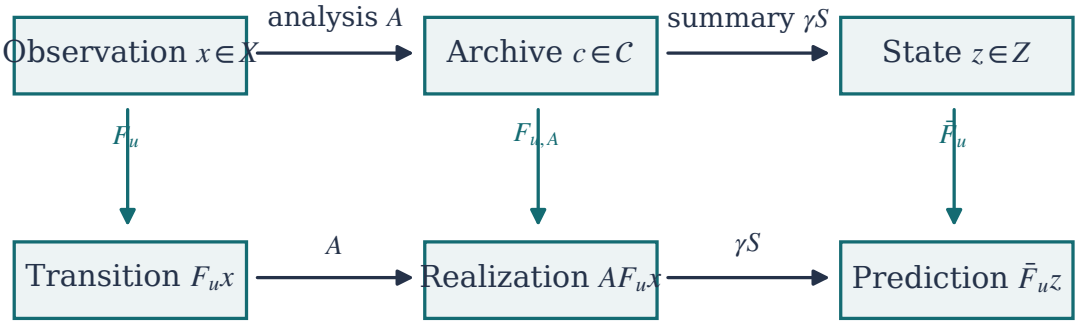
# Large Signal Model

## A theoretical atlas

Realization, predictive closure, memory
and certified probability domains




Left square: coordinate consistency. Right square: predictive closure.

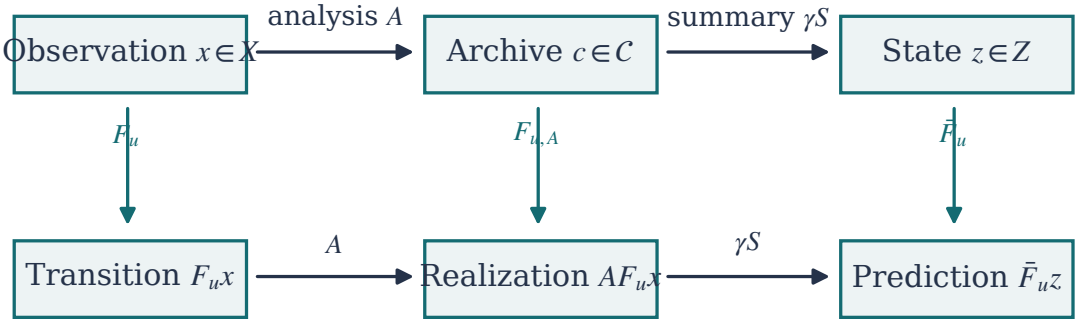
Komorebi P.

A visual companion to the mathematical volume. Exact identities, geometric obstructions and worked rational examples connect the theory to its formal proofs.

Orientation

# Four contracts, one model

An observation, its archive, a predictive summary and a probability state carry different mathematical obligations. The theory connects them through maps whose domains and identities are explicit.




Left square: coordinate consistency. Right square: predictive closure.

| Contract | Defining identity | Question answered |
|---|---|---|
| Reconstruction | $SA = I_X$ | Can the accepted observation be recovered? |
| Transport | $F_b T_{ba} = T_{ba} F_a$ | Do coordinate descriptions agree? |
| Closure | $\gamma F_u = \bar{F}_u \gamma$ | Does the summary determine its own future? |
| Probability | $B_o C \subseteq C$, $\lambda \sum_o B_o = \lambda$ | Are every branch and every finite record valid? |

**Reading route.** Plates 1–4 build the representation and its predictive quotient. Plates 5–9 explain interactions, lost paths, memory and observable identification. Plates 10–16 develop behavior spaces, positivity, curved geometry and statistical guarantees. Plate 17 locates the formal proofs.

**Visual conventions.** Teal marks exact structure. Gold marks transported, refined or uncertain quantities. Red marks an obstruction or an omitted contribution. Every figure is drawn from a formula or an exact finite example. Curves of bounds are labeled as bounds.

**Scope.** The atlas selects core mechanisms from the 87 registered propositions. It introduces no new numbered theorem, fitted model, benchmark result or empirical coverage claim. Its diagrams supplement the proofs.

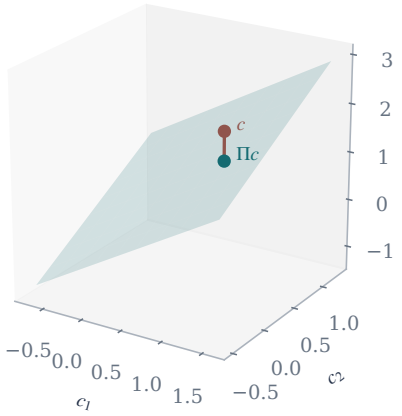
1 P1

# Realizable coefficients form a subspace

Exact storage includes a consistency condition.

For linear analysis and synthesis maps $A : X \to C$ and $S : C \to X$, the left-inverse law $SA = I_X$ gives injectivity of $A$. The coefficient projector is $\Pi = AS$, with

$$\Pi^2 = \Pi, \qquad \operatorname{im}\Pi = \operatorname{im} A, \qquad c \text{ is realizable} \iff \Pi c = c.$$




**A redundant archive.** Take $X = \mathbb{R}^2$, $C = \mathbb{R}^3$ and

$$A(x_1, x_2) = (x_1, x_2, x_1 + x_2), \qquad S(c_1, c_2, c_3) = (c_1, c_2).$$

Then $\Pi(c_1, c_2, c_3) = (c_1, c_2, c_1 + c_2)$. The shaded plane is the set of realizable coefficients. The red segment shows synthesis followed by reanalysis of an inconsistent vector.

**The mathematical point.** A left inverse determines an idempotent consistency map. Here it is oblique: $\Pi$ is not self-adjoint. Orthogonality requires an additional hypothesis.

**Scope.** For bounded linear maps on normed spaces, $\|S\|$ controls amplification of coefficient errors. Exact invertibility by itself gives no conditioning guarantee.

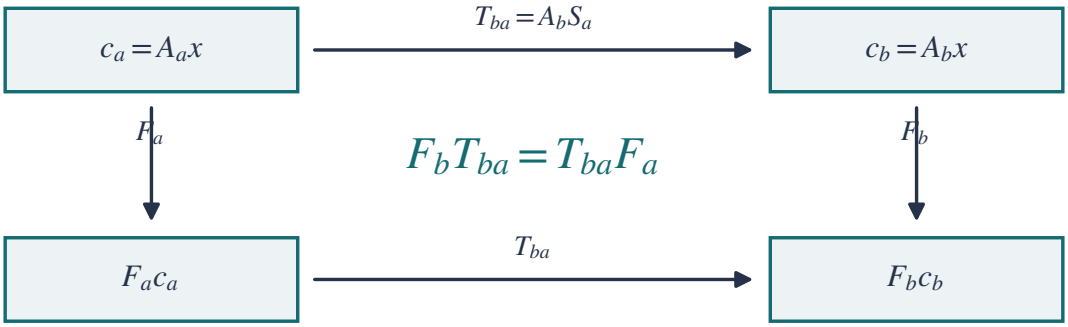
**Formal route.** `Realization.lean`; norm estimates in `Normed.lean`.

2 P2

# One operation, several coordinates

A commuting square expresses coordinate independence.

Let $\mathcal{C}_a = \operatorname{im} A_a$ and $T_{ba} = A_b S_a$. Coordinate realizations are $F_a = A_a F S_a$ on $\mathcal{C}_a$. Both paths through the square represent the same underlying operation.




$$T_{cb}T_{ba} = T_{ca}, \qquad T_{aa} = I_{\mathcal{C}_a}, \qquad F_b T_{ba} = T_{ba} F_a.$$

Conversely, a compatible family $G_a$ on the realizable coefficient spaces determines one operation. With a supplied reference chart $r$,

$$F = S_r G_r A_r, \qquad G_a = A_a F S_a.$$

**The mathematical point.** Compatibility is a descent condition: local coordinate descriptions belong to one object-space operation precisely when they agree under transport.

A fitted family's discrepancy can be measured in object space:

$$S_b(G_b A_b x - T_{ba} G_a A_a x) = S_b G_b A_b x - S_a G_a A_a x.$$

**Scope.** Zero chart discrepancy establishes agreement. It does not establish correctness of the agreed predictor. Deterministic recodings of the same observation supply the same observed information.

**Formal route.** `Realization.lean`. The formal construction uses a nonempty chart family.

3 P6–P7

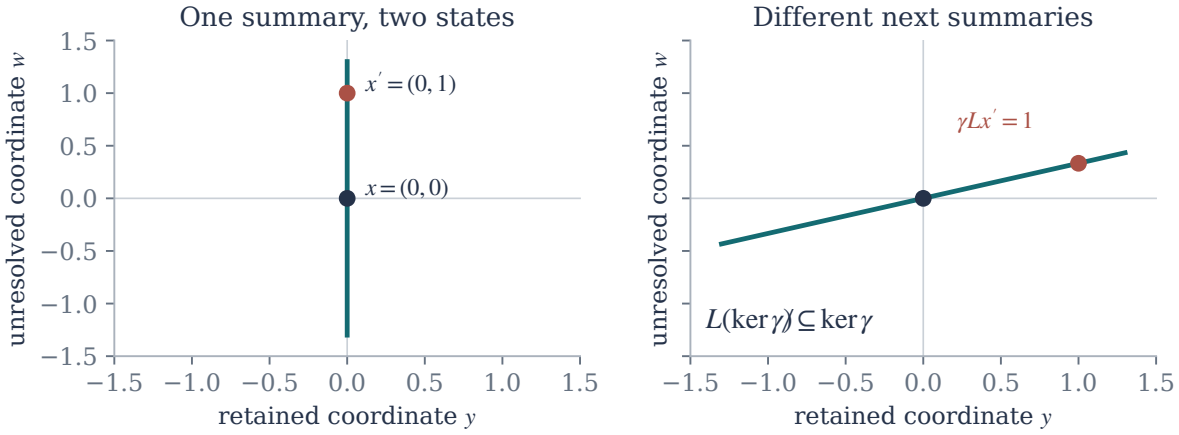
# The obstruction lives in a fiber

Equal present summaries must have equal next summaries.

A reduced map $\bar{F} : \gamma(X) \to \gamma(X)$ exists exactly when

$$\gamma x = \gamma x' \quad \Longrightarrow \quad \gamma F x = \gamma F x'.$$

For linear $L$ and linear $\gamma$, this is $L(\ker \gamma) \subseteq \ker \gamma$.




**The running rational example.** Throughout the deterministic plates, use

$$L = \begin{pmatrix} 1/2 & 1 \\ 1/4 & 1/3 \end{pmatrix}, \qquad \gamma = (1, 0), \qquad x_0 = (1, 0)^{\mathsf{T}}.$$

The two states $(0, 0)$ and $(0, 1)$ have summary zero. Their next summaries are $0$ and $1$. Thus no scalar autonomous update reproduces every state's first-coordinate dynamics.

**The mathematical point.** Predictive sufficiency is a property of entire summary fibers. A lossless archive of each full state does not make its chosen scalar summary sufficient.

**Scope.** The diagram is a counterexample for the declared operator. In a stochastic model, the corresponding condition concerns equality of next-summary distributions.

**Formal route.** `PredictiveClosure.lean`; `LinearClosure.lean`.

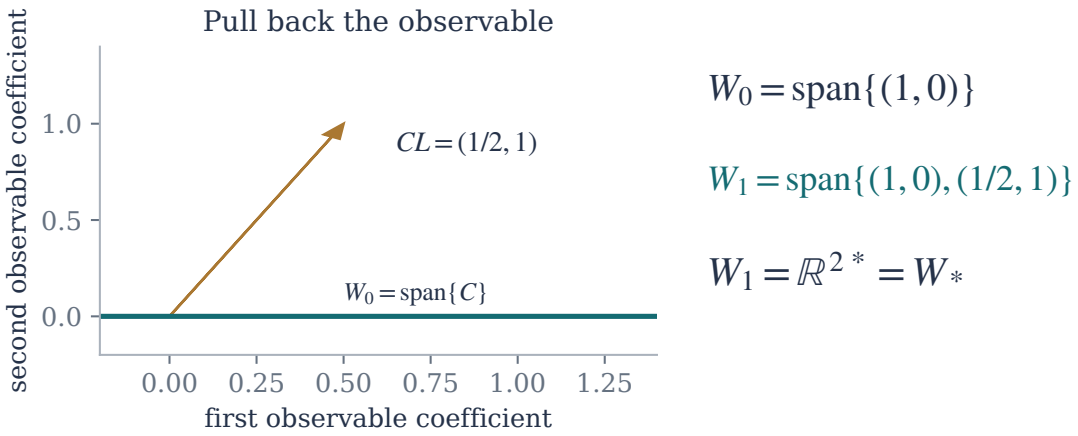
4 P8, P14

# Refinement closes an observable space

Add the directions generated by known dynamics.

For a known finite family $L_u$ on a $d$-dimensional space, start from the rows of a required observation $C_0$ and form

$$W_{n+1} = W_n + \sum_u W_n L_u, \qquad W_0 = \mathrm{row}(C_0).$$




$$W_0 = \mathrm{span}\{(1,0)\}$$

$$W_1 = \mathrm{span}\{(1,0),(1/2,1)\}$$

$$W_1 = \mathbb{R}^{2\,*} = W_*$$

Every strict inclusion raises dimension. After at most $d - \dim W_0$ strict enlargements, the limit $W_*$ is the least common invariant row space containing $W_0$.

For the running example, the rows of

$$O = \begin{pmatrix} 1 & 0 \\ 1/2 & 1 \end{pmatrix}, \qquad D = OLO^{-1} = \begin{pmatrix} 0 & 1 \\ 1/12 & 5/6 \end{pmatrix}$$

span $W_* = \mathbb{R}^{2*}$. The state $Ox_s = (y_s, y_{s+1})^{\mathsf{T}}$ therefore closes.

**The mathematical point.** The added coordinate is forced by the action of $L$ on the requested observable. In a larger system, unobserved directions unrelated to its future can remain outside $W_*$.

**Scope.** Computing $Ox_s$ uses a forward observation. It becomes an available delay state $(y_{t-1}, y_t)$ only after a one-step warm-up. Known-operator completion does not establish sensor access to arbitrary new observables.

**Formal route.** `InvariantCompletion.lean`; `PredictiveCompletion.lean`; `DelayClosure.lean`.

5 P3–P5, P20

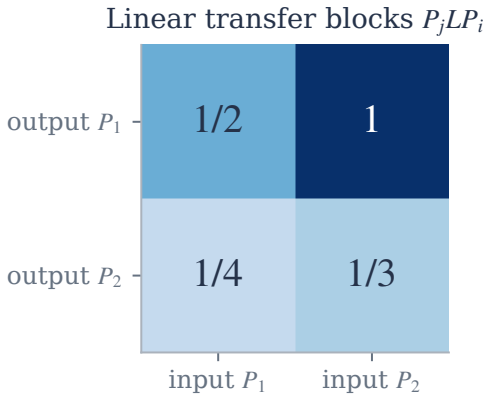
# Packets decompose operations too

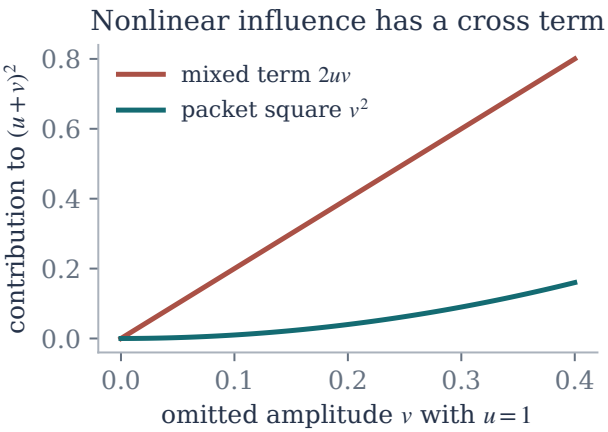
Resolution changes preserve identities when every block is retained.

For complete input and output projector families,

$$x = \sum_i P_i x, \qquad L = \sum_{j,i} Q_j L P_i.$$

Refining either family subdivides these transfer blocks without changing their total. Orthogonal packet families also give the corresponding squared-norm decomposition.




For a polynomial operation, multilinearity produces every ordered packet interaction:

$$Q_j B_m(x, \dots, x) = \sum_{i_1, \dots, i_m} Q_j B_m(P_{i_1} x, \dots, P_{i_m} x).$$

In the scalar square, the mixed contribution is $2uv$. With $u = 1$, it is linear in omitted amplitude $v$, whereas the omitted packet square is quadratic.

**The mathematical point.** A nonlinear operation acts on interactions between packets. Reconstruction energy alone does not bound an omitted packet's influence on a future quantity.

**Scope.** Complete refinement and truncation are different operations. Truncation requires an explicit omitted-interaction estimate on its declared domain. Projectors from different atlases need not commute.

**Formal route.** `PacketTree.lean`; `TreeEnergy.lean`; `Transfer.lean`; `Polynomial.lean`.

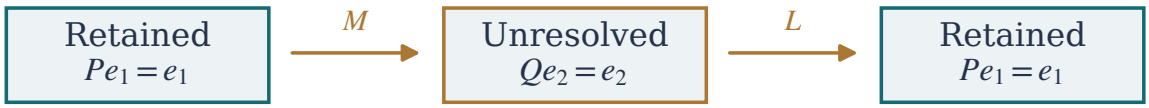
6 P7

# Compression can erase a return path

A two-step effect can survive when both compressed steps vanish.

Let $P^2 = P$ and $Q = I - P$. Inserting $I = P + Q$ gives the exact defect

$$\boxed{PLMP - (PLP)(PMP) = PLQMP.}$$




Separate compression deletes both arrows. Composition retains their return path.

Take $P = \mathrm{diag}(1, 0)$ and

$$L = M = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}.$$

Then $PLP = PMP = 0$, while $PLMP = P$. The missing term $PLQMP = P$ is the full return through the discarded coordinate.

The running example exhibits the same mechanism in $PL^2P$:

$$PL^2P - (PLP)^2 = \begin{pmatrix} 1/4 & 0 \\ 0 & 0 \end{pmatrix}.$$

Starting from $(1, 0)$, separate scalar compression predicts $y_2 = 1/4$. The exact value is $1/2$.

**The mathematical point.** The obstruction is an omitted route through the operator composition. Its size depends jointly on outgoing and returning transfer blocks.

**Scope.** The algebra requires idempotence. An orthogonal split additionally gives the familiar Hilbert-space energy interpretation.

**Formal route.** `LinearClosure.lean`; norm bounds in `Normed.lean`.

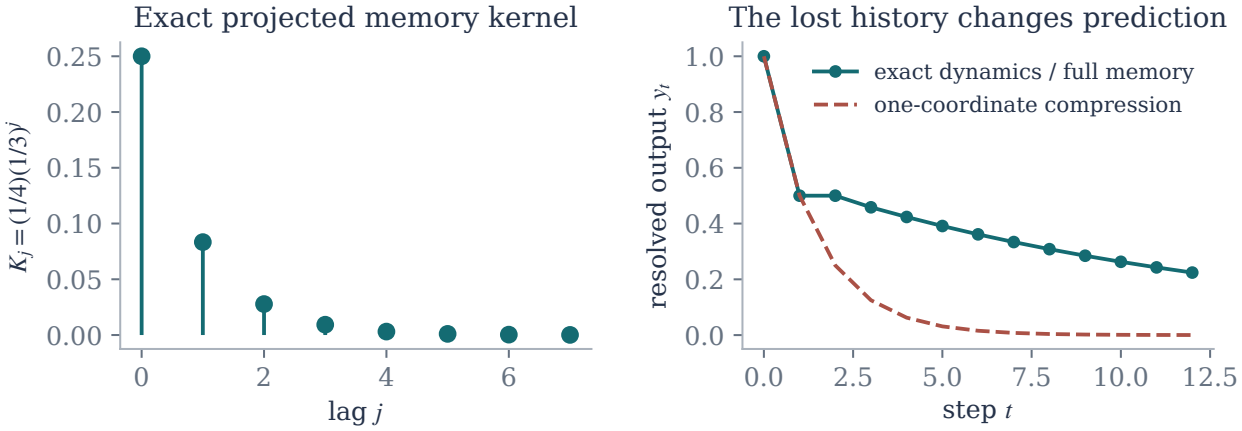
7 P10, P19

# Memory is an exact elimination

The discarded state leaves a kernel and an initial transient.

Write the split dynamics as $y_{t+1} = Ay_t + Bw_t$ and $w_{t+1} = Cy_t + Dw_t$. Eliminating $w_t$ gives

$$y_{t+1} = Ay_t + \sum_{j=0}^{t-1} BD^j C\, y_{t-1-j} + BD^t w_0.$$




For the running example, $A = 1/2$, $B = 1$, $C = 1/4$, $D = 1/3$ and $w_0 = 0$. Thus $K_j = \frac{1}{4}\,3^{-j}$. If $|y_t| \le 1$, retaining the last $H$ history terms leaves a same-history tail at most $\Delta_H = \frac{3}{8}\,3^{-H}$.

Here $0 \le y_t \le 1$ and the truncated recursion has gain

$$a_H = \frac{1}{2} + \sum_{j=0}^{H-1} K_j = \frac{7}{8} - \frac{3}{8}\,3^{-H} < 1.$$

With the same initial state, recursive propagation therefore gives

$$\sup_t |y_t - \hat{y}_t| \le \frac{\Delta_H}{1 - a_H} = \frac{3}{3^H + 3}.$$

**Scope.** P10's tail estimate uses the same resolved history. P19 supplies the separate stability argument for a predictor using its own history. For other initial states, retain the term $BD^t w_0$.

**Formal route.** `Memory.lean`; `MemoryBounds.lean`; `MemoryStability.lean`.

8 P18

# A local defect needs a propagation law

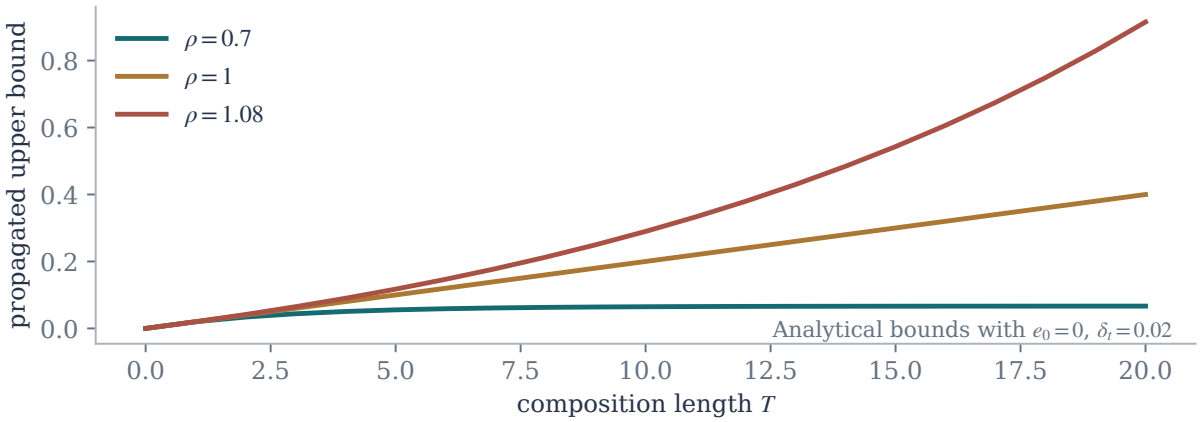
Small one-step residuals can accumulate or contract.

Allow a representation $\gamma_t$ to change with time. Define the typed defect

$$r_t(x) = \gamma_{t+1} F_t(x) - \hat{f}_t \gamma_t(x).$$

If $\hat{f}_t$ is $\lambda_t$-Lipschitz on the relevant domain and $\|r_t(x_t)\| \le \delta_t$, the prediction error satisfies

$$e_T \le \left( \prod_{j=0}^{T-1} \lambda_j \right) e_0 + \sum_{k=0}^{T-1} \left( \prod_{j=k+1}^{T-1} \lambda_j \right) \delta_k.$$




For a fixed representation $R$, a reference reduced operator $B$ and fitted operator $\widehat{B}$,

$$RL - \widehat{B}R = (RL - BR) + (B - \widehat{B})R.$$

The first term is representational nonclosure. The second is estimation error. They can enter the same propagation inequality after being bounded in the state norm.

**The mathematical point.** The suffix products determine whether earlier defects are damped or amplified. The chart shows analytical envelopes, not measured prediction errors.

**Scope.** Both the true represented state and predicted state must remain in the Lipschitz domain. Stability of each action's eigenvalues alone does not certify stability under switching.

**Formal route.** `DefectPropagation.lean`.

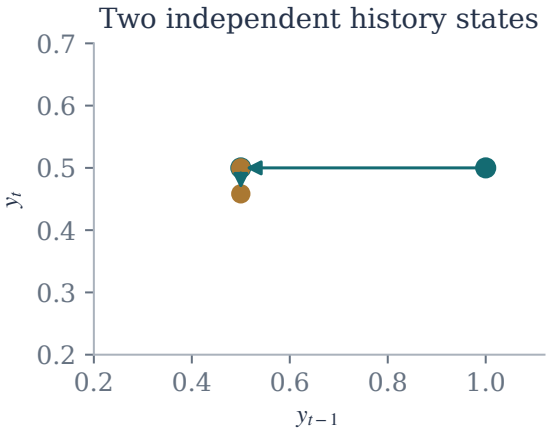
9 P14–P17

# What four observations identify

Data span, closure and coverage are separate checks.

The running example begins $1, \frac{1}{2}, \frac{1}{2}, \frac{11}{24}, \frac{61}{144}$. Its first four outputs give

$$Z_- = \begin{pmatrix} 1 & 1/2 \\ 1/2 & 1/2 \end{pmatrix}, \quad Z_+ = \begin{pmatrix} 1/2 & 1/2 \\ 1/2 & 11/24 \end{pmatrix}, \quad Z_+ Z_-^{-1} = \begin{pmatrix} 0 & 1 \\ 1/12 & 5/6 \end{pmatrix}.$$




## Input and successor windows

$\det Z_- = 1/4$

$D_2 = Z_+ Z_-^{-1}$

The determinant $\det Z_- = 1/4$ establishes coverage of the *declared* two-dimensional delay space. The resulting recurrence is

$$y_{t+1} = \frac{5}{6} y_t + \frac{1}{12} y_{t-1}, \qquad t \geq 1.$$

For general paired states, a linear fit exists exactly when $\ker Z_- \subseteq \ker Z_+$. It is unique on $W = \operatorname{im} Z_-$, and maps $W$ to itself exactly when $\operatorname{im} Z_+ \subseteq W$.

**The mathematical point.** These data identify the transition on their input span. The known generating system separately certifies that this span is the required full predictive state space.

**Scope.** Finite sampled rank does not exclude a larger unexcited state space. Finite continuation agreement does not establish unrestricted discovery. Perturbations additionally require a conditioning bound.

**Formal route.** `DelayClosure.lean`; `DataSpan.lean`; see the exact module map at the end.

10 P27–P28

# Behavior can supply the coordinates

Future tests define a predictive state without naming hidden variables.

For a controlled process, arrange population probabilities into a behavior table: columns index past histories and rows index future tests. Its full column span is the linear predictive behavior space.

| | $h_1$ | $h_2$ | $\cdots$ |
|---|---|---|---|
| $\tau_1$ | $p(h_1\tau_1)$ | $p(h_2\tau_1)$ | $\cdots$ |
| $\tau_2$ | $p(h_1\tau_2)$ | $p(h_2\tau_2)$ | $\cdots$ |
| $\vdots$ | $\vdots$ | $\vdots$ | |

Assume the full behavior space has rank at most $r$. Choose an invertible $r$-by-$r$ population core $H$. Let $G_s$ be the corresponding table with a declared action-observation symbol $s$ inserted. With a compatible row/column convention,

$$\boxed{B_s = G_s H^{-1}.}$$

The empty-test row supplies the mass functional and a covered history supplies the initial predictive state. The factorization then transports every future test through the observable coordinates.

**The mathematical point.** The dimension is determined by independent predictive behaviors. A global behavior-rank bound and an invertible core establish that the chosen coordinates cover the full behavior space.

**Two distinct obstructions.** An invertible sampled block does not establish the full-process rank bound. A fitted linear realization does not automatically preserve a probability domain. The next plates address the latter obligation.

**Scope.** The reconstruction uses population tables under the stated coverage hypotheses. No finite-sample rank test here authenticates those hypotheses or an unknown hidden-state model.

**Formal route.** `BehaviorSpace.lean`; `BehaviorMinimality.lean`; `ObservableCore.lean`; `ObservableRealization.lean`.

11 P29, P39–P40

# Probability meaning is geometric

An operator family needs a domain, a mass map and a covered initial state.

Let $C \subseteq \mathbb{R}^d$ be a convex cone and $\lambda$ a linear mass map strictly positive on $C \setminus \{0\}$. Set $K = \{z \in C : \lambda z = 1\}$. Linear observation operators must satisfy

$$B_o C \subseteq C, \qquad \lambda \sum_o B_o = \lambda, \qquad z_0 \in K.$$

Cone preservation gives positivity. The mass identity gives normalization.




$$p_o = \lambda B_o z \geq 0, \qquad \sum_o p_o = 1$$

A finite record $o_1 \cdots o_T$ has probability

$$\lambda B_{o_T} \cdots B_{o_1} z_0 \geq 0.$$

Summing over every branch repeatedly applies the mass identity and gives total probability one. A positive-probability posterior is $B_o z/(\lambda B_o z)$.

**A finite-generator witness.** Let $V$ have full row rank and normalized generator columns. Nonnegative lifts $A_o$ with $B_o V = V A_o$ and $\mathbf{1}^\mathsf{T} \sum_o A_o = \mathbf{1}^\mathsf{T}$ certify the generated cone. The initial state must be a normalized nonnegative mixture of those columns.

**The mathematical point.** All-horizon validity follows by induction from a finite exact witness. Statistical proximity to a generating process is a further obligation.

**Scope.** A zero-probability branch does not define a normalized posterior by division. A finite positive lift is sufficient, but some valid low-dimensional processes require curved domains.

**Formal route.** `CanonicalCone.lean`; `PositiveLift.lean`; `PositiveLiftTopology.lean`.

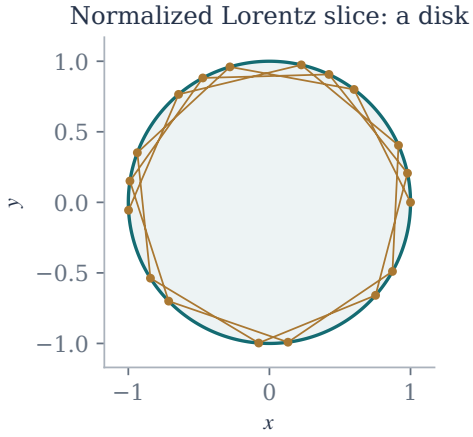
12 P46, P50–P51

# A curved domain can have a finite witness

A matrix embedding changes the certificate language.

The Lorentz cone and its normalized disk admit the representation

$$C = \{(m, x, y) : m \geq \sqrt{x^2 + y^2}\}, \qquad J(m, x, y) = \frac{1}{2}\begin{pmatrix} m + x & y \\ y & m - x \end{pmatrix}.$$




$$m \geq \sqrt{x^2 + y^2}$$

$$\det J = (m^2 - x^2 - y^2)/4$$

$$\lambda_{\pm}(J) = (m \pm \sqrt{x^2 + y^2})/2$$

Since the eigenvalues of $J$ are $(m \pm \sqrt{x^2 + y^2})/2$, membership in $C$ is equivalent to $J \succeq 0$. Also $\operatorname{tr} J = m$. For the exact spatial rotation $R = \left(\begin{smallmatrix} 3/5 & -4/5 \\ 4/5 & 3/5 \end{smallmatrix}\right)$ and $T = \operatorname{diag}(1, R)$,

$$J(Tz) = \frac{1}{5} W J(z) W^{\mathsf{T}}, \qquad W = \begin{pmatrix} 2 & -1 \\ 1 & 2 \end{pmatrix}.$$

**The mathematical point.** A rational congruence identity certifies the unchanged rotation on the curved cone. In the full P46 model, the canonical domain has no finite normalized positive-generator certificate.

**Scope.** The drawn orbit is finite. Density and the finite-generator obstruction are proved in the volume, not established by its plotted points. The full instrument also needs its observation, mass and initial-state identities. No quantum implementation is asserted.

**Formal route.** `RotationP46.lean`; `MatrixEmbedding.lean`; `MatrixRotation.lean`.

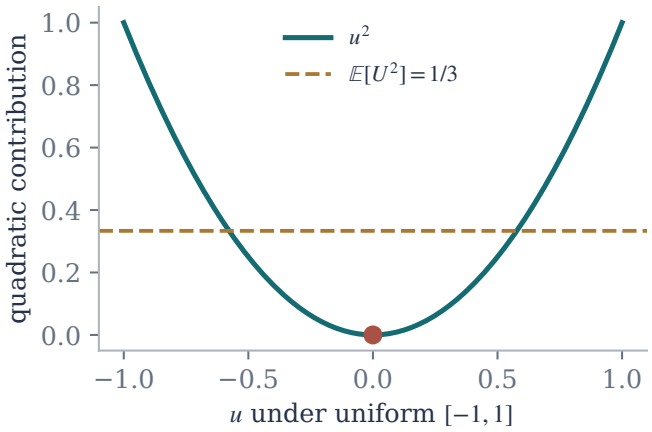
13 P56–P59

# Integration retains a covariance term

Averaging an operator before quadratic action changes the instrument.

Let $X \succeq 0$, $K(u) = K_0 + uK_1$ and $\bar{K} = \mathbb{E}[K(U)]$. For a scalar $U$ of finite second moment,

$$\mathbb{E}[K(U)XK(U)^{\mathsf{T}}] - \bar{K}X\bar{K}^{\mathsf{T}} = \mathrm{Var}(U)K_1XK_1^{\mathsf{T}} \succeq 0.$$




$$\bar{K} = \mathbb{E}[K(U)]$$

$$\mathbb{E}[KXK^T] - \bar{K}X\bar{K}^T$$

$$= \mathrm{Var}(U)\, K_1 X K_1^T$$

For the uniform probability measure $\mu(du) = du/2$ on $[-1, 1]$, take

$$K(u) = \tfrac{1}{2}\,\mathrm{diag}(1 + 3u, 1 - 3u), \qquad X = I/2.$$

Then $\bar{K} = I/2$. The averaged congruence is $I/2$, with trace one. Congruence by the averaged matrix is $I/8$, with trace $1/4$. The omitted covariance correction is $3I/8$.

**The mathematical point.** The missing term is positive semidefinite and carries three quarters of the mass in this one-bin example. With $n$ equal bins, the missing mass is $3/(4n^2)$.

**Scope.** The matrices in this example are diagonal and commute. The failure concerns the order of averaging and forming a quadratic action. A density value is interpreted relative to its declared measure.

**Formal route.** See P57–P58 in the formal coverage map for exact integral and bin identities.

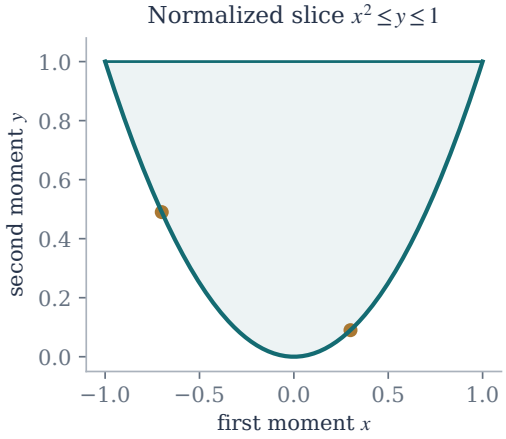
14 P72, P83

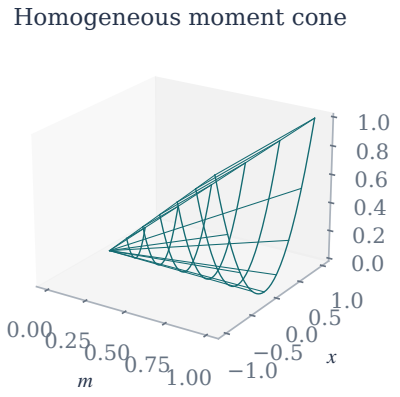
# Uncertainty on a curved domain

Parameter uncertainty and arbitrary entrywise noise have different geometry.

For $u \in [-1,1]$, set $v(u) = (1, u, u^2)^\mathsf{T}$. Its convex hull is

$$K = \{(1,x,y) : x^2 \le y \le 1\}, \qquad C = \left\{(m,x,y) : \begin{pmatrix} m & x \\ x & y \end{pmatrix} \succeq 0,\ m - y \ge 0\right\}.$$




Choose a reference probability measure on $[-1,1]$ with moments $\mathbb{E}U = \mathbb{E}U^3 = 0$, $\mathbb{E}U^2 = 1/3$ and $\mathbb{E}U^4 = 1/5$. With $L_2(u) = (3u^2 - 1)/2$, define

$$B_{a,b}(u) = v(u)(1, au, bL_2(u)).$$

Its normalized-state density is $p_{a,b}(u \mid x, y) = 1 + axu + byL_2(u)$. If $|a| + |b| \le 1 - \rho$ for $\rho \in [0,1]$, it is at least $\rho$, and integrated mass is exact.

**The mathematical point.** Every positive-density posterior is the same boundary point $v(u)$ throughout this parameter family. Uncertainty changes its probability without moving its posterior ray.

**Scope.** This protection depends on the supplied algebraic family. Free perturbations of output coordinates can leave the cone. The three-dimensional wireframe shows the homogeneous geometry, not sampled evidence.

**Formal route.** `CurvedMomentGeometry.lean`; `CurvedMomentInstrument.lean`; exact P83 map in `COVERAGE.json`.

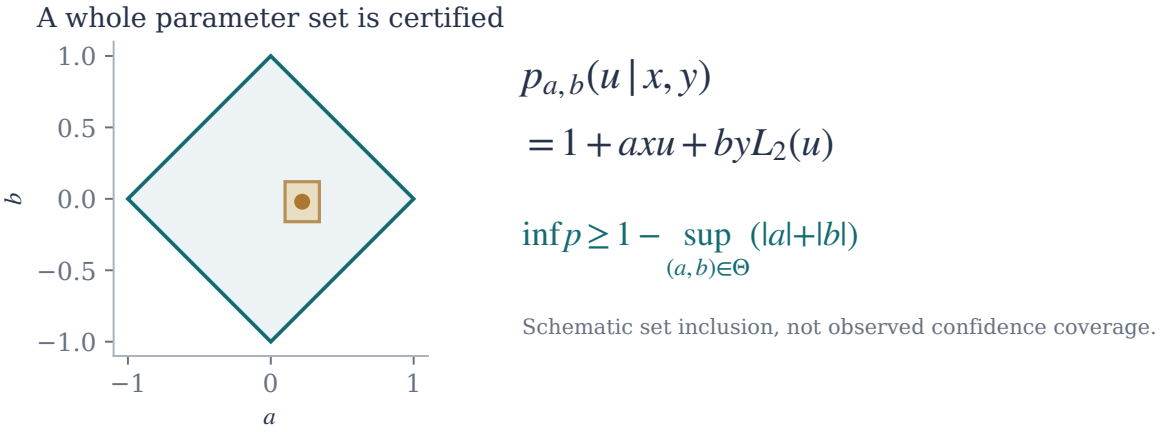
15 P66, P68, P81–P84, P87

# A confidence set can meet an exact certificate

The statistical event and the internal validity check do different work.

Assume $Y_t = \theta_*^{\mathsf{T}}\phi_t + \xi_t$ with predictable $\phi_t$, conditional sub-Gaussian noise of scale $\sigma > 0$ and $\|\theta_*\|_2 \le S$. For ridge $\eta > 0$, define $V_t = \eta I + \sum_{s\le t} \phi_s\phi_s^{\mathsf{T}}$ and $\hat\theta_t = V_t^{-1}\sum_{s\le t}\phi_s Y_s$. For $0 < \delta < 1$, with probability at least $1-\delta$, simultaneously for every finite $t$,

$$\|\hat\theta_t - \theta_*\|_{V_t} \le \sigma\sqrt{\log\frac{\det V_t}{\det(\eta I)} + 2\log(1/\delta)} + \sqrt{\eta}S.$$




In the structured moment family, a nonempty returned rectangle $\Theta$ is internally valid when its sufficient floor

$$\rho_\Theta = 1 - \sup_{(a,b)\in\Theta} (|a| + |b|)$$

is nonnegative. Statistical coverage places the true parameter in $\Theta$ on a simultaneous event. The exact domain argument applies to *every* parameter in $\Theta$.

**The mathematical point.** A data-selected time is allowed by the simultaneous confidence statement. It does not convert the model, noise or predictability assumptions into observed facts.

**Scope.** Predictable residual rejection can give a lower discrepancy certificate. Nonrejection alone supplies no upper accuracy guarantee. The drawn rectangle is schematic, not an estimated confidence set.

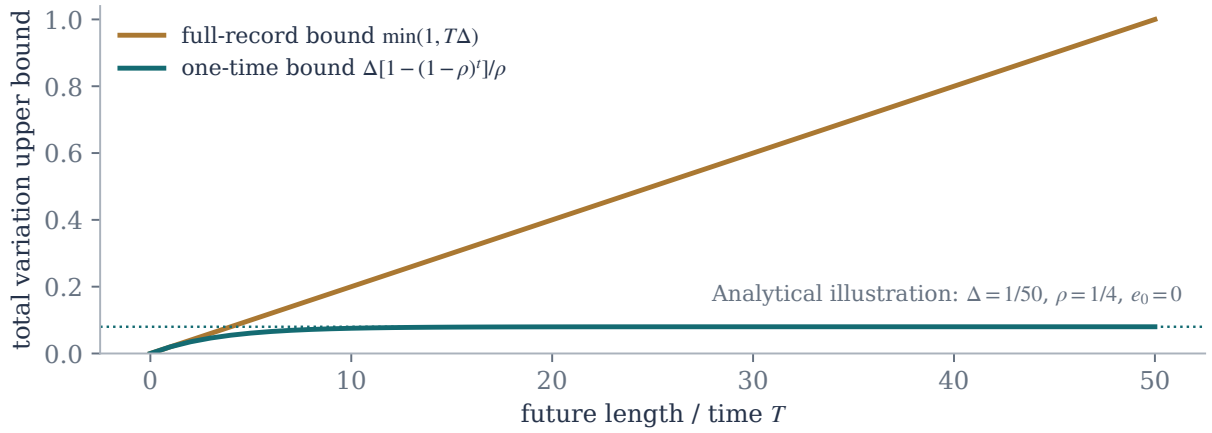
**Formal route.** `RegressionConfidence.lean`; `PredictableResidualP66.lean`; `StructuredLearningCapstone.lean`.

16 P85–P87

# Choose the future quantity before its bound

A long record and one distant observation are different objects.

Start two processes from the same history or initial law and use the same action policy. If their conditional next-observation laws differ by at most $\Delta$ at every common history, telescoping gives the record bound $\min(1, T\Delta)$ in total variation.




For Markov kernels with a common probability measure $\mu$ and $P(s,\cdot) \geq \rho\mu(\cdot)$, $0 < \rho \leq 1$, and $\sup_s \|P(s,\cdot) - Q(s,\cdot)\|_{\mathrm{TV}} \leq \Delta$, their one-time law distances satisfy

$$e_t \leq \min\left(1, (1-\rho)^t e_0 + \frac{\Delta[1-(1-\rho)^t]}{\rho}\right).$$

A length-$k$ rolling window starting at time $t$ has distance at most $\min(1, e_t + (k-1)\Delta)$.

**The mathematical point.** Contraction can keep a single marginal close indefinitely while complete-record distinguishability tends to one. Two distinct independent Bernoulli laws already exhibit that separation.

**Scope.** The plotted curves are analytical illustrations with declared constants. The marginal result needs actual Markov and minorization hypotheses. Arbitrary policies that remember the entire history require their own comparison argument.

**Formal route.** `CurvedMomentRecordTV.lean`; `MarkovMarginals.lean`; `MarkovWindows.lean`; `BernoulliRecordSeparation.lean`.

17 Navigation

# From the atlas to the proofs

The visual route points into a source-bound formal companion.

The accepted Lean release records complete mappings for all 87 registered proposition cores. Its scope is stated proposition by proposition in `Lean/COVERAGE.md` and `Lean/COVERAGE.json`. Each plate selects a mechanism from those statements.

| Plates | Proof route | Primary entry modules in `Lean/KomorebiLSM/` |
|---|---|---|
| 1–2 | Reconstruction and transport | `Realization`, `Normed` |
| 3–4 | Closure and completion | `PredictiveClosure`, `LinearClosure`, `InvariantCompletion` |
| 5–6 | Packets and return paths | `PacketTree`, `Transfer`, `Polynomial`, `LinearClosure` |
| 7–8 | Memory and error propagation | `Memory`, `MemoryBounds`, `MemoryStability`, `DefectPropagation` |
| 9–10 | Observable and behavior states | Exact P14–P17 and P27–P28 entries in `COVERAGE.json` |
| 11–12 | Probability witnesses | `PositiveLift`, `RotationP46`, `MatrixEmbedding`, `MatrixRotation` |
| 13–14 | Integration and curved moments | Exact P57–P58 and P83 entries in `COVERAGE.json` |
| 15–16 | Confidence and horizon bounds | `RegressionConfidence`, `StructuredLearningCapstone`, `MarkovMarginals` |

**A small notation key.** $A, S$ denote analysis and synthesis; $\mathcal{C} = \operatorname{im} A$ denotes realizable coefficients; $\gamma$ denotes a predictive summary; $P, Q$ denote retained and discarded projections; $W_*$ denotes the completed observable row space; $C$ denotes a probability cone in the stochastic plates; $\lambda$ denotes its mass map. Symbols local to a plate are introduced there.

**Evidence boundaries.** Written proofs, exact example calculations, kernel verification, and statistical statements are separately identified. The formal release verifies the registered mathematical cores under their hypotheses. It does not empirically authenticate a model class, noise assumption, sensor or physical process.

**Reproducibility.** The companion includes the editable LaTeX source, Python figure generator, exact rational example receipt and a proposition/module map. All illustrations are mathematical diagrams or declared analytical calculations. The original volume remains the source of the full statements and proofs.