Spectral Causal Identifiability Under Partial and Noisy Observation

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

When causal structure can still be recovered from incomplete, noisy measurements.

Technical Strapline

Identifiability criteria in terms of mode observability, spectral separation, and noise geometry.

Executive Summary (Non-Technical)

In practice, we almost never observe full system state. Most causal pipelines therefore operate under partial information and significant measurement noise.

This paper asks a strict question: under what conditions is causal structure still recoverable in that setting? The proposed answer is spectral. Recovery depends on which dynamic modes are visible and how cleanly they separate from noise.

The practical value is high: a theorem-level boundary between "causally recoverable" and "structurally ambiguous" regimes can save large experimental and modeling effort.

The paper does not claim universal identifiability. It aims to characterize exactly when identifiability is possible and when it provably fails.

Abstract

Causal identification typically requires rich interventions or full-state observability. We show that partial, noisy observations can still identify causal structure when signals admit a stable spectral decomposition with separation conditions. We derive identifiability criteria in terms of mode observability, spectral gap structure, and noise geometry, and provide constructive recovery guarantees.

1. Problem

Causal discovery typically assumes either full-state observability or the ability to perform controlled interventions. In reality, practitioners observe only a noisy projection of a latent multivariate system.

The fundamental question is: when is causal structure still identifiable from such compressed, noisy traces?

Existing theory (Spirtes et al. 2000, Pearl 2009) focuses on graphical conditions. We propose a complementary spectral approach: identifiability is governed by mode observability, spectral separation, and noise geometry rather than only graphical faithfulness.

2. Setup

2.1 Latent Structural System

Let $Z_t \in \mathbb{R}^d$ be a latent VAR(1) process:

$$Z_t = AZ_{t-1} + \epsilon_t, \quad \epsilon_t \sim \mathcal{N}(0, \Sigma_\epsilon)$$

where the causal structure is encoded in the sparsity pattern of A.

Diagonalize: $A = U \operatorname{diag}(\mu_1, \dots, \mu_d) U^{-1}$ with spectral modes $c_k(t) = u_k^\top Z_t$.

2.2 Observation Operator

We observe $Y_t = HZ_t + \xi_t$ where $H \in \mathbb{R}^{m \times d}$ with $m < d$ and $\xi_t \sim \mathcal{N}(0, \Sigma_\xi)$.

Definition 1 (Mode Observability). Mode k is observable if $\|Hu_k\| > 0$. The observability index is $o_k = \|Hu_k\|^2 / \|u_k\|^2 \in [0, 1]$.

2.3 Spectral Separation

Definition 2 (Spectral Gap Condition). The system satisfies spectral gap condition SG(δ) if $|\mu_j - \mu_k| \geq \delta$ for all $j \neq k$ with $o_j, o_k > 0$.

2.4 Identifiability Target

- **Graph identifiability:** recover the sparsity pattern of A.
- **Effect identifiability:** recover A_{ij} for specific (i, j) pairs.
- **Intervention identifiability:** recover the causal effect of $\operatorname{do}(Z_i = z)$.

3. Main Theorem

Theorem Candidate 1 (Spectral Causal Identifiability). Suppose SG(δ) holds and at least r modes are observable with $o_k \geq o_{\min} > 0$. Then:

1. If $r \geq d$, the full causal graph of A is identifiable from $\{Y_t\}$ up to estimation error $O(\sigma_\xi / (o_{\min} \cdot \delta))$.
2. If $r < d$, the causal graph restricted to observable modes is identifiable, and unobservable modes induce bounded confounding bias $\leq C \sum_{k: o_k = 0} |\mu_k| / \delta$.

Theorem Candidate 2 (Impossibility Boundary). If $o_k = 0$ for mode k and $|\mu_k - \mu_j| < \delta/2$ for some observable j, then the causal effect mediated by mode k is not identifiable from observations alone.

4. Proof Sketch

1. **Spectral observation equations.** Project VAR dynamics into mode space: $\tilde{c}_k(t) = \mu_k \tilde{c}_k(t-1) + \tilde{\epsilon}_k(t) + \tilde{\xi}_k(t)$ where tildes denote observed projections.
2. **Invertibility.** Under SG(δ) and $o_k > 0$, the transfer function $(zI - A)^{-1}$ restricted to observable modes is injective on mode parameters.
3. **Stability.** Perturbation theory for eigenvalues (Weyl, Davis-Kahan) gives estimation error bounds proportional to $\sigma_\xi/(o_{\min} \cdot \delta)$.
4. **Impossibility.** When modes collide spectrally and one is unobservable, the corresponding column of A is not separable from the noise-floor contribution.

5. Empirics/Simulation

5.1 Synthetic VAR Systems

- Generate VAR(1) with planted spectral structure ($d = 20$, $m = 8$).
- Sweep observability ($o_{\min}$) and spectral gap (δ).
- Report graph recovery F1 score.

5.2 Baselines

- Compare against PC algorithm, FCI, and LASSO-Granger.
- Report advantage in low-observability regimes.

5.3 Financial Application

- Equity sector factors as latent modes.
- Observed: sector ETF returns (partial projection of full factor space).
- Test: which causal links between factors are recoverable?

6. Limits

- **Nonlinear causality:** the VAR(1) assumption is restrictive; nonlinear extensions need further work.
- **Collapsed modes:** when eigenvalues cluster, SG(δ) fails and identifiability degrades.
- **Adversarial confounding:** strategic noise placement can make specific causal links unrecoverable.
- **Finite-sample regime:** practical recovery quality depends on trajectory length relative to d/m.

7. Related Work

- **Causal discovery:** Spirtes-Glymour-Scheines (2000), Pearl (2009), FCI (Zhang 2008).
- **Latent confounders:** Silva-Scheines (2006), Anandkumar et al. (2012).
- **System identification:** Ljung (1999), spectral methods for linear systems.

- **Spectral Granger:** our existing fin_spectral_granger topic — the present paper extends from bivariate Granger to full structural graph recovery.

8. Cross-Paper Connections

- **Minimal Sufficient State (paper 4):** which causal modes need observation for decision optimality? Paper 4 gives K^*; the present paper says whether those K^* modes are identifiable. Together they answer: "can I observe enough to make a good decision?"
- **Ergodic Control (paper 7):** if the controller acts on modes it cannot causally identify, the ergodic guarantee degrades. The observability index o_k from this paper constrains which modes the controller in paper 7 can reliably act on.
- **Phase Transitions (paper 5):** as the spectral gap δ crosses a critical threshold, identifiability undergoes a phase transition — from full recovery to partial recovery. This connects to the cusp behavior in paper 5.
- **Universality (paper 1):** if the latent system belongs to a universality class, its spectral exponent s constrains the eigenvalue gap structure, which in turn determines the identifiability regime via SG(δ).

9. Optimal Intervention Design

9.1 Problem Statement

Given a budget B for interventions, choose an intervention matrix ΔH to maximize the number of identifiable causal links. Formally:

$$\max_{\Delta H : \|\Delta H\|_F \le B} |\{(i,j) : A_{ij} \text{ is identifiable under } H + \Delta H\}|$$

9.2 Greedy Mode-Targeting Strategy

Observation: identifiability requires $o_k > 0$ and SG(δ). Since δ is fixed by A, the design variable is o_k.

Rank unobservable modes by their contribution to confounding bias: $b_k = |\mu_k|/\delta$. Prioritize modes with largest b_k.

For each such mode, choose ΔH to make row k of HU nonzero. The cost is $\|\Delta h\| = 1/\|u_k\|$ to achieve $o_k > 0$.

Greedy algorithm: 1. Sort unobservable modes by descending b_k. 2. While budget remains, add an observation row targeting the highest-bias unobservable mode. 3. Each addition reduces confounding bias by b_k and may unlock new causal links.

9.3 Sample Complexity

Theorem Candidate 3 (Sample Complexity for Graph Recovery). Under SG(δ) with all modes observable ($o_k \ge o_{\min}$), the number of time-series observations needed to recover the sparsity pattern of A with probability $1 - \delta_{\text{fail}}$ is:

$$T \geq C \cdot \frac{d^2 \log(d/\delta_{\text{fail}})}{o_{\min}^2 \cdot \delta^2 \cdot \mathrm{SNR}^2}$$

where $\mathrm{SNR} = \lambda_{\min}(\Sigma_\epsilon)/\lambda_{\max}(\Sigma_\xi)$.

This makes the dependence on observability, spectral gap, and noise explicit. In the low-observability regime ($o_{\min} \to 0$), the sample complexity diverges — matching the impossibility boundary of Theorem 2.

10. Outlook

- **Active interventions:** the greedy strategy (Section 9.2) provides a practical starting point; optimal design with combinatorial constraints is future work.
- **Online causal control:** adaptive policies that simultaneously learn causal structure and optimize. The sample complexity result (Theorem 3) gives the exploration phase length.
- **Lean formalization:** the impossibility boundary (Theorem 2) is a clean target for formal verification. The sample complexity lower bound is also natural.
- **Bridge to decision theory:** combine with the Minimal Sufficient Spectral State paper (paper 4) to determine which causal modes need observation for decision optimality.
- **Empirical priority:** test the greedy intervention strategy on synthetic VAR systems and validate the predicted sample complexity scaling.