The present paper then sharpens that broader pattern in one precise way: here the primary invariant is a closed

visible law, whereas in the bridge theorems below the canonical visible object is a quotient or test-class survivor.

G.1 Bridge A: structural versus practical identifiability

Let Θ be a smooth parameter manifold, let \mathcal{E} be an admissible experiment class, and for each $E \in \mathcal{E}$ let

$$p_E(z | \theta)$$

be the observational law. Define structural equivalence relative to \mathcal{E} by

$$\theta \sim \theta' \iff p_E(\cdot | \theta) = p_E(\cdot | \theta') \text{ for all } E \in \mathcal{E}.$$

Assume that near a regular point the equivalence classes form a smooth quotient manifold

$$q : \Theta \rightarrow \bar{\Theta} := \Theta / \sim .$$

For a fixed experiment E , let $I_E(\theta)$ denote the Fisher information bilinear form on $T_\theta \Theta$.

Theorem G.1 . *Under the regular quotient hypothesis, each experiment $E \in \mathcal{E}$ factors uniquely through the quotient: there exists a reduced statistical model*

$$\bar{p}_E(z | \bar{\theta}), \quad \bar{\theta} \in \bar{\Theta},$$

such that

$$p_E(z | \theta) = \bar{p}_E(z | q(\theta)).$$

Moreover:

1. for every vertical vector $v \in \ker dq_\theta$, the score annihilates,

$$\partial_v \log p_E(z | \theta) = 0;$$

2. the Fisher form vanishes on vertical directions and descends uniquely to a bilinear form \bar{I}_E on $T_{q(\theta)} \bar{\Theta}$ satisfying

$$I_E(\theta)(u, w) = \bar{I}_E(q(\theta))(dq_\theta u, dq_\theta w);$$

3. structural nonidentifiability is exactly the vertical degeneracy of q , while practical nonidentifiability for a given experiment is degeneracy or strong anisotropy of the descended form \bar{I}_E on the quotient tangent.

Proof. By definition, each $p_E(\cdot | \theta)$ is constant on every structural fibre, so it depends only on $q(\theta)$; this gives the unique reduced model. Differentiating the factorisation along a vertical vector $v \in \ker dq_\theta$ gives the score annihilation. The Fisher identity then shows that $I_E(\theta)(v, \cdot) = 0$, so $I_E(\theta)$ depends only on quotient tangent classes and therefore descends uniquely to \bar{I}_E .

The theorem is local rather than global. It says that the correct canonical object is not a matrix written in some auxiliary horizontal complement, but the descended bilinear form on the structural quotient itself. A direct consequence is that enlarging the admissible experiment class can only refine the quotient and remove vertical directions; it cannot create a new structural fibre.

Proposition G.2 . *Let $\mathcal{E}_1 \subseteq \mathcal{E}_2$ be two experiment classes, and assume both regular quotients exist near the point of interest:*

$$q_i : \Theta \rightarrow \bar{\Theta}_i = \Theta / \sim_i, \quad i = 1, 2.$$

Then there exists a unique local surjection $\pi : \bar{\Theta}_2 \rightarrow \bar{\Theta}_1$ such that

$$q_1 = \pi \circ q_2.$$

For every experiment $E \in \mathcal{E}_1$, the descended Fisher forms satisfy the pullback relation

$$\bar{I}_E^{(2)} = \pi^* \bar{I}_E^{(1)}.$$

Hence enlarging the experiment class can only sharpen the quotient geometry.

Proof. If two parameters are equivalent for the larger class \mathcal{E}_2 , they are equivalent for the smaller class \mathcal{E}_1 . Therefore every \sim_2 -class sits inside a \sim_1 -class, which induces the unique local map π with $q_1 = \pi \circ q_2$. The Fisher pullback identity then follows from Theorem G.1 applied to the factorisation of the statistical model through the two quotients.

A canonical toy model is the scalar decay law

$$\dot{x} = -kx, \quad Y_i = cx_0 e^{-kt_i} + \varepsilon_i, \quad \varepsilon_i \sim \mathcal{N}(0, \sigma^2),$$

with parameters (c, x_0, k) . The full observational law depends only on $a = cx_0$ and k , so the structural fibres are

$$(c, x_0, k) \sim (\lambda c, x_0/\lambda, k), \quad \lambda > 0.$$

The quotient is therefore two-dimensional with coordinates (a, k) , and the quotient Fisher form is exactly the Fisher matrix of the reduced mean model ae^{-kt_i} . This shows in one line that a direction can be structurally invisible upstairs while the remaining quotient directions are still experiment-dependent in visibility.

In the present quotient setting this split becomes finite-dimensional rather than merely local.

Theorem G.3 (Exact inverse split for the hidden-load class). Let $T \succeq 0$ and let $S := \text{Ran}(T)$. Let X be a visible law in the support-preserving class beneath T , so that by Theorem 6.4

$$X = T^{1/2}(I_S + \Lambda)^{-1}T^{1/2}, \quad \Lambda \succeq 0 \text{ on } S.$$

Then:

- (i) conditional on the tangent ceiling T , the hidden-load operator is uniquely identified by X and equals

$$\Lambda = (T|_S)^{1/2}(X|_S)^{-1}(T|_S)^{1/2} - I_S; \quad (\text{G.1})$$

moreover,

$$\text{rank}(\Lambda) = \text{rank}(T - X), \quad \log \det(I_S + \Lambda) = \log \text{pdet}(T) - \log \text{pdet}(X);$$

- (ii) without an independently fixed ceiling, X alone does not identify a unique hidden load: for every $T' \succeq X$ with $\text{Ran}(T') = S$, there exists a unique $\Lambda_{T'} \succeq 0$ on S such that

$$X = T'^{1/2}(I_S + \Lambda_{T'})^{-1}T'^{1/2}, \quad \Lambda_{T'} = (T'|_S)^{1/2}(X|_S)^{-1}(T'|_S)^{1/2} - I_S.$$

Hence X determines a full ceiling cone of compatible ceiling-load explanations.

Proof. Part (i) is the explicit recovery formula already built into Theorem 6.4. The rank and determinant identities are Corollary 6.7.

For part (ii), fix $T' \succeq X$ on the same support S and define $\Lambda_{T'}$ by the displayed formula. Then

$$T'^{-1/2}XT'^{-1/2} \preceq I_S.$$

Since inversion reverses Loewner order on positive definite operators,

$$(T'^{-1/2}XT'^{-1/2})^{-1} \succeq I_S,$$

so $\Lambda_{T'} \succeq 0$. Rearranging gives the stated representation of X , and uniqueness follows because $T'^{-1/2}XT'^{-1/2} = (I_S + \Lambda_{T'})^{-1}$.

Combined with Corollary 6.5, this gives the order reversal at fixed ceiling:

$$X_1 \preceq X_0 \iff \Lambda_1 \succeq \Lambda_0.$$

It also gives the zero-load boundary $\Lambda = 0 \iff X = T$.

Proposition G.4 (Gram rigidity and hidden-factor gauge). Fix T and X and let Λ be the uniquely identified hidden-load operator from Theorem G.3. Write $r := \text{rank}(\Lambda)$. Then every minimal hidden realisation is a Gram factorisation

$$\Lambda = BB^\top, \quad B \in \mathbb{R}^{|S| \times r},$$

and if $B, B' \in \mathbb{R}^{|S| \times r}$ both have full column rank and satisfy $BB^\top = B'B'^\top = \Lambda$, then

$$B' = BQ$$

for some orthogonal matrix $Q \in O(r)$. Consequently (T, X) identifies the hidden Gram operator and its intrinsic invariants, but not a unique hidden basis, hidden factorisation, or hidden sparsity pattern.

Proof. The existence of a minimal factorisation follows from Theorem 6.6. If $BB^\top = B'B'^\top$ with both factors of column rank r , then the linear map sending the columns of B to those of B' preserves the Euclidean inner product induced by the common Gram matrix. Hence it extends to an orthogonal matrix $Q \in O(r)$ with $B' = BQ$.

G.2 Bridge B: CME versus CLE as a visible quadratic shadow theorem

Consider a reaction network on \mathbb{N}^d with reaction jumps $\nu_r \in \mathbb{Z}^d$, propensities $a_r(x) \geq 0$, drift

$$b(x) = \sum_{r=1}^R a_r(x) \nu_r,$$

and quadratic fluctuation matrix

$$D(x) = \sum_{r=1}^R a_r(x) \nu_r \nu_r^\top.$$

Let $\Pi \in \mathbb{R}^{m \times d}$ be a linear observation map and write

$$y = \Pi x, \quad b_\Pi(x) = \Pi b(x), \quad D_\Pi(x) = \Pi D(x) \Pi^\top.$$

Define the pullback algebra

$$\mathcal{F}_\Pi = \{f : \exists \varphi \text{ with } f = \varphi \circ \Pi\}.$$

Theorem G.5 . Let L_{CME} be the jump generator

$$(L_{\text{CME}}f)(x) = \sum_{r=1}^R a_r(x) [f(x + v_r) - f(x)],$$

and let L_{CLE} be the diffusion generator

$$(L_{\text{CLE}}g)(x) = b(x) \cdot \nabla g(x) + \frac{1}{2} D(x) : \nabla^2 g(x).$$

Then the following hold.

1. If φ is any polynomial on observation space of degree at most two, then on the pullback observable $f = \varphi \circ \Pi$ one has the exact identity

$$L_{\text{CME}}f(x) = b_{\Pi}(x) \cdot \nabla \varphi(\Pi x) + \frac{1}{2} D_{\Pi}(x) : \nabla^2 \varphi(\Pi x).$$

In particular, CME and CLE have the same visible drift and the same visible quadratic shadow on that test class.

2. The projected quadratic theory closes on observation space if and only if both b_{Π} and D_{Π} are fibre-constant on the fibres of Π .
3. For general C^3 pullback observables, the first exact discrepancy between the two generators is the third-order Taylor remainder. Thus the quadratic visible shadow is exact, and the first forgotten local structure begins at cubic order, with boundary effects contributing a separate residue on the lattice state space.

Proof. For $f = \varphi \circ \Pi$ with $\deg \varphi \leq 2$, expand

$$\varphi(\Pi x + \Pi v_r) - \varphi(\Pi x) = \nabla \varphi(\Pi x) \cdot (\Pi v_r) + \frac{1}{2} (\Pi v_r)^{\top} \nabla^2 \varphi(\Pi x) (\Pi v_r),$$

which is exact because all third derivatives vanish. Summing over r gives the stated formula. Closure is equivalent to the right-hand side depending only on $y = \Pi x$, which is exactly fibre-constancy of b_{Π} and D_{Π} . For general C^3 pullbacks, Taylor's theorem with integral remainder shows that the omitted term begins at order three.

A simple exact example is

$$\emptyset \xrightarrow{\alpha} X, \quad X \xrightarrow{\mu} \emptyset, \quad X \xrightarrow{\kappa} Y, \quad Y \xrightarrow{\mu} \emptyset,$$

observed through $\Pi = (1 \ 1)$ so that $n = x + y$. Then the internal conversion $X \rightarrow Y$ is invisible to the observation and one finds

$$b_{\Pi}(x, y) = \alpha - \mu(x + y) = \alpha - \mu n, \quad D_{\Pi}(x, y) = \alpha + \mu(x + y) = \alpha + \mu n.$$

Hence both visible coefficients are fibre-constant in n , so the quadratic visible theory closes exactly even though the underlying network is not itself visible.

G.3 Bridge C: hidden elimination in dissipative input-output systems

Consider a linear co-located dissipative system

$$\dot{z} = Az + Gu, \quad y = G^{\top} z,$$

with state $z \in \mathbb{R}^n$, input $u \in \mathbb{R}^m$, and dissipation inequality

$$A + A^\top \leq 0.$$

Write the corresponding transfer function on the open right half-plane as

$$H(s) = G^\top (sI - A)^{-1} G, \quad \Re s > 0.$$

Split the state into resolved and hidden blocks,

$$z = (r, h), \quad A = \begin{pmatrix} A_{rr} & A_{rh} \\ A_{hr} & A_{hh} \end{pmatrix}, \quad G = \begin{pmatrix} G_r \\ G_h \end{pmatrix}.$$

Theorem G.6. *Assume zero initial state and $\Re s > 0$.*

1. *The full transfer function is positive real:*

$$H(s) + H(s)^* \geq 0.$$

2. *Exact elimination of the hidden block gives the reduced dynamic Schur-complement representation*

$$H(s) = D_{\text{eff}}(s) + C_{\text{eff}}(s) \Sigma_r(s)^{-1} B_{\text{eff}}(s),$$

where

$$\Sigma_r(s) = sI - A_{rr} - A_{rh}(sI - A_{hh})^{-1}A_{hr},$$

$$B_{\text{eff}}(s) = G_r + A_{rh}(sI - A_{hh})^{-1}G_h,$$

$$C_{\text{eff}}(s) = G_r^\top + G_h^\top (sI - A_{hh})^{-1}A_{hr},$$

$$D_{\text{eff}}(s) = G_h^\top (sI - A_{hh})^{-1}G_h.$$

In particular, the reduced transfer remains positive real.

3. *In the gradient special case $A = -B$ with $B = B^\top \geq 0$ and hidden variables not directly forced or observed, the hidden contribution becomes a symmetric positive kernel in time, and the reduced transfer is a matrix Stieltjes function. Thus the passive-transfer bridge contains a symmetric kernel subcase and a genuinely dynamic positive-real case.*

Proof. For any vector ξ , set $X = (sI - A)^{-1}G\xi$. Then

$$G\xi = (sI - A)X.$$

Hence

$$2 \operatorname{Re}(\xi^* H(s) \xi) = 2 \operatorname{Re}(X^* G \xi) = 2 \operatorname{Re}(s) \|X\|^2 - X^*(A + A^\top)X \geq 0,$$

which proves positive-realness. The reduced formula follows by block elimination of the hidden variable in the resolvent equation

$$(sI - A) \begin{pmatrix} r \\ h \end{pmatrix} = \begin{pmatrix} G_r \\ G_h \end{pmatrix} u.$$

Substituting the eliminated hidden block into the resolved equation and then into the output formula yields the stated Schur-complement representation. In the gradient case, symmetry of B forces the hidden contribution to be a symmetric positive kernel after inverse Laplace

transform, and the transfer therefore lies in the Stieltjes subclass.

A two-state example already shows the mechanism. Let

$$A = \begin{pmatrix} -a & c \\ c & -d \end{pmatrix}, \quad G = \begin{pmatrix} 1 \\ g \end{pmatrix}, \quad a, d > 0, \quad |c| \leq \sqrt{ad}.$$

Then $A + A^\top \leq 0$, the transfer $H(s) = G^\top (sI - A)^{-1} G$ is positive real on $\Re s > 0$, and eliminating the second state gives the dynamic Schur-complement formula of Theorem G.6. When $g = 0$ the hidden state is only internally coupled and the reduced transfer becomes the resolved-port special case.

G.4 Synthesis and boundary

The three bridge theorems above exhibit two species of hidden-to-visible reduction.

Proposition G.7 . *The identifiability and CME/CLE bridges are second-order visible-medium reductions: after removing structurally invisible directions or restricting to a visible pullback class, the canonical visible object is a positive bilinear or quadratic survivor on the active quotient. The dissipative MZ/ROM bridge is a passive-transfer reduction: after hidden elimination, the canonical visible object is a positive-real reduced transfer, with a symmetric kernel subcase in gradient systems. The present paper is stronger in a different direction: it proves a closed visible law on the positive cone, with the tangent Donsker–Varadhan object playing the role of a lower-rung survivor beneath that law.*

Proof. The first two assertions are Theorem G.1 and Theorem G.5. The passive-transfer assertion is Theorem G.6. The final statement follows from the main body of the paper: the quotient precision itself is closed under marginal elimination, whereas the tangent Donsker–Varadhan sector reappears only as the first quotient-visible law in the H_0 -adapted tangent regime.

We close by stating explicitly what this appendix does *not* prove. It does not prove a universal abstract theorem unifying all hidden-to-visible reductions. It does not identify one single invariant class containing both the second-order visible-medium species and the passive-transfer species. It does not claim that every reduction problem in neighbouring fields must fall into one of the three bridge forms above. What it does provide is a theorem-safe bridge bank showing that the present quotient programme sits inside a wider and already nontrivial hidden-to-visible methodology.

Acknowledgements

No external funding was received. All work was produced independently without affiliation. Author’s email is contact@nomogenetics.com.

References

- [1] J. R. Dunkley. *Finite Observation: A Spectral Theory of Nonequilibrium Visibility*. Preprint, 2026.
- [2] 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**:1–47, 1975.
- [3] 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**:279–301, 1975.
- [4] M. D. Donsker and S. R. S. Varadhan. Asymptotic evaluation of certain Markov process expectations for large time. III. *Communications on Pure and Applied Mathematics*, **29**(4):389–461, 1976.
- [5] R. S. Ellis. *Entropy, Large Deviations, and Statistical Mechanics*. Grundlehren der mathematischen Wissenschaften, Vol. 271, Springer, 1985. Reprinted as Classics in Mathematics, 2006.
- [6] H. Touchette. The large deviation approach to statistical mechanics. *Physics Reports*, **478**:1–69, 2009.
- [7] L. Bertini, A. Faggionato, and D. Gabrielli. Large deviations of the empirical flow for continuous time Markov chains. *Annales de l’I.H.P. Probabilités et Statistiques*, **51**(3):867–900, 2015.
- [8] J. Schnakenberg. Network theory of microscopic and macroscopic behavior of master equation systems. *Reviews of Modern Physics*, **48**:571–585, 1976.
- [9] J. L. Lebowitz and H. Spohn. A Gallavotti–Cohen-type symmetry in the large deviation functional for stochastic dynamics. *Journal of Statistical Physics*, **95**:333–365, 1999.
- [10] U. Seifert. Stochastic thermodynamics, fluctuation theorems and molecular machines. *Reports on Progress in Physics*, **75**:126001, 2012.
- [11] M. Esposito. Stochastic thermodynamics under coarse graining. *Physical Review E*, **85**:041125, 2012.
- [12] C. Maes. Frenesy: time-symmetric dynamical activity in nonequilibria. *Physics Reports*, **850**:1–33, 2020.
- [13] I. Schur. Über Potenzreihen, die im Innern des Einheitskreises beschränkt sind. *Journal für die reine und angewandte Mathematik*, **147**:205–232, 1917.
- [14] F. Zhang, editor. *The Schur Complement and Its Applications*. Springer, New York, 2005.
- [15] E. V. Haynsworth. On the Schur complement. *Basel Mathematical Notes*, BMN 20, 1968.
- [16] W. N. Anderson, Jr. Shorted operators. *SIAM Journal on Applied Mathematics*, **20**(3):520–525, 1971.
- [17] W. N. Anderson, Jr. and G. E. Trapp. Shorted operators. II. *SIAM Journal on Applied Mathematics*, **28**(1):60–71, 1975.
- [18] R. Bhatia. *Positive Definite Matrices*. Princeton University Press, 2007.
- [19] S. Demko, W. F. Moss, and P. W. Smith. Decay rates for inverses of band matrices. *Mathematics of Computation*, **43**(168):491–499, 1984.
- [20] P. G. Doyle and J. L. Snell. *Random Walks and Electric Networks*. Carus Mathematical Monographs, Vol. 22, Mathematical Association of America, 1984.

- [21] D. M. Malioutov, J. K. Johnson, and A. S. Willsky. Walk-sums and belief propagation in Gaussian graphical models. *Journal of Machine Learning Research*, **7**:2031–2064, 2006.
- [22] V. Chandrasekaran, P. A. Parrilo, and A. S. Willsky. Latent variable graphical model selection via convex optimization. *The Annals of Statistics*, **40**(4):1935–1967, 2012.
- [23] W. N. Anderson, Jr. and R. J. Duffin. Series and parallel addition of matrices. *Journal of Mathematical Analysis and Applications*, **26**(3):576–594, 1969.
- [24] K. Löwner. Über monotone Matrixfunktionen. *Mathematische Zeitschrift*, **38**:177–216, 1934.
- [25] T. Ando. Concavity of certain maps on positive definite matrices and applications to Hadamard products. *Linear Algebra and Its Applications*, **26**:203–241, 1979.
- [26] F. Hansen and G. K. Pedersen. Jensen’s inequality for operators and Löwner’s theorem. *Mathematische Annalen*, **258**(3):229–241, 1982.
- [27] R. Bhatia. *Matrix Analysis*. Graduate Texts in Mathematics, Vol. 169, Springer, 1997.
- [28] F. Kubo and T. Ando. Means of positive linear operators. *Mathematische Annalen*, **246**:205–224, 1980.
- [29] W. Pusz and S. L. Woronowicz. Functional calculus for sesquilinear forms and the purification map. *Reports on Mathematical Physics*, **8**(2):159–170, 1975.
- [30] L. P. Kadanoff. Scaling laws for Ising models near T_c . *Physics Physique Fizika*, **2**:263–272, 1966.
- [31] K. G. Wilson and J. Kogut. The renormalization group and the ε expansion. *Physics Reports*, **12**(2):75–200, 1974.
- [32] J. Polchinski. Renormalization and effective Lagrangians. *Nuclear Physics B*, **231**(2):269–295, 1984.
- [33] D. Carlson, E. Haynsworth, and T. Markham. A generalization of the Schur complement by means of the Moore–Penrose inverse. *SIAM Journal on Applied Mathematics*, **26**(1):169–175, 1974.
- [34] K. Fan. Maximum properties and inequalities for the eigenvalues of completely continuous operators. *Proceedings of the National Academy of Sciences*, **37**:760–766, 1951.
- [35] R. A. Horn and C. R. Johnson. *Matrix Analysis*. Cambridge University Press, second edition, 2013.
- [36] A. P. Dempster. Covariance selection. *Biometrics*, **28**:157–175, 1972.
- [37] S. L. Lauritzen. *Graphical Models*. Oxford Statistical Science Series, No. 17, Oxford University Press, 1996.
- [38] R. Grone, C. R. Johnson, E. M. Sá, and H. Wolkowicz. Positive definite completions of partial Hermitian matrices. *Linear Algebra and Its Applications*, **58**:109–124, 1984.
- [39] E. B. Bogomol’nyi. Stability of classical solutions. *Soviet Journal of Nuclear Physics*, **24**(4):449–454, 1976.
- [40] J. Williamson. On the algebraic problem concerning the normal forms of linear dynamical systems. *American Journal of Mathematics*, **58**(1):141–163, 1936.
- [41] R. A. Horn and C. R. Johnson. *Topics in Matrix Analysis*. Cambridge University Press, 1991.
- [42] A. George and J. W. H. Liu. *Computer Solution of Large Sparse Positive Definite Systems*. Prentice-Hall, Englewood Cliffs, NJ, 1981.

- [43] M. Benzi and G. H. Golub. Bounds for the entries of matrix functions with applications to preconditioning. *BIT Numerical Mathematics*, **39**(3):417–438, 1999.
- [44] M. Benzi and N. Razouk. Decay bounds and $O(n)$ algorithms for approximating functions of sparse matrices. *Electronic Transactions on Numerical Analysis*, **28**:16–39, 2007.
- [45] T. J. Rivlin. *Chebyshev Polynomials: From Approximation Theory to Algebra and Number Theory*. Wiley, 1974; second edition, 1990.
- [46] L. N. Trefethen. *Approximation Theory and Approximation Practice*. SIAM, Philadelphia, 2013.
- [47] M. Rudelson and R. Vershynin. Sampling from large matrices: an approach through geometric functional analysis. *Journal of the ACM*, **54**(4), Article 21, 2007.
- [48] J. G. Kemeny and J. L. Snell. *Finite Markov Chains*. Springer-Verlag, 1976. Originally published by D. Van Nostrand, 1960.
- [49] P. Buchholz. Exact and ordinary lumpability in finite Markov chains. *Journal of Applied Probability*, **31**:59–75, 1994.
- [50] G. Teza and A. L. Stella. Exact coarse graining preserves entropy production out of equilibrium. *Physical Review Letters*, **125**:110601, 2020.
- [51] Y. Weiss and W. T. Freeman. Correctness of belief propagation in Gaussian graphical models of arbitrary topology. *Neural Computation*, **13**(10):2173–2200, 2001.
- [52] C. R. Rao. Information and accuracy attainable in the estimation of statistical parameters. *Bulletin of the Calcutta Mathematical Society*, **37**:81–91, 1945.
- [53] S. Amari and H. Nagaoka. *Methods of Information Geometry*. Translations of Mathematical Monographs, Vol. 191, American Mathematical Society, 2000.
- [54] N. Ay, J. Jost, H. V. Lê, and L. Schwachhöfer. *Information Geometry*. Ergebnisse der Mathematik und ihrer Grenzgebiete, Vol. 64, Springer, 2017.
- [55] C. Lüders, J. Thewes, and M. Aßmann. Real time $g^{(2)}$ monitoring with 100 kHz sampling rate. *Optics Express*, **26**(19):24854–24863, 2018. doi: 10.1364/OE.26.024854.