The Minimal Sufficient Spectral State for Decision and Risk

Dr. Tamás Nagy

tnagyphd@gmail.com

Draft

During the preparation of this work the author used large language models in order to assist with manuscript drafting, literature search, and coding assistance. After using these tools, the author reviewed and edited the content as needed and takes full responsibility for the content of the published article.

Reader-Friendly Subtitle

How much state is truly needed for good decisions under uncertainty.

Technical Strapline

Necessary-and-sufficient characterization of information-complete spectral compression for value and risk.

Executive Summary (Non-Technical)

Many decision systems carry too much state. That increases complexity, cost, and fragility without always improving outcomes.

This paper asks for the minimum: the smallest spectral state that still preserves optimal decision value and calibrated risk behavior for a target decision class.

If such a state can be characterized exactly, we gain a rigorous compression principle: not heuristic simplification, but theorem-backed sufficiency.

The paper does not claim one universal minimal state for all tasks. It characterizes minimality relative to an explicit decision and risk functional family.

Abstract

Decision systems often carry excessive state, obscuring both theory and implementation. We define and characterize the minimal sufficient spectral state needed to preserve optimal decision value and calibrated risk estimation. The main theorem gives necessary and sufficient conditions under which a reduced spectral state is information-complete for a target decision class.

1. Problem

Decision and risk pipelines routinely carry full distributional state when only a fraction is decision-relevant. This overhead increases latency, storage, and fragility. The question is: what is the minimal spectral state that preserves decision optimality and risk calibration?

This is not a compression heuristic. We seek a theorem: a necessary-and-sufficient characterization of the boundary between lossless and lossy spectral state reduction for a given decision class.

2. Setup

2.1 Decision and Risk Functional Families

Let P be a probability distribution over portfolio losses. Define: - Decision class $\mathcal{D} = \{d : \mathcal{P} \to \mathcal{A}\}$: maps from distributions to actions. - Risk class $\mathcal{R} = \{\rho : \mathcal{P} \to \mathbb{R}\}$: coherent or convex risk measures.

2.2 Spectral State and Compression

The full spectral state is $\mathbf{A}(P) = (A_0, A_1, \ldots, A_{N-1})$, the COS coefficients of the CDF of P.

Definition 1 (Spectral Compression). The K-truncation map is:

$$\Pi_K : \mathbf{A} \mapsto (A_0, A_1, \ldots, A_{K-1}, 0, 0, \ldots)$$

The compressed distribution $P_K = \Pi_K(P)$ retains only the first K spectral modes.

2.3 Sufficiency Criterion

Definition 2 (Decision-Risk Sufficiency). The compressed state $\Pi_K(P)$ is $(\mathcal{D}, \mathcal{R})$-sufficient if: - $d(\Pi_K(P)) = d(P)$ for all $d \in \mathcal{D}$, and - $|\rho(\Pi_K(P)) - \rho(P)| \leq \epsilon$ for all $\rho \in \mathcal{R}$.

Connection to existing kernel: the URRT (Universal/MainTheorem.lean) already proves $N = \Theta(\log(1/\epsilon)/\log \rho)$ parameters suffice for CDF approximation. The present paper lifts this from CDF approximation to decision-functional preservation.

3. Main Theorem

Theorem Candidate 1 (Minimal Sufficient Spectral State). Let $\mathcal{R}$ be the class of all Lipschitz-continuous spectral risk measures with Lipschitz constant L_ρ. Then the minimal K^* satisfying $(\mathcal{D}, \mathcal{R})$-sufficiency at tolerance ϵ is:

$$K^* = \left\lceil \frac{\log(L_\rho/\epsilon)}{\log \rho} \right\rceil$$

where $\rho > 1$ is the spectral gap of the underlying distribution family.

Moreover, $K^* - 1$ modes are not sufficient: there exists $\rho \in \mathcal{R}$ with $|\rho(\Pi_{K^*-1}(P)) - \rho(P)| > \epsilon$.

Corollary (Decision Independence). If $d \in \mathcal{D}$ depends only on quantiles and spectral risk measures, then the same K^* is sufficient for decision optimality.

4. Proof Sketch

1. **Upper bound (sufficiency).** From URRT, the CDF error under K-truncation satisfies $\|F - F_K\|_\infty \le C\rho^{-K}$. For Lipschitz risk measures, $|\rho(P) - \rho(P_K)| \le L_\rho \|F - F_K\|_\infty \le L_\rho C \rho^{-K}$. Setting $L_\rho C \rho^{-K} \le \epsilon$ gives the upper bound.
2. **Lower bound (necessity).** Construct a pair of distributions P, P' agreeing on modes $1, \ldots, K^* - 1$ but differing on mode K^*, with $|\rho(P) - \rho(P')| > \epsilon$ for a specific $\rho \in \mathcal{R}$. This follows from the entropy lower bound in Universal/EntropyLowerBound.lean.
3. **Decision transfer.** If decisions depend on P only through risk measures and quantiles, the CDF-level sufficiency propagates to decision-level sufficiency.

5. Empirics/Simulation

5.1 Portfolio Risk

- 100-asset portfolio with lognormal marginals.
- Compare VaR/ES accuracy at $K = 5, 10, 20, 50, 128$ modes.
- Report the empirical K^* where decision quality saturates.

5.2 Allocation Quality

- Run Markowitz optimization with full vs compressed state.
- Measure utility loss, Sharpe ratio difference, and turnover.

5.3 Compute Savings

- Report wall-clock time for full vs minimal-state risk engine on 1000-scenario stress grid.

6. Limits

- **Task dependence:** K^* depends on $\mathcal{R}$ and $\mathcal{D}$; different tasks may have different minimal states.
- **Nonstationary distributions:** ρ may change over time, requiring adaptive K.
- **Non-Lipschitz functionals:** discontinuous decision rules (hard thresholds) are not covered.

7. Related Work

- **Sufficient statistics:** Fisher (1922), Halmos-Savage (1949).
- **Information bottleneck:** Tishby et al. (1999).
- **POMDP compression:** Pineau et al. (2003), belief-space methods.
- **URRT:** our Universal paper — the CDF-level compression result that this paper extends to decisions.
- **Risk coding theorem:** our Risk Information paper — the information-theoretic foundation.

8. Cross-Paper Connections

- **Decision Functional Approximation (paper 8):** the K^* derived here is the same approximation budget that governs the universal approximation rate ρ^{-K} in paper 8. Together, papers 4 and 8 say: "you need exactly K^* modes, and with K^* modes any Lipschitz decision functional is approximated at rate ϵ."
- **Causal Identifiability (paper 3):** K^* modes must be observable (paper 3 sense) for the minimal sufficient state to be constructible from data. If K^* exceeds the number of identifiable modes, the decision-optimal compression is unreachable.
- **Phase Transitions (paper 5):** the spectral gap ρ that determines K^* can itself undergo a phase transition as model complexity changes. At criticality, $\rho \to 1$ and $K^* \to \infty$ — the minimal state diverges.
- **Nonlinear FTAP (paper 2):** in the nonlinear pricing setting, the minimal mode budget for consistent pricing is a special case of K^* with $\mathcal{R}$ replaced by the class of nonlinear pricing functionals $\{\Pi_k\}$.

9. Adaptive K Monitoring Protocol

9.1 Motivation

In practice, the spectral gap ρ is not constant. Regime changes, structural breaks, and distribution drift can change ρ, which in turn changes the minimal K^*. A deployed system needs to track K^* online.

9.2 Algorithm

Input: streaming spectral coefficients $A_0(t), A_1(t), \ldots$ at each time step.

1. **Estimate spectral gap:** fit $|A_k|$ to geometric decay $C\rho^{-k}$ over a rolling window of W observations. Extract $\hat{\rho}(t)$.
2. **Compute current K^*:**

$$K^*(t) = \left\lceil \frac{\log(L_\rho/\epsilon)}{\log \hat{\rho}(t)} \right\rceil$$

3. **Regime detection:** if $|K^*(t) - K^*(t-1)| > \Delta_{\text{crit}}$, flag a regime transition. Common causes: volatility regime change, structural break, liquidity shock.
4. **State resize:** expand or contract the spectral state vector. When K^* increases, new modes are initialized from the current distribution estimate. When K^* decreases, trailing modes are dropped with explicit loss accounting.

9.3 Safety Margin

In practice, set $K_{\text{deploy}} = K^*(t) + K_{\text{buffer}}$ where $K_{\text{buffer}} = \lceil 2/\log \hat{\rho} \rceil$ provides one order of magnitude safety beyond the sufficiency threshold.

10. Explicit Portfolio Walkthrough

10.1 Setup

A 5-asset equity portfolio with lognormal marginals. The joint distribution is described by 5 mean returns, 5 volatilities, and 10 correlation parameters (25 total).

Spectral COS coefficients: $\mathbf{A} = (A_0, A_1, \dots, A_{N-1})$ with $N = 128$.

Measured spectral gap: $\hat{\rho} = 3.2$ (typical for lognormal with moderate correlation).

10.2 VaR/CVaR Sufficiency

For $\text{CVaR}_{0.99}$, Lipschitz constant $L = 1/(1 - 0.99) = 100$. Target tolerance $\epsilon = 0.01$ (1bp of portfolio value).

$$K^* = \left\lceil \frac{\log(100/0.01)}{\log 3.2} \right\rceil = \left\lceil \frac{\log 10000}{1.163} \right\rceil = \left\lceil \frac{9.21}{1.163} \right\rceil = 8$$

Result: 8 spectral modes suffice to preserve CVaR at 1bp accuracy, versus 128 modes in the full representation — a 16x compression.

10.3 Markowitz Sufficiency

The Markowitz allocation depends on the distribution through mean and covariance, both of which are finite-mode functionals. For a 5-asset portfolio, the allocation is a Lipschitz function of the first $\sim d^2/2 = 15$ spectral modes in the multivariate COS expansion.

This implies $K^*_{\text{Markowitz}} \leq 15$, consistent with the risk-measure result.

10.4 Compute Savings

On a 1000-scenario stress grid: - Full state ($K = 128$): $O(128 \times 1000)$ evaluations per step. - Minimal state ($K = 8$): $O(8 \times 1000)$ evaluations per step. - Speedup: $\sim 16\times$ with guaranteed $\leq$ 1bp error.

11. Outlook

- **Auditable deployment:** the theorem provides a certificate for state budget choices. The portfolio walkthrough (Section 10) demonstrates the audit trail for a concrete deployment.
- **Adaptive K:** the monitoring protocol (Section 9) enables online state management.
- **Lean formalization:** the upper bound follows directly from existing URRT Lean chain; lower bound extends EntropyLowerBound.lean to functional loss. The portfolio sufficiency computation is a natural numerical validation target.
- **Integration with URRT engine:** the existing COS engine already supports variable K; wrapping it with the adaptive monitoring protocol is a near-term engineering task.