

Generic Separation of Axis-Normalized Latent-Source Representations by Higher-Order Cumulants

CausalSmith*

September 1, 2026

Abstract

We study a population identification question for a bivariate latent linear non-Gaussian model in which the number m of middle source slots and an axis normalization are maintained inputs. The representation-level target asks whether a truncated joint-cumulant vector through order L can also arise from the opposite axis-normalized arrow convention with the same m . At $L = 2m + 2$, the unordered finite loading-slope support is generically recovered and the opposite cumulant-map fiber is empty, while the same-arrow parametrization retains ambiguity beyond middle-slot relabelling. The closure of the opposite-arrow compatibility locus has codimension exactly one in each arrow image variety. For $m \geq 3$, generic real opposite-fiber exclusion already holds at order $2m + 1$; loading-support recovery uses the apolar order $2m + 2$. There exist Euclidean-open parameter neighborhoods whose intersections with the feasible regions are nonempty, sharpening the population geometry of the maintained representation class.

1 Introduction

Econometric models often become directionally informative only after the researcher imposes structure beyond covariance restrictions. In linear non-Gaussian models, higher-order cumulants encode the loading directions of independent shocks and can therefore restrict which recursive representation is compatible with the population distribution. This mechanism underlies LiNGAM and its latent-variable extensions [Shimizu et al., 2006, Hoyer et al., 2008, Tashiro et al., 2014, Maeda and Shimizu, 2020], and is closely connected to independent component analysis and cumulant-based tensor decompositions [Comon, 1994, Hyvärinen and Oja, 2000, Cardoso, 1999, Kolda and Bader, 2009].

This paper isolates a population question for a fixed number m of known or externally specified middle source slots and the two axis-normalized bivariate latent linear non-Gaussian representation classes in Definition 1 and Definition 2. One convention tags the representation as $X \rightarrow Y$, the other as $Y \rightarrow X$. The observed object is the finite vector of joint cumulants $T_L(P)$ in Definition 3. The target is whether this vector has an opposite-arrow cumulant-map preimage with the same source count and normalization. It serves settings in which an institutional model, factor design, or separate model-selection analysis fixes the number of common shock slots and the researcher evaluates population compatibility of the maintained reverse representation.

*Machine-generated by the CausalSmith research pipeline. The displayed formal statements and proofs are Lean-verified under the paper’s theorem-local citation interface; where a result from the literature is a formal dependency, it enters as a published input rather than a Lean-checked step, and cited work otherwise supplies framing and positioning.

Econometric target and decision. The maintained model class contains two competing structural representations with the same externally fixed m . The population decision rejects a convention when the observed truncated cumulant vector has an empty fiber under its polynomial cumulant map. The resulting exclusion restriction is representation-level and comes with an exact description of the compatibility set and the residual ambiguity within the surviving convention. These questions matter before estimation: they establish the population target for a cumulant-based procedure and identify the algebraic surface near which such a procedure would be weak.

Contributions. The paper establishes three facts about this target. First, the apolar order $2m + 2$ generically recovers the loading-slope support and excludes the opposite fiber. Second, opposite-arrow compatibility has a codimension-one closure, substantially larger than a dimension count based on the m middle slots might suggest. Third, for $m \geq 3$, generic real opposite-fiber exclusion requires one less cumulant order than support recovery. A separate fiber result characterizes the remaining same-arrow parameter ambiguity.

The main result at the apolar order $K = 2m + 2$ is [Theorem 1](#). On the stated relatively Zariski-open dense parameter loci, and on the Euclidean-open parameter neighborhoods supplied by the theorem whose intersections with the real feasible regions are nonempty, the order- K cumulants recover the unordered finite loading-slope support through binary-form apolar equations and rule out a full representation in the opposite arrow convention. Its exact domain is the pair of polynomial cumulant maps and their parameter fibers.

The same-arrow corollary in [Theorem 2](#) adds the corresponding fiber behavior on the generic loci of the apolar theorem. The same cumulants preserve the unordered loading-slope multiset while leaving multiple same-arrow parametrizations beyond the admissible middle-source relabellings in [Definition 9](#). In particular, a direct slot can be exchanged with a middle slot while preserving the cumulant vector and remaining outside the admissible label orbit. Thus generic arrow separation delivers more information than covariance equivalence while retaining precisely characterized structural-parameter ambiguity.

The third result, [Theorem 3](#), describes the failure set at the same apolar order. The cumulant vectors that have a generic representation for one arrow and at least one full representation for the other form an exceptional compatibility locus whose closure has codimension exactly one inside each arrow image variety. For $m \geq 2$, the exceptional compatibility closure has codimension exactly one in each arrow image variety. This codimension-one statement is important for interpretation: generic exclusion holds off the exceptional set, whose size makes proximity to compatibility surfaces empirically relevant.

The final identification result concerns the cumulant information needed for generic real opposite-arrow exclusion. [Theorem 4](#) shows that, for $m \geq 3$, order $2m + 1$ already suffices for generic separation over the real feasible regions of [Definition 10](#). This improves the real information-order bound relative to the apolar support-recovery order, which remains the relevant threshold when the target includes recovery of the loading support.

The paper is closest to the cumulant-based latent LiNGAM literature. [Chen et al. \[2025b\]](#) develop a directional testing construction from a higher-cumulant rank deficiency at order $2m + 3$. [Schkoda et al. \[2024\]](#) give recursive higher-cumulant rank conditions for unobserved confounding. [Tramontano et al. \[2024, 2025\]](#) target causal-effect identification from cumulants. Our one-order improvement pertains specifically to fixed-slot, axis-normalized population fiber exclusion; the comparison below aligns each result with its statistical or algebraic target.

The comparison with [Chen et al. \[2025b\]](#) turns on distinct targets. In their result, m counts latent confounders in a real statistical testing problem, and the order- $2m + 3$ rank condition supports a causal-direction test. Here m counts externally fixed middle source slots, the axis normalization is maintained, and the conclusion is generic emptiness of an opposite polynomial-map fiber together with support-recovery and codimension-one geometry. The difference in order is therefore indexed to these respective targets and assumptions.

Table 1: Auditable comparison with the closest cumulant-based approaches

Work	Source-count and normalization	Population object	Scope of conclusion	Statistical output
Classical LiNGAM/ICA	model-specific dimension and mixing structure	observational distribution or mixing representation	ordering or mixing identification under the stated statistical model	estimators or algorithms are central
Chen et al. [2025b]	m enters a latent-confounder testing construction	higher-cumulant rank condition	directional testing at order $2m + 3$ in a real statistical model	test construction
Schkoda et al. [2024]	latent-confounding structure encoded by recursive rank restrictions	higher-cumulant tensors	rank conditions for unobserved confounding	identification conditions
Tramontano et al. [2024, 2025] This paper	model-specific latent structure externally fixed m , fixed source slots, and axis normalization	cumulants relevant to causal effects truncated population cumulant vector	causal-effect identification complex-generic opposite-fiber exclusion and support recovery at $2m + 2$; codimension-one compatibility; separate real-generic exclusion at $2m + 1$ for $m \geq 3$	constructive identification methodology population fiber characterization

A one-middle-slot interpretation. When $m = 1$, the forward convention can be written

$$X = S_0 + S_1, \quad Y = \gamma S_0 + \rho S_1 + S_2.$$

An econometrician may regard S_1 as one externally maintained common-shock slot and S_0, S_2 as the two axis-normalized innovation slots. If both loadings of S_1 are nonzero, it has the usual latent-confounder interpretation; the theorem uses the more general source-slot representation. At order four, [Theorem 1](#) says that, off its stated algebraic exceptional set, the resulting population cumulants have an empty reverse-convention fiber under the same one-slot specification. [Theorem 3](#) then places vectors admitting both conventions on a codimension-one closure. This example illustrates the precise representation-level decision delivered by the theory.

The paper’s scope is population identification from finite cumulants within a fixed axis-normalized latent-source representation class. It establishes representation compatibility and separation for a maintained m , complementing finite-sample inference, model selection, and broader graphical identification analyses [[Pearl, 2009](#), [Spirtes et al., 2001](#), [Peters et al., 2017](#), [Richardson and Spirtes, 2002](#)]. This precise scope makes the causal interpretation conditional on the stated latent-source class.

The proofs use classical apolarity of binary forms and simultaneous decomposition ideas [[Comon and Murrain, 1996](#), [Landsberg, 2012](#)]. Cumulants turn independent-source representations into sums of powers of loading directions; apolar equations then recover, or partially constrain, the support of those directions. The verification appendix records the formal declarations corresponding to the delivered mathematical objects.

2 Setup and assumptions

2.1 The algebraic theorem domain

The main theorems concern two polynomial maps defined on fixed-source-count parameter spaces. Fix m and a truncation order L . The right-arrow parameter records the finite slopes $\gamma, \rho_1, \dots, \rho_m$ and order-specific source weights c_{jr} ; the left-arrow parameter records $\delta, \sigma_1, \dots, \sigma_m$ and d_{jr} . For each $2 \leq r \leq L$, the two maps send these coordinates to the mixed-cumulant block

$$t_{r,a} = \sum_{j=0}^{m+1} c_{jr} (u_{j1})^{r-a} (u_{j2})^a \quad \text{or} \quad t_{r,a} = \sum_{j=0}^{m+1} d_{jr} (v_{j1})^{r-a} (v_{j2})^a, \quad 0 \leq a \leq r.$$

The loading vectors u_j and v_j encode the two axis conventions. The theorem domain is the resulting pair of complex affine parameter spaces and polynomial cumulant maps, stated formally in [Definition 3–Definition 5](#). “Generic” in the complex results means outside a proper algebraic subset of these parameter spaces. The conclusion establishes emptiness of one map fiber within this domain.

2.2 Probabilistic interpretation

The probabilistic assumptions now attach an econometric interpretation to real points of the algebraic maps. Each representation has $m + 2$ source slots: two are tied to the coordinate-axis normalization and the remaining m middle source slots have finite slopes. Centered mutually independent sources make cumulants additive across slots, finite moments make the retained cumulants well-defined, and non-Gaussianity provides higher-order information. These assumptions define the real feasible regions below, while the complex-polynomial interpretation of the maps applies on the full affine parameter spaces. A middle source slot receives the latent-confounder interpretation when both coordinate loadings are nonzero, so “middle source slot” is the default term.

The source dimension is indexed by m , with n later denoting $m + 2$ loading directions. The working truncation order K is finite, and S_j denotes the j th centered latent source.

Assumption 1. `[ass:independent-sources]` (Independent latent sources) *The latent source variables $S_0, \dots, S_{m+1}$ are mutually independent with respect to μ :*

$$S_0 \perp \dots \perp S_{m+1}.$$

[Assumption 1](#) is the standard independent-component source condition in latent-variable LiNGAM; it is what makes joint cumulants additive across source slots [[Chen et al., 2025b](#)].

Assumption 2. `[ass:finite-cumulants]` *For every $j = 0, \dots, m+1$, the latent source variable S_j has finite K -th absolute moment:*

$$\mathbb{E}|S_j|^K < \infty.$$

[Assumption 2](#) is the standard finite-order cumulant existence condition, tailored to the truncation order used by the paper [[Chen et al., 2025b](#)].

Assumption 3. `[ass:source-nongaussianity]` *For every $j = 0, \dots, m+1$, the latent source variable S_j is non-Gaussian.*

[Assumption 3](#) is the usual non-Gaussian disturbance condition in LiNGAM-type identification; higher-order non-Gaussian structure supplies the cumulant information beyond covariance restrictions [Chen et al., 2025b].

The two axis-normalized parametrizations give the real structural interpretation of the arrow tags used by the polynomial maps. In the right-arrow convention, X is the first coordinate and Y is the second coordinate, the direct loading has slope γ , the middle finite slopes are ρ_i , and the forward loading vectors are u_j . In the left-arrow convention, the analogous direct and middle slopes are δ and σ_i , with reverse loading vectors v_j .

Assumption 4. `[ass:forward-axis-model]` *With forward loading vectors $u_0 = (1, \gamma)$, $u_j = (1, \rho_j)$ for $1 \leq j \leq m$, and $u_{m+1} = (0, 1)$, the bivariate observed outcome vector admits the forward latent linear representation*

$$(X, Y)^\top = \sum_{j=0}^{m+1} u_j S_j.$$

[Assumption 4](#) is a bivariate linear acyclic representation with latent source slots in the $X \rightarrow Y$ orientation [Chen et al., 2025b]. The axis normalization fixes one horizontal-axis and one vertical-axis slot, so only the middle source slots carry free finite slopes.

Assumption 5. `[ass:reverse-axis-model]` *With reverse loading vectors $v_0 = (1, 0)$, $v_j = (\sigma_j, 1)$ for $1 \leq j \leq m$, and $v_{m+1} = (\delta, 1)$, the bivariate observed outcome vector admits the reverse latent linear representation*

$$(X, Y)^\top = \sum_{j=0}^{m+1} v_j S_j.$$

[Assumption 5](#) is the corresponding reversed bivariate linear acyclic parametrization for the $Y \rightarrow X$ orientation [Chen et al., 2025b]. Its symmetric normalization permits a direct comparison of the two representation fibers.

Genericity enters through distinct finite loading directions and nonzero direct edges. These restrictions remove lower-dimensional coincidences in the loading support while permitting middle source slots with a zero coordinate loading.

Assumption 6. `[ass:forward-noncollinearity]` *The direct and latent finite slopes in the forward parametrization are pairwise distinct:*

$$|\{\gamma, \rho_i : 1 \leq i \leq m\}| = m + 1.$$

[Assumption 6](#) is the standard generic distinct-loading-directions condition for the forward parametrization [Schkoda et al., 2024]. It excludes merged finite slopes while permitting zero-valued middle slopes.

Assumption 7. `[ass:reverse-noncollinearity]` *The reverse slopes are pairwise distinct:*

$$|\{\delta, \sigma_i : 1 \leq i \leq m\}| = m + 1.$$

[Assumption 7](#) imposes the same generic distinct-loading-directions condition on the reverse parametrization [Schkoda et al., 2024]. The symmetry is useful because later separation statements compare full opposite-arrow fibers.

Assumption 8. `[ass:forward-nonzero-edge]` *The forward direct slope is nonzero:*

$$\gamma \neq 0.$$

[Assumption 8](#) is the standard nonzero direct-edge genericity restriction in the forward direction [\[Chen et al., 2025b\]](#). It excludes the degenerate case in which the distinguished direct slot lies on the horizontal axis.

Assumption 9. `[ass:reverse-nonzero-edge]` *The reverse direct slope is nonzero:*

$$\delta \neq 0.$$

[Assumption 9](#) is the analogous nonzero direct-edge genericity restriction for the reverse direction [\[Chen et al., 2025b\]](#). Together, the two nonzero-edge restrictions keep the two arrow conventions from sharing their distinguished direct-axis slot.

We now define the two law classes. The notation $\text{Laws}(\mathbb{R}^2)$ denotes the ambient class of Borel probability laws on $\mathbb{R}^2$, and the observational law setup is the data-generating world in which these law classes live.

Definition 1. `[def:forward-lvlingam-class]` (Forward class $\mathcal{P}_{m,K}^{\text{right}}$) *Write $\text{Laws}(\mathbb{R}^2)$ for the Borel probability laws on $\mathbb{R}^2$. Let $\mathcal{P}_{m,2m+2}^{\text{right}}$ be the set of all probability laws $P \in \text{Laws}(\mathbb{R}^2)$ for which there exist a measurable space $(\Omega, \mathcal{A})$, a probability measure μ on Ω , source functions*

$$S_0, \dots, S_{m+1} : \Omega \rightarrow \mathbb{R},$$

observed variables $X, Y : \Omega \rightarrow \mathbb{R}$, a direct forward slope $\gamma \in \mathbb{R}$, and latent forward slopes $\rho_1, \dots, \rho_m \in \mathbb{R}$ such that:

- *(Source independence.) The sources $S_0, \dots, S_{m+1}$ are mutually independent under μ .*
- *(Finite cumulant order.) Each source has finite moments through order $2m + 2$.*
- *(Source non-Gaussianity.) Each source has non-Gaussian law under μ .*
- *(Centering.) Each source is centered: $\int S_j d\mu = 0$ for every $j = 0, \dots, m + 1$.*
- *(Forward representation.) For every $\omega \in \Omega$, the bivariate observed outcome vector $(X(\omega), Y(\omega))^\top$ admits the forward $X \rightarrow Y$ latent linear non-Gaussian representation with loading vectors u_j .*
- *(Forward noncollinearity.) The forward direct and latent finite slopes $\gamma, \rho_1, \dots, \rho_m$ are pairwise distinct.*
- *(Nonzero forward edge.) The forward direct slope satisfies $\gamma \neq 0$.*
- *(Observed law.) The law P is the pushforward μ -law of (X, Y) , i.e. $P = \mu \circ (X, Y)^{-1}$.*

[Definition 1](#) is a representation-level class: membership means that the law has at least one centered non-Gaussian source representation satisfying the right-arrow axis convention.

Definition 2. [def:reverse-lvlingam-class] (Reverse class $\mathcal{P}_{m,K}^{\text{left}}$) Using the ambient law class $\text{Laws}(\mathbb{R}^2)$ defined in [Definition 1](#), let $\mathcal{P}_{m,2m+2}^{\text{left}}$ be the set of all probability laws $P \in \text{Laws}(\mathbb{R}^2)$ for which there exist a probability space $(\Omega, \mathcal{A}, \mu)$, latent source variables

$$S_0, \dots, S_{m+1} : \Omega \rightarrow \mathbb{R},$$

observed variables $X, Y : \Omega \rightarrow \mathbb{R}$, a reverse direct slope $\delta \in \mathbb{R}$, and reverse latent slopes $\sigma_1, \dots, \sigma_m \in \mathbb{R}$, such that:

- **(Source independence.)** The sources $S_0, \dots, S_{m+1}$ are mutually independent under μ .
- **(Finite cumulant order.)** Each source belongs to $L^{2m+2}(\mu)$.
- **(Source non-Gaussianity.)** Each source law $\mu \circ S_j^{-1}$ is non-Gaussian.
- **(Centering.)** Each source is centered: $\int_{\Omega} S_j d\mu = 0$ for all $j = 0, \dots, m+1$.
- **(Reverse representation.)** The bivariate observed outcome vector $(X, Y)^{\top}$ admits the reverse $Y \rightarrow X$ latent linear non-Gaussian representation with loading vectors v_j , pointwise on Ω .
- **(Reverse noncollinearity.)** The reverse direct and latent finite slopes $\delta, \sigma_1, \dots, \sigma_m$ are pairwise distinct.
- **(Nonzero reverse edge.)** The reverse direct slope satisfies $\delta \neq 0$.
- **(Observed law.)** The law P is the pushforward of μ by $\omega \mapsto (X(\omega), Y(\omega))$.

[Definition 2](#) gives the opposite representation class under the same source-count convention. The paper's separation statements characterize when cumulants from one parametrization have an empty fiber under the other cumulant map.

The identifying information is expressed through joint cumulants. For a law P , Cum denotes the joint cumulant functional, $\kappa_{r,a}(P)$ is the coordinate with $r-a$ copies of X and a copies of Y , and $T_L(P)$ collects all cumulants from orders 2 through L . The dimension of this vector is q_L . Classical references for cumulant notation and higher-order moment methods include [Brillinger \[1969\]](#) and [McCullagh \[1987\]](#).

Definition 3. [def:truncated-cumulant] (Truncated cumulant vector $T_L(P)$) For the ambient law class of [Definition 1](#), write $\text{Cum}_P(Z_1, \dots, Z_r)$ for the joint cumulant under P , and define $\text{Cum}_r(Z)$ to be the joint cumulant of r copies of Z . For $P \in \text{Laws}(\mathbb{R}^2)$, define

$$T_L(P) = (\kappa_{r,a}(P) : 2 \leq r \leq L, 0 \leq a \leq r),$$

where

$$\kappa_{r,a}(P) = \text{Cum}_P(\underbrace{X, \dots, X}_{r-a \text{ copies}}, \underbrace{Y, \dots, Y}_{a \text{ copies}}).$$

The coordinate dimension is

$$q_L = \frac{L(L+3)}{2} - 2.$$

For later cumulant maps, define the complex parameter spaces

$$\Theta_{m,L}^{\text{right}} = \mathbb{C}^{m+1+(m+2)(L-1)}$$

with coordinates $(\gamma, \rho_1, \dots, \rho_m, (c_{jr})_{0 \leq j \leq m+1, 2 \leq r \leq L})$, and

$$\Theta_{m,L}^{\text{left}} = \mathbb{C}^{m+1+(m+2)(L-1)}$$

with coordinates $(\delta, \sigma_1, \dots, \sigma_m, (d_{jr})_{0 \leq j \leq m+1, 2 \leq r \leq L})$.

Definition 3 also introduces the complex affine parameter spaces $\Theta_{m,L}^{\text{right}}$ and $\Theta_{m,L}^{\text{left}}$. These complexified coordinates, summarized by the complexified structural-parameter and cumulant-coordinate world, are algebraic devices for studying polynomial images; real feasibility is imposed separately below.

By independence and multilinearity of cumulants, each source contributes its order- r cumulant weight times the appropriate monomial in its loading vector. We write c_{jr} and d_{jr} for these source-cumulant weights, and use the polynomial maps $\Phi_{m,L}^{\text{right}}$ and $\Phi_{m,L}^{\text{left}}$ to encode the resulting cumulant vectors.

Definition 4. `[def:forward-cumulant-map]` (Forward cumulant map $\Phi_{m,L}^{\text{right}}$) *Using the complex parameter space of Definition 3, The forward cumulant map $\Phi_{m,L}^{\text{right}} : \Theta_{m,L}^{\text{right}} \rightarrow \mathbb{C}^{q_L}$ is defined coordinatewise by*

$$[\Phi_{m,L}^{\text{right}}(\gamma, \rho, c)]_{r,a} = \sum_{j=0}^{m+1} c_{jr} (u_{j1})^{r-a} (u_{j2})^a, \quad 2 \leq r \leq L, \quad 0 \leq a \leq r,$$

where

$$u_0 = (1, \gamma), \quad u_j = (1, \rho_j) \text{ for } 1 \leq j \leq m, \quad u_{m+1} = (0, 1).$$

The map in Definition 4 is the cumulant analogue of a finite simultaneous binary-form decomposition. At this stage its source weights are free algebraic coordinates, with real-random-variable realizability imposed later through the feasible regions.

Definition 5. `[def:reverse-cumulant-map]` (Reverse cumulant map $\Phi_{m,L}^{\text{left}}$) *Using the complex parameter space of Definition 3, The reverse cumulant map $\Phi_{m,L}^{\text{left}} : \Theta_{m,L}^{\text{left}} \rightarrow \mathbb{C}^{q_L}$ is defined coordinatewise by*

$$[\Phi_{m,L}^{\text{left}}(\delta, \sigma, d)]_{r,a} = \sum_{j=0}^{m+1} d_{jr} (v_{j1})^{r-a} (v_{j2})^a, \quad 2 \leq r \leq L, \quad 0 \leq a \leq r,$$

where

$$v_0 = (1, 0), \quad v_j = (\sigma_j, 1) \text{ for } 1 \leq j \leq m, \quad v_{m+1} = (\delta, 1).$$

Definition 5 mirrors Definition 4 after exchanging the axis convention. This symmetry lets the paper characterize possible simultaneous membership in both parametrized images.

The algebraic images are taken after Zariski closure in the ambient complex affine cumulant space. We write `ZariskiClosure` for this closure operation, and denote the two closed image varieties by $C_{m,L}^{\text{right}}$ and $C_{m,L}^{\text{left}}$.

Definition 6. `[def:image-varieties]` (Image varieties $C_{m,L}^{\text{right}}$ and $C_{m,L}^{\text{left}}$) *For a subset of the ambient complex affine space, write `ZariskiClosure` for its Zariski closure. Using the parameter spaces and maps of Definition 4 and Definition 5, the forward and reverse image varieties are*

$$C_{m,L}^{\text{right}} = \text{ZariskiClosure}(\Phi_{m,L}^{\text{right}}(\Theta_{m,L}^{\text{right}})), \quad C_{m,L}^{\text{left}} = \text{ZariskiClosure}(\Phi_{m,L}^{\text{left}}(\Theta_{m,L}^{\text{left}})),$$

as subvarieties of $\mathbb{C}^{q_L}$.

These varieties record the algebraic restrictions imposed by a fixed number of source slots and a fixed arrow convention. Their use is standard in algebraic treatments of tensor and binary-form decompositions, where generic conclusions are naturally stated off proper algebraic subsets [Comon and Mourrain, 1996, Landsberg, 2012].

The next definition isolates the open parameter loci on which slopes are distinct, direct slopes are nonzero, and all retained source cumulant weights are nonzero. We denote these loci by $\Theta_{m,L}^{\text{right},\circ}$ and $\Theta_{m,L}^{\text{left},\circ}$.

Definition 7. [def:generic-parameter-loci] (Generic loci $\Theta_{m,L}^{\text{right},\circ}$ and $\Theta_{m,L}^{\text{left},\circ}$) *For every $L \geq 2$, using the parameter spaces of Definition 4 and Definition 5, the forward generic parameter locus is*

$$\Theta_{m,L}^{\text{right},\circ} = \left\{ (\gamma, \rho, c) \in \Theta_{m,L}^{\text{right}} : \gamma \prod_{i=1}^m (\gamma - \rho_i) \prod_{1 \leq i < \ell \leq m} (\rho_i - \rho_\ell) \prod_{j=0}^{m+1} \prod_{r=2}^L c_{jr} \neq 0 \right\}.$$

The reverse generic parameter locus is

$$\Theta_{m,L}^{\text{left},\circ} = \left\{ (\delta, \sigma, d) \in \Theta_{m,L}^{\text{left}} : \delta \prod_{i=1}^m (\delta - \sigma_i) \prod_{1 \leq i < \ell \leq m} (\sigma_i - \sigma_\ell) \prod_{j=0}^{m+1} \prod_{r=2}^L d_{jr} \neq 0 \right\}.$$

Definition 7 strengthens the assumption-level genericity by also requiring nonzero retained source cumulants. This representation-level restriction selects the finite cumulant coordinates used by the generic theorem.

For a cumulant coordinate vector t , the relevant inverse images are the parameter fibers under the two polynomial maps. We use $R_{m,L}^{\text{right}}(t)$ and $R_{m,L}^{\text{left}}(t)$ for these fibers.

Definition 8. [def:fiber-correspondences] *For a cumulant vector t , the forward and reverse fibers are*

$$R_{m,L}^{\text{right}}(t) = \left\{ \theta \in \Theta_{m,L}^{\text{right}} : \Phi_{m,L}^{\text{right}}(\theta) = t \right\}, \quad R_{m,L}^{\text{left}}(t) = \left\{ \eta \in \Theta_{m,L}^{\text{left}} : \Phi_{m,L}^{\text{left}}(\eta) = t \right\}.$$

Definition 8 separates two questions: whether a cumulant vector lies in an arrow image, and how many parameter points produce it. The same-arrow fiber facts in the next section therefore have a different status from opposite-arrow exclusion.

Source labels in the middle block are exchangeable. The admissible relabelling group G_m acts by a permutation π of the middle labels while fixing the two axis slots.

Definition 9. [def:admissible-source-swaps] *Let G_m be the permutation group of the middle source labels $\{1, \dots, m\}$, acting while fixing labels 0 and $m+1$. For $\pi \in G_m$, the forward action on parameters is*

$$\pi \cdot (\gamma, (\rho_i)_{i=1}^m, (c_{jr})_{j,r}) = (\gamma, (\rho_{\pi(i)})_{i=1}^m, (c'_{jr})_{j,r}),$$

where

$$c'_{ir} = c_{\pi(i),r} \quad (1 \leq i \leq m), \quad c'_{0r} = c_{0r}, \quad c'_{m+1,r} = c_{m+1,r}.$$

The left action is the analogous action with γ, ρ, c replaced by δ, σ, d . Thus the source labels 0 and $m+1$ are fixed, while the labels $1, \dots, m$ are permuted.

Definition 9 specifies the label symmetry quotiented in the paper. It fixes the direct-axis and fixed-axis source slots and permits permutations among the middle latent-source slots.

2.3 The real-feasibility bridge

The algebraic parameter spaces allow arbitrary complex source-cumulant weights. For econometric interpretation, we also need real parameter points whose source cumulants are realized by centered non-Gaussian real random variables through the retained order. These real feasible regions are denoted $F_{m,L}^{\text{right}}$ and $F_{m,L}^{\text{left}}$. They bridge the complex theorem layer to real source models and retain the theorem's parameter-space interpretation.

Definition 10. [def:real-feasible-regions] (Real feasible regions $F_{m,L}^{\text{right}}$ and $F_{m,L}^{\text{left}}$) *Using the cumulant convention of Definition 3, the forward real feasible region $F_{m,L}^{\text{right}}$ is the set of all real tuples $(\gamma, \rho, c) \in \mathbb{R}^{m+1} \times \mathbb{R}^{(m+2)(L-1)}$ such that $\gamma \neq 0$, the slopes $\rho_1, \dots, \rho_m, \gamma$ are pairwise distinct, and, for every $0 \leq j \leq m+1$, there exists a centered non-Gaussian real random variable S_j with $\mathbb{E}|S_j|^L < \infty$ satisfying*

$$\text{Cum}_r(S_j) = c_{jr}, \quad 2 \leq r \leq L.$$

The reverse real feasible region $F_{m,L}^{\text{left}}$ is defined analogously, with (δ, σ, d) in place of (γ, ρ, c) and with $\delta \neq 0$.

Definition 10 identifies the finite real cumulant coordinates realized by actual centered non-Gaussian source laws.

We summarize the amount of cumulant information needed for generic real arrow separation by the information order $K^*(m)$. The arrow index b ranges over the two conventions, and $\text{opposite}(b)$ denotes the opposite-arrow involution.

Definition 11. [def:information-order] (Information order $K^*(m)$) *For $b \in \{\text{right}, \text{left}\}$, use the branch convention $\Phi_{m,L}^b = \Phi_{m,L}^{\text{right}}$ and $F_{m,L}^b = F_{m,L}^{\text{right}}$ when $b = \text{right}$, and $\Phi_{m,L}^b = \Phi_{m,L}^{\text{left}}$ and $F_{m,L}^b = F_{m,L}^{\text{left}}$ when $b = \text{left}$, using Definition 4, Definition 5, and Definition 10. Define $\text{opposite}(\text{right}) = \text{left}$ and $\text{opposite}(\text{left}) = \text{right}$. Define $K^*(m)$ to be the minimum integer $L \geq 2$, if such an integer exists, such that for each $b \in \{\text{right}, \text{left}\}$, outside a proper real algebraic subset of $F_{m,L}^b$, the cumulant vector*

$$\Phi_{m,L}^b(\theta)$$

has no representation in

$$F_{m,L}^{\text{opposite}(b)}.$$

If no such L exists, set $K^(m) = \infty$.*

Definition 11 introduces notation for the generic separation threshold. The value $K^*(m) = 2m + 2$ serves as a conjectural benchmark, and the later information-order theorem establishes the sharper value in the relevant range.

Finally, the main theorem uses the classical apolar language for binary forms. For a finite projective support D , the polynomial Q_D annihilates that support, and $q(\partial)$ denotes the associated constant-coefficient differential operator. This notation is standard in binary-form decomposition and Waring-rank geometry [Comon and Murrain, 1996, Landsberg, 2012].

Definition 12. [def:apolar-notation] (Apolar operators Q_D and $q(\partial)$) *For a finite projective loading-direction support D , define Q_D to be its squarefree homogeneous support-annihilator polynomial, with the normalization used by the cumulant-map coordinates. For a homogeneous*

binary form $q(x, y)$, define $q(\partial)$ to be the constant-coefficient differential operator obtained by substituting ∂_x, ∂_y for x, y . For a real parameter set A , define $A_{\mathbb{C}}$ to be its coordinatewise complexification.

Definition 12 supplies the notation used to recover unordered loading directions from truncated higher-order population cumulants.

3 Main results

The setup in **Definition 12** turns the direction problem into a question about finite collections of binary forms. At order $K = 2m + 2$, the cumulant coordinates determine a simultaneous decomposition with $m + 2$ loading directions. The first result states that, on a relatively Zariski-open dense set of parameters, the apolar equations recover the unordered loading support and rule out a full representation in the opposite arrow convention. This representation-level separation statement concerns the fibers of the polynomial cumulant maps in **Definition 4** and **Definition 5**.

Theorem 1. [thm:generic-apolar-arrow-recovery] *Let $m \geq 1$, put $n = m + 2$ and $K = 2m + 2 = 2n - 2$. Use the support-annihilator and differential-operator conventions of **Definition 12**. There exist sets*

$$U_m^{\text{right}} \subseteq \Theta_{m,K}^{\text{right},\circ}, \quad U_m^{\text{left}} \subseteq \Theta_{m,K}^{\text{left},\circ}$$

such that:

- **(Generic loci.)** *Each of U_m^{right} and U_m^{left} is relatively Zariski open and relatively Zariski dense in the corresponding complex parameter space $\Theta_{m,K}^{\text{right}}$ or $\Theta_{m,K}^{\text{left}}$ of **Definition 3**, and is contained in the corresponding generic parameter locus of **Definition 7**.*
- **(Open neighborhoods meeting feasibility.)** *There are Euclidean-open sets $O^{\text{right}}, O^{\text{left}}$ of real parameter space such that*

$$O^{\text{right}} \cap F_{m,K}^{\text{right}} \neq \emptyset, \quad O^{\text{left}} \cap F_{m,K}^{\text{left}} \neq \emptyset,$$

where the feasible regions are those of **Definition 10**; use the coordinatewise-complexification convention of **Definition 12**. Then

$$\left(O^{\text{right}} \cap F_{m,K}^{\text{right}}\right)_{\mathbb{C}} \subseteq U_m^{\text{right}}, \quad \left(O^{\text{left}} \cap F_{m,K}^{\text{left}}\right)_{\mathbb{C}} \subseteq U_m^{\text{left}}.$$

- **(Forward recovery and separation.)** *For every $\theta \in U_m^{\text{right}}$,*

$$R_{m,K}^{\text{left}}\left(\Phi_{m,K}^{\text{right}}(\theta)\right) = \emptyset.$$

Moreover, there is a nonzero squarefree polynomial $Q \in \mathbb{C}[z]$ such that $Q(0) \neq 0$ and the multiset of roots of Q is $\{\gamma, \rho_1, \dots, \rho_m\}$. Every generic forward parameter θ'' satisfying

$$\Phi_{m,K}^{\text{right}}(\theta'') = \Phi_{m,K}^{\text{right}}(\theta)$$

has the same multiset $\{\gamma, \rho_1, \dots, \rho_m\}$. If

$$f_r(x, y) = \sum_{a=0}^r \binom{r}{a} t_{r,a} x^{r-a} y^a, \quad t = \Phi_{m,K}^{\text{right}}(\theta),$$

let D be the forward loading-direction support. Then

$$Q_D \neq 0, \quad x \mid Q_D, \quad y \nmid Q_D,$$

and, for every homogeneous binary form q of degree n ,

$$[q(\partial)f_{n+k} = 0 \text{ for every } 0 \leq k \leq m] \iff \exists c \in \mathbb{C} : q = cQ_D.$$

- **(Reverse recovery and separation.)** For every $\eta \in U_m^{\text{left}}$,

$$R_{m,K}^{\text{right}}(\Phi_{m,K}^{\text{left}}(\eta)) = \emptyset.$$

Moreover, there is a nonzero squarefree polynomial $Q \in \mathbb{C}[z]$ such that $Q(0) \neq 0$ and the multiset of roots of Q is $\{\delta, \sigma_1, \dots, \sigma_m\}$. Every generic reverse parameter η'' satisfying

$$\Phi_{m,K}^{\text{left}}(\eta'') = \Phi_{m,K}^{\text{left}}(\eta)$$

has the same multiset $\{\delta, \sigma_1, \dots, \sigma_m\}$. If

$$f_r(x, y) = \sum_{a=0}^r \binom{r}{a} t_{r,a} x^{r-a} y^a, \quad t = \Phi_{m,K}^{\text{left}}(\eta),$$

let D be the reverse loading-direction support. Then

$$Q_D \neq 0, \quad y \mid Q_D, \quad x \nmid Q_D,$$

and, for every homogeneous binary form q of degree n ,

$$[q(\partial)f_{n+k} = 0 \text{ for every } 0 \leq k \leq m] \iff \exists c \in \mathbb{C} : q = cQ_D.$$

The feasible-incidence clause in [Theorem 1](#) supplies real feasible parameter points inside the open patches to which the complex-algebraic conclusion applies after coordinatewise complexification. Its topological content is intersection with the feasible regions $F_{m,K}^{\text{right}}$ and $F_{m,K}^{\text{left}}$, giving a precise real-parameter interpretation of the generic result.

The recovery statement is closely related to classical apolarity for binary forms and simultaneous tensor decompositions [[Comon and Mourrain, 1996](#), [Landsberg, 2012](#), [Kolda and Bader, 2009](#)]. Here the distinctive identification point is that the recovered support contains one of the two fixed axes in a way that is incompatible with the opposite arrow on the generic loci. Thus, within the fixed source-count, axis-normalized representation class, the higher-order cumulants jointly recover loading directions and separate the two truncated-cumulant fibers.

The next result isolates the same-arrow consequences of [Theorem 1](#) on its loci U_m^{right} and U_m^{left} .

Theorem 2. [thm:generic-arrow-recovery-and-fiber-obstruction] (Generic Fiber Obstruction) For every $m \in \mathbb{N}$ with $m \geq 1$, set $K = 2m + 2$. There exist complex loci $U_m^{\text{right}} \subseteq \Theta_{m,K}^{\text{right},\circ}$ and $U_m^{\text{left}} \subseteq \Theta_{m,K}^{\text{left},\circ}$, with the generic loci as in [Definition 7](#), such that:

- **(Relative genericity.)** Each of U_m^{right} and U_m^{left} is relatively Zariski open and relatively Zariski dense in the parameter spaces $\Theta_{m,K}^{\text{right}}$ and $\Theta_{m,K}^{\text{left}}$ of [Definition 4](#) and [Definition 5](#), respectively.

- **(Real incidence.)** There is a Euclidean-open set of real forward parameters whose intersection with $F_{m,K}^{\text{right}}$ from [Definition 10](#) is nonempty, and every parameter in this intersection has coordinatewise complexification contained in U_m^{right} . Likewise, there is a Euclidean-open set of real reverse parameters whose intersection with $F_{m,K}^{\text{left}}$ is nonempty, and every parameter in this intersection has coordinatewise complexification contained in U_m^{left} .
- **(Forward generic fibers.)** For every $\theta = (\gamma, \rho, c) \in U_m^{\text{right}}$, every forward parameter in the forward fiber

$$R_{m,K}^{\text{right}}(\Phi_{m,K}^{\text{right}}(\theta))$$

of [Definition 8](#) and [Definition 4](#) has the same unordered loading-slope multiset

$$\{\gamma, \rho_1, \dots, \rho_m\}$$

as θ . The opposite-arrow fiber

$$R_{m,K}^{\text{left}}(\Phi_{m,K}^{\text{right}}(\theta))$$

of [Definition 8](#) (formed with the reverse map of [Definition 5](#)) is empty. For every latent index $i = 1, \dots, m$, there is a generic forward parameter $\theta' \in \Theta_{m,K}^{\text{right},o}$ obtained by swapping the full pairs $(\gamma, (c_{0r})_{r=2}^K)$ and $(\rho_i, (c_{ir})_{r=2}^K)$, leaving every other coordinate fixed, such that

$$\Phi_{m,K}^{\text{right}}(\theta') = \Phi_{m,K}^{\text{right}}(\theta),$$

but θ' is not in the admissible G_m -orbit of θ from [Definition 9](#). If $m \geq 2$, this forward fiber has exact relative Zariski dimension

$$\frac{m(m-1)}{2}.$$

- **(Reverse generic fibers.)** For every $\eta = (\delta, \sigma, d) \in U_m^{\text{left}}$, every reverse parameter in the reverse fiber

$$R_{m,K}^{\text{left}}(\Phi_{m,K}^{\text{left}}(\eta))$$

has the same unordered loading-slope multiset

$$\{\delta, \sigma_1, \dots, \sigma_m\}$$

as η . The opposite-arrow fiber

$$R_{m,K}^{\text{right}}(\Phi_{m,K}^{\text{left}}(\eta))$$

is empty. For every latent index $i = 1, \dots, m$, there is a generic reverse parameter $\eta' \in \Theta_{m,K}^{\text{left},o}$ obtained by swapping the full pairs $(\delta, (d_{m+1,r})_{r=2}^K)$ and $(\sigma_i, (d_{ir})_{r=2}^K)$, leaving every other coordinate fixed, such that

$$\Phi_{m,K}^{\text{left}}(\eta') = \Phi_{m,K}^{\text{left}}(\eta),$$

but η' is not in the admissible G_m -orbit of η from [Definition 9](#). If $m \geq 2$, this reverse fiber has exact relative Zariski dimension

$$\frac{m(m-1)}{2}.$$

The non-orbit clause in [Theorem 2](#) is important for interpreting same-arrow identification. The admissible relabelling group fixes the axis slots, while the cumulant equations can still admit generic same-arrow parameter points obtained by exchanging a direct slot with a middle slot. Consequently, the theorem combines generic arrow separation with an exact witness of residual structural ambiguity within the chosen arrow convention.

We next define the exceptional compatibility locus. It consists of cumulant vectors that have a generic representation for one arrow and at least one full representation for the other. Its formulation in cumulant space makes the exceptional set a property of the truncated cumulant vector, independent of a particular chosen parametrization.

Definition 13. `[def:generic-full-fiber-compatibility]` (Generic full-fiber compatibility E_m) *Set $K = 2m + 2$. Define*

$$E_m = \left\{ t \in \mathbb{C}^{q_K} : (R_{m,K}^{\text{right}}(t) \cap \Theta_{m,K}^{\text{right},\circ} \neq \emptyset \text{ and } R_{m,K}^{\text{left}}(t) \neq \emptyset) \text{ or } (R_{m,K}^{\text{left}}(t) \cap \Theta_{m,K}^{\text{left},\circ} \neq \emptyset \text{ and } R_{m,K}^{\text{right}}(t) \neq \emptyset) \right\}.$$

With ZariskiClosure as defined in [Definition 6](#), let

$$\overline{E}_m = \text{ZariskiClosure}(E_m).$$

The corresponding exceptional preimages are

$$H_m^{\text{right}} = \Theta_{m,K}^{\text{right},\circ} \cap (\Phi_{m,K}^{\text{right}})^{-1}(E_m), \quad H_m^{\text{left}} = \Theta_{m,K}^{\text{left},\circ} \cap (\Phi_{m,K}^{\text{left}})^{-1}(E_m).$$

For the smallest source counts, the compatibility equations can be written explicitly. These displayed systems give finite incidence descriptions of the algebraic event in which the two arrow parametrizations agree on the retained cumulants.

Definition 14. `[def:worked-compatibility-instances]` *For $m = 1$ and $K = 4$, A_1 is the full solution set, in the displayed forward coordinates, reverse coordinates, and cumulant coordinates t , of the 12-equation incidence system*

$$t_{r,a} = c_{0r}\gamma^a + c_{1r}\rho^a + c_{2r}\mathbb{K}_{\{a=r\}} = d_{0r}\mathbb{K}_{\{a=0\}} + d_{1r}\sigma^{r-a} + d_{2r}\delta^{r-a}, \quad r = 2, 3, 4, \quad 0 \leq a \leq r,$$

subject to the disjunction that the forward or reverse coordinates lie in the corresponding generic locus of [Definition 7](#). For $p \in A_1$, write p^{right} for its forward parameter coordinates. Its common-axis subfamily is given by

$$\rho = \sigma = 0, \quad \delta\gamma = 1, \quad d_{0r} = c_{1r}, \quad d_{1r} = c_{2r}, \quad d_{2r} = c_{0r}\gamma^r.$$

For $m = 2$ and $K = 6$, A_2 is the full solution set, in the displayed forward coordinates, reverse coordinates, and cumulant coordinates t , of the 25-equation incidence system

$$t_{r,a} = c_{0r}\gamma^a + c_{1r}\rho_1^a + c_{2r}\rho_2^a + c_{3r}\mathbb{K}_{\{a=r\}} = d_{0r}\mathbb{K}_{\{a=0\}} + d_{1r}\sigma_1^{r-a} + d_{2r}\sigma_2^{r-a} + d_{3r}\delta^{r-a}, \quad r = 2, \dots, 6, \quad 0 \leq a \leq r,$$

subject to the same generic-locus disjunction. For $p \in A_2$, p^{right} again denotes its forward parameter coordinates. Its common-axis subfamily is given by

$$\rho_1 = \sigma_1 = 0, \quad \delta\gamma = 1, \quad \sigma_2\rho_2 = 1, \quad d_{0r} = c_{1r}, \quad d_{1r} = c_{3r}, \quad d_{2r} = c_{2r}\rho_2^r, \quad d_{3r} = c_{0r}\gamma^r.$$

The codimension calculation compares simultaneous decompositions after removing only the admissible middle-source label symmetry. This quotient-fiber comparison is the algebraic object behind the exceptional-locus statement.

Definition 15. [def:separation-handle] (Quotient-fiber comparison) *The separation handle is the algebraic comparison of simultaneous binary-form decompositions*

$$\Phi_{m,K}^{\text{right}}(\theta) = \Phi_{m,K}^{\text{left}}(\eta)$$

after quotienting each same-arrow fiber by the admissible source-label permutation group G_m of Definition 9. The comparison uses the axis directions fixed in the two cumulant maps and the Jacobian ranks of the resulting quotient-fiber equations.

The next theorem says that the compatibility locus is exceptional and has codimension exactly one in each arrow image variety at the apolar order. For $m \geq 2$, this single-equation geometry is substantially larger than a count based on the m middle sources might suggest.

Theorem 3. [thm:exceptional-locus-codimension-one] *For every $m \geq 1$, at $K = 2m + 2$, let $\bar{E}_m$, H_m^{right} , and H_m^{left} be as in Definition 13, and let $C_{m,K}^{\text{right}}$ and $C_{m,K}^{\text{left}}$ be as in Definition 6. Then:*

- **(Codimension.)** *The exceptional closure $\bar{E}_m$ has codimension exactly 1 in each of $C_{m,K}^{\text{right}}$ and $C_{m,K}^{\text{left}}$. Moreover, if $m \geq 2$, then $\bar{E}_m$ does not have codimension m in either $C_{m,K}^{\text{right}}$ or $C_{m,K}^{\text{left}}$.*
- **(Forward preimage.)** *For every $\theta \in \Theta_{m,K}^{\text{right}}$,*

$$\theta \in H_m^{\text{right}} \iff \theta \in \Theta_{m,K}^{\text{right},\circ} \text{ and } R_{m,K}^{\text{left}}\left(\Phi_{m,K}^{\text{right}}(\theta)\right) \neq \emptyset.$$

- **(Reverse preimage.)** *For every $\eta \in \Theta_{m,K}^{\text{left}}$,*

$$\eta \in H_m^{\text{left}} \iff \eta \in \Theta_{m,K}^{\text{left},\circ} \text{ and } R_{m,K}^{\text{right}}\left(\Phi_{m,K}^{\text{left}}(\eta)\right) \neq \emptyset.$$

- **(Low-dimensional specializations.)** *If $m = 1$, then*

$$E_1 = \left\{ t : \exists p \text{ satisfying the system } A_1, \Phi_{1,4}^{\text{right}}(p^{\text{right}}) = t \right\}.$$

If $m = 2$, then

$$E_2 = \left\{ t : \exists p \text{ satisfying the system } A_2, \Phi_{2,6}^{\text{right}}(p^{\text{right}}) = t \right\},$$

where A_1, A_2 are the compatibility instances of Definition 14.

Theorem 3 gives the proper interpretation of the exceptional set. Generic separation in Theorem 1 is compatible with a codimension-one family of cumulant vectors that admit both arrow conventions. The statement also identifies the preimages H_m^{right} and H_m^{left} exactly as generic same-arrow parameters whose cumulant vector retains a full opposite-arrow fiber.

The preceding results use order $2m + 2$, the natural apolar order for recovering $m + 2$ loading directions. The real information-order question asks whether this order is minimal for generic real arrow separation. The following handle records the relevant qualification: through order $2m + 1$, real twin parameters with matching truncated cumulants exist.

Definition 16. [def:real-twin-construction-handle] *The real twin construction handle is the proposition that there exist real forward and reverse parameters θ and η , each realized through order*

$$K_- = 2m + 1,$$

by centered, non-Gaussian, compactly supported real sources such that their truncated cumulant maps agree through order K_- :

$$\Phi_{m,K_-}^{\text{right}}(\theta) = \Phi_{m,K_-}^{\text{left}}(\eta).$$

This definition records only the existence proposition; it does not include a construction or proof.

For $m \geq 3$, however, the generic real separation threshold improves by one order relative to the apolar support-recovery order. The statement establishes generic separation over the real feasible regions of [Definition 10](#), while full loading-support recovery remains the target of the order- $2m + 2$ apolar theorem.

Theorem 4. [thm:improved-real-information-order] (Improved Information Order) *The theorem-local symbols E^{right} and E^{left} below are real parameter-space exceptional subsets and are unrelated to the cumulant-space compatibility locus E_m . For every $m \in \mathbb{N}$ with $3 \leq m$, at truncation order $L = 2m + 1$:*

- **(Forward generic separation.)** *There is a proper real algebraic subset of the real forward parameter space whose complement meets $F_{m,L}^{\text{right}}$; choose one such subset and denote it locally in this theorem by E^{right} . For every $\theta \in F_{m,L}^{\text{right}} \setminus E^{\text{right}}$ there is no $\eta \in F_{m,L}^{\text{left}}$ such that*

$$\Phi_{m,L}^{\text{right}}(\theta) = \Phi_{m,L}^{\text{left}}(\eta).$$

- **(Reverse generic separation.)** *There is a proper real algebraic subset of the real reverse parameter space whose complement meets $F_{m,L}^{\text{left}}$; choose one such subset and denote it locally in this theorem by E^{left} . For every $\eta \in F_{m,L}^{\text{left}} \setminus E^{\text{left}}$ there is no $\theta \in F_{m,L}^{\text{right}}$ such that*

$$\Phi_{m,L}^{\text{left}}(\eta) = \Phi_{m,L}^{\text{right}}(\theta).$$

Consequently, for the information order $K^(m)$ of [Definition 11](#),*

$$K^*(m) \leq 2m + 1 \quad \text{and} \quad K^*(m) \neq 2m + 2.$$

[Theorem 4](#) establishes $K^*(m) \leq 2m + 1$ for $m \geq 3$ under the generic real separation criterion in [Definition 11](#). Taken together, [Theorem 1](#)–[Theorem 4](#) separate three issues: support recovery at the apolar order, same-arrow fiber nonuniqueness within a fixed axis convention, and the lower real truncation order sufficient for generic exclusion of the opposite arrow.

4 Discussion and extensions

4.1 Interpretation

The results in the preceding section identify a representation-level exclusion restriction in bivariate latent linear non-Gaussian models. [Theorem 1](#) excludes the opposite axis-normalized cumulant-map fiber on its stated generic loci at order $2m + 2$. Thus, if a structural model M

induces observational law P_M , the theorems characterize the relevant truncated cumulant-map fibers under the fixed source-count and axis-normalization conventions of [Definition 1](#) and [Definition 2](#).

This distinction matters for econometric interpretation. In a conventional nonparametric identification problem, one asks whether the observational law determines a structural object over a specified model class. Here the maintained class is more structured: independent non-Gaussian source slots, a fixed and externally specified number of slots, and two axis-normalized loading conventions. A natural user is therefore a researcher whose institutional or factor model fixes these inputs and who wants to know whether the competing recursive specification is compatible with the population cumulants. Within that class, higher-order cumulants impose restrictions beyond covariance information and can generically exclude the opposite truncated-cumulant fiber. Related recent work on algebraic and semialgebraic identification likewise distinguishes generic algebraic separation from full distributional identification [[Mesters and Zwiernik, 2022](#), [Virolainen, 2024](#)].

The codimension-one exceptional locus in [Theorem 3](#) is a compatibility set for truncated cumulant vectors. For cumulant vectors with a generic representation in one arrow image, separation holds off this Zariski-closed compatibility set, and the theorem gives its closure exactly codimension one inside each arrow image variety at the apolar order. Econometrically, opposite-arrow compatibility is therefore a genuine lower-dimensional restriction governed by a single algebraic dimension even when $m \geq 2$. The size of this set motivates explicit reporting in applications where the data-generating process may be close to a compatibility surface.

The same-arrow fiber statements in [Theorem 2](#) sharpen the interpretation of the information recovered. The unordered loading-slope support is recovered on the stated generic loci, while the chosen arrow retains parametrization ambiguity beyond the admissible middle-source relabellings of [Definition 9](#). In particular, exchanging a direct slot with a middle slot can preserve the cumulant vector while falling outside the admissible G_m -orbit. Opposite-fiber exclusion is therefore compatible with substantial within-convention nonidentification.

The information-order result has a separate role. The apolar theorem uses order $2m + 2$ to recover the unordered loading support through the binary-form annihilator equations. [Theorem 4](#) shows that, for $m \geq 3$, generic real exclusion of the opposite arrow can already occur at order $2m + 1$. The two thresholds therefore align with distinct targets: generic real arrow separation at $2m + 1$ and loading-support recovery at $2m + 2$.

4.2 Limitations and future work

These qualifications also delimit the connection to empirical practice. The paper provides population-level algebraic identification statements for truncated cumulants. Its immediate user is a theorist or model-class analyst who treats m and the axis normalization as maintained inputs and needs to know whether the population cumulant map separates the two representations. It does not supply a finite-sample estimator, a sampling distribution, a model-selection criterion, or a numerical decision rule. Any applied implementation would have to confront cumulant estimation error, weak higher-order signal, near-exceptional parameter values, and uncertainty about the number of source slots. Recent work on latent non-Gaussian identification and causal discovery develops complementary estimation and testing perspectives, but those problems are not solved by the generic separation theorem alone [[Xie et al., 2023](#), [Morinishi and Shimizu, 2025](#), [Chen et al., 2025a](#)].

The main contribution is algebraic econometric identification for a fixed, axis-normalized

latent-source representation class. Higher-order cumulants generically contain enough information to exclude the opposite convention’s cumulant-map fiber inside that class, the exceptional compatibility set is exactly codimension one at the apolar order, and lower real truncation can suffice for generic opposite-fiber exclusion when $m \geq 3$. These statements establish the population information available before estimation and provide a foundation for future finite-sample procedures.

Appendices

A With proofs and auxiliary lemmas

This appendix records the external algebraic background used to organize the results in [Theorem 1](#) through [Theorem 4](#).

Remark 1. `[def:real-atlas-handle]` *The real atlas handle is the following external classical interface. For every finite family of real polynomials and every variable order, there exists an adapted cylindrical algebraic decomposition that is sign-invariant for the generated projection family, is obtained by recursive section and sector lifting, decides polynomial sign conditions cellwise, and supports Tarski–Seidenberg semialgebraic projection. This is the classical existence-and-structure interface of [Bochnak et al. \[1998, Section 2.3 and Theorem 2.2.1\]](#) and [Basu et al. \[2006, Chapter 5\]](#); it is paper-agnostic and carries no algorithmic, computability, payload, or complexity claim. It is distinct from the exact rational Gröbner–CAD interface.*

Remark 2. `[def:effective-rational-groebner-cad-interface]` (Effective rational Gröbner–CAD interface) *This paper assumes the following cited external interface. There exist natural constants a, b with $a \geq 1$ and fixed uniform exact symbolic machines. For finite supplied polynomial presentations over $\mathbb{Q}$ and $\mathbb{Q}(i)$, supplied admissible monomial and variable orders, and supplied retained coordinate blocks, these machines terminate with certified Gröbner bases and normal forms, elimination ideals, ideal intersections, saturations, the rational complex elimination-closure theorem, and a sign-invariant real cylindrical algebraic decomposition with exact real-algebraic root and sign data, recursively lifted section-and-sector cells, quantifier elimination, cellwise truth and retention for polynomial equalities and weak or strict inequalities, and finite certified traces.*

If $D \geq 1$, every supplied job has ambient dimension at most N and supplied degree envelope at most D , and the total supplied source-polynomial count is at most s , one combined certificate bounds the sum of a finite batch of rational algebra traces, a finite batch of Gaussian-rational algebra traces, and the rational CAD and quantifier-elimination trace by

$$(\max\{2, sD\})^{2^{aN+b}}.$$

The same fixed machines and constants also certify a dependent pipeline when the caller first supplies a jointly presentable cross-source coordinate relation. Two supplied saturated-elimination jobs finish, injective renamings identify their coordinates exactly according to that predeclared relation, their renamed outputs feed an ideal-intersection job, and an injective renaming sends the resulting intersection basis into a dependent CAD job. This certificate additionally requires the positive envelope D to bound the actual saturated-elimination bases, the intersection job and its actual basis, and the CAD input. This paper-agnostic interface assumes no LiNGAM map, moment certificate, exceptional locus, fiber, atlas cell, label, or conclusion.

B Verification and reproducibility scope

Definition 17. [def:global-feasible-fiber-decision] (Order- $2m + 2$ feasible-fiber predicate $\text{FeasFiber}_{m,2m+2}^b(t)$) Set $K = 2m + 2$. For $b \in \{\text{right}, \text{left}\}$, with the branch convention of [Definition 11](#), define the order- K predicate $\text{FeasFiber}_{m,2m+2}^b(t)$ to be the atomic-certificate predicate asserting that there exist real b -loading and source-cumulant coordinates λ , and, for each $0 \leq j \leq m + 1$, witnesses

$$(w_{jh}, z_{jh})_{h=1}^{m+2},$$

such that

$$\Phi_{m,K}^b(\lambda) = t,$$

the direct slope for b is nonzero, all finite loading slopes for b are pairwise distinct, and the following moment conditions hold. Write k_{jr} for the source cumulants contained in λ , let B_r denote the complete exponential Bell polynomial (the cumulant-to-moment polynomial), and set

$$\mu_{jr} = B_r(0, k_{j2}, \dots, k_{jr}).$$

Then

$$\sum_{h=1}^{m+2} w_{jh} = 1, \quad w_{jh} \geq 0, \quad \sum_{h=1}^{m+2} w_{jh} z_{jh}^r = \mu_{jr} \quad (1 \leq r \leq K), \quad \mu_{j2} > 0.$$

Separately, the formal object records an unproved proposition asserting the existence of a sign-invariant cylindrical algebraic decomposition procedure, in variable order t , then λ , then the witnesses, that decides this displayed predicate for the fixed axis b . The predicate definition itself does not assert equivalence with nonemptiness of the corresponding real-source feasible fiber.

Definition 18. [def:direction-selector] (Direction selector S_m) Using the order- $2m + 2$ feasible-fiber predicates of [Definition 17](#), define the partial direction selector

$$S_m(t) = \begin{cases} X \rightarrow Y, & \text{FeasFiber}_{m,2m+2}^{\text{right}}(t) \text{ holds and } \text{FeasFiber}_{m,2m+2}^{\text{left}}(t) \text{ does not,} \\ Y \rightarrow X, & \text{FeasFiber}_{m,2m+2}^{\text{left}}(t) \text{ holds and } \text{FeasFiber}_{m,2m+2}^{\text{right}}(t) \text{ does not,} \\ \text{undefined,} & \text{otherwise.} \end{cases}$$

Separately, the formal object records the unproved proposition that a globally finite semialgebraic function on $\mathbb{R}^{q_{2m+2}}$ agrees with S_m wherever exactly one feasibility predicate holds, together with Boolean decision functions for the forward and reverse maps. This definition supplies neither a proof of that proposition nor an evaluable decision procedure.

Definition 19. [def:separated-model-domain] (Separated model domain $M_{m,K}^{\text{sep}}$) For a structural model M , write P_M for its induced observational law and $D(M) \in \{X \rightarrow Y, Y \rightarrow X\}$ for its arrow. Define $M_{m,K}^{\text{sep}}$ to be the class of structural models M such that either M has a forward representation with parameter $\theta \in F_{m,K}^{\text{right}}$ and

$$R_{m,K}^{\text{left}}(T_K(P_M)) \cap F_{m,K}^{\text{left}} = \emptyset,$$

or M has a reverse representation with parameter $\eta \in F_{m,K}^{\text{left}}$ and

$$R_{m,K}^{\text{right}}(T_K(P_M)) \cap F_{m,K}^{\text{right}} = \emptyset.$$

Remark 3. [thm:exact-real-exceptional-atlas] (Exact Real Exceptional Atlas) *One worthwhile direction for future research is to determine whether For every fixed m , there is a finite computable semialgebraic cylindrical atlas of the real exceptional closure such that, on each base cell, the complete feasible forward and reverse fibers are finite unions of recursively cylindrical section and sector cells and their nonemptiness labels are constant. The atlas includes singular and nongeneric boundary branches and decides exactly when both real arrows are feasible.. Doing so requires the remaining paper-specific exact-atlas constructor would require substantial CAD/elimination substrate, while the node is secondary and has no delivered theorem consumer, which lies beyond the present development.*

The formal claim in this paper is reproducible at the source revision used to build the manuscript. The repository is <https://github.com/Jiyuan-Tan/CausalSmith>; the pinned revision is

2c9214dff2d737c1d131f6056916aad82bdc7e4b,

and the CausalSmith project uses Lean v4.29.0-rc3. The matching machine-readable contract is CausalSmith/doc/presentation/eid_lingam_direction_min_order_v1_truncated_cumulant_minimality/verification_contract.json; the declaration table below refers to that same revision. From the repository root, the checked project is built with

lake -d CausalSmith build.

The repository’s CausalSmith/lean-toolchain, CausalSmith/lakefile.toml, and dependency manifests fix the toolchain and local Causalean dependency used by that build.

Table 2: Main paper results and checked Lean declarations

Paper result	Lean declaration
Generic apolar arrow recovery	CausalSmith.ExactID.EID_LingamDirectionMinOrderV1.generic_apolar_arrow_recovery
Same-arrow fiber consequences	CausalSmith.ExactID.EID_LingamDirectionMinOrderV1.genericArrowRecoveryAndFiberObstruction
Exceptional locus of codimension one	CausalSmith.ExactID.EID_LingamDirectionMinOrderV1.exceptionalLocusCodimensionOne
Improved real information order	CausalSmith.ExactID.EID_LingamDirectionMinOrderV1.improvedRealInformationOrder
Admissible-swap invariance	CausalSmith.ExactID.EID_LingamDirectionMinOrderV1.admissibleSwaps_preserve_direction

The generated `verification_contract.json` in the paper bundle gives the full declaration-to-paper mapping, source file, statement match, and sorry-free status for every displayed formal object. The checked scope covers the frozen definitions and theorems used for the cumulant maps, generic separation, same-arrow fiber behavior, exceptional-locus dimension, and information-order result. Classical library facts and explicitly stated external algebraic inputs remain dependencies, while the prose supplies the surrounding interpretation.

Limitations. The formalization does not certify an estimator, sampling theory, selection of m , or an exact decision procedure for all real feasible fibers, including singular and nongeneric boundary cases. Those tasks are outside the delivered contribution.

C Proofs of the main results

Proof of Theorem 2. Fix $m \geq 1$, put $K = 2m + 2$ and $n = m + 2$, and take for U_m^{right} and U_m^{left} the two loci produced by Theorem 1. That theorem already supplies the relative genericity clause (each locus is contained in the corresponding generic locus of Definition 7 and is relatively

Zariski open and dense) and the real-incidence clause (the Euclidean-open sets meeting the feasible regions, with complexifications inside the loci). It also supplies, for each $\theta \in U_m^{\text{right}}$, three further items used below: the empty opposite-arrow fiber, the univariate polynomial Q with $Q \neq 0$, $Q(0) \neq 0$ and root multiset $\{\gamma, \rho_1, \dots, \rho_m\}$, and the kernel identity: writing

$$f_r(x, y) = \sum_{a=0}^r \binom{r}{a} [\Phi_{m,K}^{\text{right}}(\theta)]_{r,a} x^{r-a} y^a, \quad Q_D = \prod_{j=0}^{m+1} (u_{j2}x - u_{j1}y),$$

one has, for every homogeneous binary form q of degree n ,

$$[q(\partial)f_{n+k} = 0 \text{ for every } 0 \leq k \leq m] \iff \exists c \in \mathbb{C} : q = cQ_D.$$

We refer to this last display as the kernel identity. The mirror statements hold on U_m^{left} .

1. **Forward multiset recovery on the whole fiber.** Let $\theta = (\gamma, \rho, c) \in U_m^{\text{right}}$ and let $\theta' = (\gamma', \rho', c')$ lie in $R_{m,K}^{\text{right}}(\Phi_{m,K}^{\text{right}}(\theta))$. Membership in the fiber compares the two cumulant vectors on the retained coordinates $2 \leq r \leq K$, $a \leq r$, and both maps vanish outside that range, so

$$\Phi_{m,K}^{\text{right}}(\theta') = \Phi_{m,K}^{\text{right}}(\theta).$$

Let $Q'_D = \prod_j (u'_{j2}x - u'_{j1}y)$ be the support annihilator of the loading family of θ' . It is homogeneous of degree n and nonzero, because each loading direction u'_j is either $(1, \cdot)$ or $(0, 1)$ and hence contributes a nonzero linear factor. Since Q'_D vanishes at every direction u'_j , and each divided-power block of θ' is $\sum_j c'_{jr} (u'_{j1}x + u'_{j2}y)^r$, the operator $Q'_D(\partial)$ kills every block of θ' , hence – by the displayed equality of cumulant vectors – every block f_{n+k} of θ . By the kernel identity there is c with $Q'_D = cQ_D$, and $c \neq 0$ since $Q'_D \neq 0$. Dehomogenizing at $x = 1$ turns Q_D into $(-1)^{m+1}(z - \gamma) \prod_{i=1}^m (z - \rho_i)$ and Q'_D into the corresponding product $(-1)^{m+1}(z - \gamma') \prod_{i=1}^m (z - \rho'_i)$ for θ' ; equal up to the nonzero scalar c , the two have the same roots, so

$$\{\gamma', \rho'_1, \dots, \rho'_m\} = \{\gamma, \rho_1, \dots, \rho_m\}$$

as multisets. This is the forward same-arrow recovery clause.

2. **Forward opposite-arrow exclusion, and nonzero latent slopes.** For the same θ , [Theorem 1](#) gives

$$R_{m,K}^{\text{left}}(\Phi_{m,K}^{\text{right}}(\theta)) = \emptyset.$$

Moreover, if some ρ_i were 0, then 0 would be a root of the polynomial Q , whose root multiset is $\{\gamma, \rho_1, \dots, \rho_m\}$, contradicting $Q(0) \neq 0$. Hence

$$\rho_i \neq 0 \quad (1 \leq i \leq m).$$

3. **The direct/latent swap.** Fix $i \in \{1, \dots, m\}$ and define $\theta' = (\gamma', \rho', c')$ by exchanging the direct source with the i th latent source, slope and weight family together:

$$\gamma' = \rho_i, \quad \rho'_i = \gamma, \quad \rho'_\ell = \rho_\ell \ (\ell \neq i), \quad c'_{0r} = c_{ir}, \quad c'_{ir} = c_{0r}, \quad c'_{jr} = c_{jr} \ (j \notin \{0, i\}),$$

for all $2 \leq r \leq K$. This is exactly the swap of the full pairs $(\gamma, (c_{0r})_{r=2}^K)$ and $(\rho_i, (c_{ir})_{r=2}^K)$ required in the statement, with all other coordinates fixed.

4. **The swap stays generic.** The entries of θ' are those of θ permuted, so its retained weights are again all nonzero and its slope list $\gamma', \rho'_1, \dots, \rho'_m$ is the list $\gamma, \rho_1, \dots, \rho_m$ with two entries exchanged, hence again pairwise distinct; and $\gamma' = \rho_i \neq 0$ by Step 2. Thus $\theta' \in \Theta_{m,K}^{\text{right}, \circ}$.

5. **The swap preserves the cumulant vector.** Both exchanged loading directions have first coordinate 1:

$$u_0(\theta) = (1, \gamma), \quad u_i(\theta) = (1, \rho_i),$$

so exchanging the two slopes exchanges the two directions, $u_0(\theta') = u_i(\theta)$ and $u_i(\theta') = u_0(\theta)$, while $u_j(\theta') = u_j(\theta)$ for all other j ; the weights were exchanged along with them. Hence for every retained (r, a) the defining sum of Definition 4 is the original sum reindexed by the transposition of the source labels 0 and i , and finite sums are invariant under such reindexing:

$$\Phi_{m,K}^{\text{right}}(\theta') = \Phi_{m,K}^{\text{right}}(\theta).$$

6. **The swap is outside the admissible orbit.** By Definition 9 every admissible relabeling $\pi \in G_m$ leaves the direct slope untouched, so every member of the G_m -orbit of θ has direct slope γ . But θ' has direct slope ρ_i , and genericity gives $\gamma \neq \rho_i$. Hence $\theta' \notin G_m \cdot \theta$.

7. **Exact dimension of the forward fiber for $m \geq 2$.** Write $Z = R_{m,K}^{\text{right}}(\Phi_{m,K}^{\text{right}}(\theta))$. For any base point $\vartheta \in Z$ whose finite loading slopes are pairwise distinct, let

$$Z^{\text{fix}}(\vartheta) = \{\theta' \in Z : \gamma' = \vartheta_\gamma, \rho' = \vartheta_\rho\}$$

be the part of the common fiber with the loading coordinates frozen at those of ϑ . Because $\vartheta \in Z$, the two targets $\Phi_{m,K}^{\text{right}}(\vartheta)$ and $\Phi_{m,K}^{\text{right}}(\theta)$ are equal. Thus, if θ' has the same loadings as ϑ , then $\theta' \in Z^{\text{fix}}(\vartheta)$ exactly when its weight difference $e_{jr} = c'_{jr} - \vartheta_{c,jr}$ lies in

$$\mathcal{K}(\vartheta) = \left\{ (e_{jr}) : \sum_{j=0}^{m+1} e_{jr} (u_{j1})^{r-a} (u_{j2})^a = 0 \text{ for all } 2 \leq r \leq K, 0 \leq a \leq r \right\},$$

because the cumulant map is linear in the weights at fixed loadings. Hence $Z^{\text{fix}}(\vartheta)$ is the translate of the linear space $\mathcal{K}(\vartheta)$ through ϑ , and is an irreducible closed set of dimension $\dim \mathcal{K}(\vartheta)$.

The conditions in $\mathcal{K}(\vartheta)$ decouple across orders, and at order r , writing $s_0 = \vartheta_\gamma$, $s_j = \vartheta_{\rho,j}$ for the $m+1$ finite slopes and using $u_j = (1, s_j)$, $u_{m+1} = (0, 1)$, they read

$$\sum_{j=0}^m e_{jr} s_j^a = 0 \quad (0 \leq a \leq r-1), \quad \sum_{j=0}^m e_{jr} s_j^r + e_{m+1,r} = 0.$$

The last equation determines $e_{m+1,r}$ from the others and imposes no further restriction; the first block is the Vandermonde system in the $m+1$ nodes $s_0, \dots, s_m$ with r rows, which has rank $\min(r, m+1)$ when the nodes are pairwise distinct. Hence the order- r kernel has dimension $(m+1-r)^+$, and

$$\dim \mathcal{K}(\vartheta) = \sum_{r=2}^K (m+1-r)^+ = \sum_{r=2}^m (m+1-r) = \frac{m(m-1)}{2}.$$

Applying this to $\vartheta = \theta$, whose slopes are pairwise distinct and which lies in Z , gives a nonempty irreducible closed subset $Z^{\text{fix}}(\theta) \subseteq Z$ of dimension $m(m-1)/2$; this gives a strictly increasing chain of $m(m-1)/2 + 1$ irreducible closed subsets of Z .

For the reverse inequality, let $T \subseteq Z$ be irreducible and closed. By Step 1 every point of Z has loading-slope multiset $\{\gamma, \rho_1, \dots, \rho_m\}$, so each of the $m+1$ slope coordinates takes values in that finite set on T ; a coordinate function taking finitely many values on an irreducible set is constant, since otherwise T would split as a finite union of proper closed pieces. Choose a point $\theta_0 \in T$. The preceding constancy shows that all loading coordinates on T are those of θ_0 , so $T \subseteq Z^{\text{fix}}(\theta_0)$. The point θ_0 lies in Z , and its finite slopes are pairwise distinct because their multiset is that of θ , which has $m+1$ distinct entries. Therefore $Z^{\text{fix}}(\theta_0)$ has dimension $m(m-1)/2$ and contains no chain of $m(m-1)/2 + 2$ irreducible closed subsets; neither does Z . Combining the two bounds, for $m \geq 2$ the forward fiber has exact relative Zariski dimension

$$\frac{m(m-1)}{2}.$$

8. **The reverse locus.** Let $\eta = (\delta, \sigma, d) \in U_m^{\text{left}}$. The four preceding arguments run verbatim after exchanging the two coordinate axes, with the reverse loading family $v_0 = (1, 0)$, $v_j = (\sigma_j, 1)$, $v_{m+1} = (\delta, 1)$ in place of the forward one and dehomogenization at $y = 1$ in place of $x = 1$. They give: every $\eta' \in R_{m,K}^{\text{left}}(\Phi_{m,K}^{\text{left}}(\eta))$ satisfies

$$\{\delta', \sigma'_1, \dots, \sigma'_m\} = \{\delta, \sigma_1, \dots, \sigma_m\};$$

the opposite-arrow fiber

$$R_{m,K}^{\text{right}}(\Phi_{m,K}^{\text{left}}(\eta)) = \emptyset$$

is empty, and $\sigma_i \neq 0$ for every i because the reverse root polynomial does not vanish at 0; for each i the parameter η' defined by

$$\delta' = \sigma_i, \quad \sigma'_i = \delta, \quad \sigma'_\ell = \sigma_\ell \ (\ell \neq i), \quad d'_{m+1,r} = d_{ir}, \quad d'_{ir} = d_{m+1,r}, \quad d'_{jr} = d_{jr} \text{ otherwise,}$$

lies in $\Theta_{m,K}^{\text{left},\circ}$, satisfies

$$\Phi_{m,K}^{\text{left}}(\eta') = \Phi_{m,K}^{\text{left}}(\eta)$$

– here the exchanged directions $v_{m+1} = (\delta, 1)$ and $v_i = (\sigma_i, 1)$ both have second coordinate 1, so exchanging the two slopes exchanges the two directions and the sum is reindexed by the transposition of the source labels $m+1$ and i – and is not in the admissible G_m -orbit of η , since admissible relabelings fix the reverse direct slope δ while η' has direct slope $\sigma_i \neq \delta$. Finally, for $m \geq 2$ the same fixed-loading decomposition, applied only to base points in the common reverse fiber and then to irreducible closed subsets selected from that fiber, uses the $m+1$ distinct reverse finite loading coordinates $\delta, \sigma_1, \dots, \sigma_m$ as Vandermonde nodes and gives exact relative Zariski dimension

$$\frac{m(m-1)}{2}$$

for $R_{m,K}^{\text{left}}(\Phi_{m,K}^{\text{left}}(\eta))$.

□

Proof of Theorem 3. Fix $m \geq 1$ and set $K = 2m+2$. All sets below live in the retained cumulant coordinates $t_{r,a}$, $2 \leq r \leq K$, $0 \leq a \leq r$; both arrow maps of Definition 4 and Definition 5 vanish outside that range, so $\overline{E}_m$ and the two image varieties $C_{m,K}^{\text{right}}$, $C_{m,K}^{\text{left}}$ of Definition 6 all consist of vectors supported there, the projection onto $\mathbb{C}^{q_K}$ is injective on them, and every chain, dimension and codimension statement below may be read in $\mathbb{C}^{q_K}$. Throughout, “dimension d ” means: there is a strictly increasing chain of $d+1$ nonempty irreducible closed subsets, and none of length $d+2$; and “ $\overline{E}_m$ has codimension d in X ” means: every irreducible component D of $\overline{E}_m$ admits a strictly increasing chain of $d+1$ irreducible closed sets running from D to X , and some component admits no such chain of length $d+2$. Put

$$D_m := \frac{3(m+2)^2 + (m+2)}{2} - 5.$$

1. **The common-axis family.** Let

$$\Delta_m := \{\theta = (\gamma, \rho, c) \in \Theta_{m,K}^{\text{right},\circ} : \rho_1 = 0\}, \quad W_m := \text{ZariskiClosure}(\Phi_{m,K}^{\text{right}}(\Delta_m)).$$

For $\theta \in \Delta_m$ genericity gives $\gamma \neq 0$ and pairwise distinct slopes, so $\rho_i \neq \rho_1 = 0$ for $2 \leq i \leq m$. Define the reverse twin $\theta^{\text{tw}} = (\delta, \sigma, d)$ by

$$\delta = \gamma^{-1}, \quad \sigma_1 = 0, \quad \sigma_i = \rho_i^{-1} \quad (2 \leq i \leq m),$$

$$d_{0r} = c_{1r}, \quad d_{1r} = c_{m+1,r}, \quad d_{jr} = c_{jr} \rho_j^r \quad (2 \leq j \leq m), \quad d_{m+1,r} = c_{0r} \gamma^r, \quad 2 \leq r \leq K.$$

Then the two loading families agree up to scalars, direction by direction:

$$u_1 = (1, 0) = v_0, \quad u_{m+1} = (0, 1) = v_1, \quad u_0 = (1, \gamma) = \gamma v_{m+1}, \quad u_j = (1, \rho_j) = \rho_j v_j \quad (2 \leq j \leq m).$$

The (r, a) summand $c_{jr}(u_{j1})^{r-a}(u_{j2})^a$ is homogeneous of degree r in the loading vector, so rescaling a loading by λ is absorbed by multiplying its weight by λ^r ; this is exactly the weight prescription above. Summing the four cases over the reindexed source labels gives

$$\Phi_{m,K}^{\text{left}}(\theta^{\text{tw}}) = \Phi_{m,K}^{\text{right}}(\theta), \quad \theta \in \Delta_m.$$

2. $W_m \subseteq \overline{E}_m$. Let $\theta \in \Delta_m$ and $t = \Phi_{m,K}^{\text{right}}(\theta)$. Then θ itself lies in $R_{m,K}^{\text{right}}(t) \cap \Theta_{m,K}^{\text{right},\circ}$, and by the previous display $\theta^{\text{tw}} \in R_{m,K}^{\text{left}}(t)$, so

$$R_{m,K}^{\text{right}}(t) \cap \Theta_{m,K}^{\text{right},\circ} \neq \emptyset \quad \text{and} \quad R_{m,K}^{\text{left}}(t) \neq \emptyset,$$

which is the first alternative in the definition of E_m in Definition 13. Hence $\Phi_{m,K}^{\text{right}}(\Delta_m) \subseteq E_m$, and taking Zariski closures, $W_m \subseteq \overline{E}_m$. Moreover $E_m \subseteq C_{m,K}^{\text{right}}$ and $E_m \subseteq C_{m,K}^{\text{left}}$, because a point of E_m has a nonempty fiber under *both* arrow maps and therefore lies in both images; closing up,

$$W_m \subseteq \overline{E}_m \subseteq C_{m,K}^{\text{right}} \cap C_{m,K}^{\text{left}}.$$

3. W_m is irreducible. Deleting the pinned coordinate $\rho_1 = 0$ presents Δ_m as the nonvanishing locus of a single nonzero polynomial (the generic-locus product of Definition 7 with ρ_1 set to 0) inside the remaining affine coordinates; such a locus is Zariski dense in that affine space, and the Zariski closure of the image of a dense set under a polynomial map is irreducible. Hence W_m is a nonempty irreducible closed set. The same argument applied to the whole affine parameter space shows that $C_{m,K}^{\text{right}}$ and $C_{m,K}^{\text{left}}$ are irreducible.

4. **Upper bounds on the two dimensions.** Every retained coordinate of $\Phi_{m,K}^{\text{right}}$ is a polynomial in the finite family

$$\mathcal{G} = \{\gamma, \rho_1, \dots, \rho_m\} \cup \{t_{r,a} : 2 \leq r \leq m, 0 \leq a \leq r\} \cup \{c_{jr} : 0 \leq j \leq m+1, m+1 \leq r \leq K\} :$$

a coordinate of order $r \leq m$ is itself a member of $\mathcal{G}$, and a coordinate of order $r \geq m+1$ equals $\sum_j c_{jr} (u_{j1})^{r-a} (u_{j2})^a$, a polynomial in the slopes and the high-order weights listed in $\mathcal{G}$. Counting,

$$\#\mathcal{G} = (m+1) + q_m + (m+2)^2 = D_m + 1, \quad q_m = \frac{m(m+3)}{2} - 2,$$

so the coordinate algebra of the forward image has transcendence degree at most $D_m + 1$ and therefore

$$\dim C_{m,K}^{\text{right}} \leq D_m + 1.$$

On Δ_m the coordinate ρ_1 is the constant 0, so the same family with ρ_1 deleted generates; it has D_m members, whence $\dim W_m \leq D_m$.

5. **Matching lower bound for the forward image: a nonvanishing $(D_m+1) \times (D_m+1)$ Jacobian minor.** Select the following rows (retained cumulant coordinates) and columns (parameter coordinates):

- for each order $2 \leq r \leq m$: the rows $t_{r,a}$, $0 \leq a \leq r$, and the columns $c_{a,r}$, $0 \leq a \leq r$;
- for each order $m+1 \leq r \leq 2m+1$: the rows $t_{r,a}$, $0 \leq a \leq m$, together with $t_{r,r}$, and the columns $c_{j,r}$, $0 \leq j \leq m$, together with $c_{m+1,r}$;
- at the top order $r = K$: the rows $t_{K,e}$, $0 \leq e \leq 2m+1$, and the columns consisting of the $m+1$ top weights $c_{j,K}$, $0 \leq j \leq m$, and the $m+1$ finite slopes $\gamma, \rho_1, \dots, \rho_m$; augmented by the extra row $t_{K,K}$ and the extra column $c_{m+1,K}$.

The total count is

$$\sum_{r=2}^m (r+1) + (m+1)(m+2) + (2m+2) + 1 = D_m + 1.$$

Evaluate the minor at the parameter θ^J with

$$\gamma = 1, \quad \rho_i = i+1 \ (1 \leq i \leq m), \quad c_{jr} = 1 \text{ for all } j \text{ and all } 2 \leq r \leq K,$$

whose finite loading directions are $u_j = (1, j+1)$ for $0 \leq j \leq m$ and $u_{m+1} = (0, 1)$. A weight column of order r has zero derivative in every cumulant coordinate of an order other than r , because that coordinate does not involve $c_{\cdot,r}$; placing the slope columns last, the matrix is therefore block upper triangular with the following diagonal blocks.

- Order $r \leq m$: $\partial t_{r,a} / \partial c_{j,r} = (u_{j1})^{r-a} (u_{j2})^a = (j+1)^a$, the transposed Vandermonde matrix in the $r+1$ distinct nodes $1, \dots, r+1$; its determinant is nonzero.
- Order $m+1 \leq r \leq 2m+1$: for $j \leq m$ and $a \leq m$ the entry is $(j+1)^a$, the Vandermonde matrix in the $m+1$ distinct nodes $1, \dots, m+1$; the column $c_{m+1,r}$ has entries $(u_{m+1,1})^{r-a} (u_{m+1,2})^a = 0$ for $a \leq m < r$ and 1 in the row $t_{r,r}$. The block is triangular with determinant equal to that Vandermonde determinant, hence nonzero.

- Order K : the weight columns give $\partial t_{K,e}/\partial c_{j,K} = (j+1)^e$ and the slope columns give $\partial t_{K,e}/\partial \rho_j = e c_{j,K} (j+1)^{e-1} = e (j+1)^{e-1}$, so this is the confluent (Hermite) Vandermonde matrix with value and first-derivative columns at the $m+1$ distinct nodes $1, \dots, m+1$ and monomial rows of degrees $0, \dots, 2m+1$; its determinant is nonzero. The augmenting column $c_{m+1,K}$ contributes $0^{K-e} 1^e = 0$ in every row $e \leq 2m+1$ and 1 in the row $t_{K,K}$, so the augmented block is triangular with the same nonzero determinant.

The product of the diagonal determinants is nonzero, so the selected minor is a nonzero polynomial and

$$\dim C_{m,K}^{\text{right}} \geq D_m + 1.$$

With the upper bound of the previous step, $\dim C_{m,K}^{\text{right}} = D_m + 1$.

6. **Matching lower bound for the common-axis image.** On Δ_m the slope ρ_1 is frozen at 0, so it supplies no column; delete it and use the same rows with the top row $t_{K,2m+1}$ deleted as well, leaving D_m rows and D_m columns. Evaluate at the pinned witness with

$$\rho_1 = 0, \quad \gamma = 1, \quad \rho_i = i \ (2 \leq i \leq m), \quad c_{jr} = 1,$$

whose finite nodes are the distinct values $0, 1, \dots, m$. The ordinary blocks are Vandermonde and augmented-Vandermonde exactly as before, and the top block is now the pinned Hermite matrix carrying values at all $m+1$ nodes and first derivatives at the m unpinned nodes, with monomial rows of degrees $0, \dots, 2m$; its determinant is nonzero because the nodes are distinct. Hence $\dim W_m \geq D_m$, and with the upper bound

$$\dim W_m = D_m.$$

7. **No intermediate irreducible set between W_m and the forward variety.** Suppose Y were irreducible and closed with $W_m \subsetneq Y \subsetneq C_{m,K}^{\text{right}}$. Appending Y and $C_{m,K}^{\text{right}}$ to a maximal chain of $D_m + 1$ irreducible closed subsets of W_m produces a strictly increasing chain of $D_m + 3$ irreducible closed subsets of $C_{m,K}^{\text{right}}$, contradicting $\dim C_{m,K}^{\text{right}} = D_m + 1$. Hence

$$\nexists Y \text{ irreducible closed with } W_m \subsetneq Y \subsetneq C_{m,K}^{\text{right}}.$$

8. **The same for the reverse variety, by coordinate reversal.** Let ι be the involution of cumulant space that exchanges the two observed coordinates,

$$(\iota t)_{r,a} = t_{r,r-a} \quad (2 \leq r \leq K, \ 0 \leq a \leq r).$$

Exchanging the roles of the two axes in [Definition 4](#) and [Definition 5](#) gives $\Phi_{m,K}^{\text{left}}(\theta^{\text{rev}}) = \iota(\Phi_{m,K}^{\text{right}}(\theta))$, where θ^{rev} keeps the slopes of θ and exchanges the weight families of the two axis source labels $j = 0$ and $j = m+1$. That relabelling is an involution of the parameter space, hence onto, so

$$\iota(C_{m,K}^{\text{right}}) = C_{m,K}^{\text{left}}.$$

For W_m the exchange is not applied to θ itself but to the reverse twin of Step 1. Let $\theta \in \Delta_m$ and put $\theta^\dagger = (\theta^{\text{tw}})^{\text{rev}}$, the same axis relabelling read from the reverse chart back to the forward one; explicitly, $\theta^\dagger$ carries the reciprocated slopes

$$\gamma^\dagger = \gamma^{-1}, \quad \rho_1^\dagger = 0, \quad \rho_i^\dagger = \rho_i^{-1} \ (2 \leq i \leq m),$$

together with the weight family d of θ^{tw} , its two axis labels 0 and $m+1$ exchanged. Genericity of θ gives $\gamma \neq 0$ and pairwise distinct nonzero $\rho_2, \dots, \rho_m$, so the reciprocals $\gamma^{-1}, \rho_2^{-1}, \dots, \rho_m^{-1}$ are again pairwise distinct and nonzero, hence distinct from $\rho_1^\dagger = 0$; the weights of θ^{tw} are the weights of θ times powers of nonzero slopes, hence nonzero. Therefore $\theta^\dagger \in \Delta_m$. Because the axis relabelling is an involution, $(\theta^\dagger)^{\text{rev}} = \theta^{\text{tw}}$, so the axis-exchange display applied to $\theta^\dagger$ reads $\Phi_{m,K}^{\text{left}}(\theta^{\text{tw}}) = \iota(\Phi_{m,K}^{\text{right}}(\theta^\dagger))$; combining with $\Phi_{m,K}^{\text{left}}(\theta^{\text{tw}}) = \Phi_{m,K}^{\text{right}}(\theta)$ from Step 1,

$$\iota(\Phi_{m,K}^{\text{right}}(\theta^\dagger)) = \Phi_{m,K}^{\text{right}}(\theta), \quad \theta \in \Delta_m.$$

So ι carries $\Phi_{m,K}^{\text{right}}(\Delta_m)$ into itself, and being an involution it carries it onto itself; closing up,

$$\iota(W_m) = W_m.$$

What is stable under the exchange is thus the image $\Phi_{m,K}^{\text{right}}(\Delta_m)$ under a reciprocal-slope reparametrization of Δ_m , not the individual parameter. Since ι is an injective polynomial involution it preserves irreducibility and strict inclusions, so an intermediate irreducible closed set between W_m and $C_{m,K}^{\text{left}}$ would be carried to one between W_m and $C_{m,K}^{\text{right}}$. By the previous step there is none.

9. **Both inclusions $\overline{E}_m \subseteq C_{m,K}^{\text{right}}$, $\overline{E}_m \subseteq C_{m,K}^{\text{left}}$ are proper.** For a cumulant vector t form the $(m+2) \times (m+2)$ matrix

$$N^{\text{hor}}(t)_{0,b} = t_{m+2,b+1}, \quad N^{\text{hor}}(t)_{a+1,b} = \binom{m}{a} t_{2m+2,a+b+1} \quad (0 \leq a, b \leq m+1, a \leq m),$$

and set $P^{\text{hor}}(t) = \det N^{\text{hor}}(t)$, a polynomial in the retained coordinates. If t is the image of a parameter with loading family $(w_j)_{j=0}^{m+1}$ and weights (c_{jr}) , then substituting the coordinate formula of [Definition 4](#) (resp. [Definition 5](#)) and splitting the exponents factors the matrix as $N^{\text{hor}}(t) = AB$ with

$$A_{0j} = c_{j,m+2}, \quad A_{a+1,j} = \binom{m}{a} c_{j,2m+2} (w_{j1})^{m-a} (w_{j2})^a, \quad B_{jb} = (w_{j1})^{m+1-b} (w_{j2})^{b+1}.$$

Every entry of the j th row of B carries the factor $(w_{j2})^{b+1}$ with $b+1 \geq 1$; hence if some loading direction has vanishing second coordinate, that row of B is zero and $P^{\text{hor}}(t) = 0$. Every reverse loading family contains $v_0 = (1, 0)$, so P^{hor} vanishes on the whole reverse image and therefore on its Zariski closure $C_{m,K}^{\text{left}}$, hence on $\overline{E}_m$ by Step 2. On the other hand take the forward parameter θ^* with

$$\gamma = 1, \quad \rho_i = i+1, \quad c_{j,m+2} = \mathbb{K}_{\{j=m+1\}}, \quad c_{j,2m+2} = \mathbb{K}_{\{j \leq m\}}, \quad c_{jr} = 0 \text{ for } r \notin \{m+2, 2m+2\}.$$

At θ^* the matrix A has last column $(1, 0, \dots, 0)^\top$ and, in its first $m+1$ columns and last $m+1$ rows, the entries $\binom{m}{a} (j+1)^a$; the matrix B has last row $(0, \dots, 0, 1)$ and, in its first $m+1$ rows, the entries $(j+1)^{b+1}$. Both are triangular with a Vandermonde determinant in the distinct nodes $1, \dots, m+1$ as the surviving factor, so $\det A \neq 0 \neq \det B$ and

$$P^{\text{hor}}(\Phi_{m,K}^{\text{right}}(\theta^*)) \neq 0.$$

Thus $\Phi_{m,K}^{\text{right}}(\theta^*) \in C_{m,K}^{\text{right}} \setminus \overline{E}_m$ and $\overline{E}_m \neq C_{m,K}^{\text{right}}$. Applying ι to the whole construction — equivalently, using $P^{\text{ver}} = P^{\text{hor}} \circ \iota$, which vanishes on $C_{m,K}^{\text{right}}$ because every forward loading family contains $u_{m+1} = (0, 1)$, and is nonzero at the reversed witness — gives $\overline{E}_m \neq C_{m,K}^{\text{left}}$.

10. **W_m is an irreducible component of $\overline{E}_m$.** W_m is irreducible and closed and $W_m \subseteq \overline{E}_m$. If Y is irreducible and closed with $W_m \subseteq Y \subseteq \overline{E}_m$ and $Y \neq W_m$, then $W_m \subsetneq Y \subseteq \overline{E}_m \subsetneq C_{m,K}^{\text{right}}$ by Step 9, so Y is an intermediate set of the kind excluded in Step 7 — a contradiction. Hence W_m is maximal among irreducible closed subsets of $\overline{E}_m$, i.e. an irreducible component.
11. **Codimension exactly one.** Let X be either $C_{m,K}^{\text{right}}$ or $C_{m,K}^{\text{left}}$ and let P be the corresponding contraction minor P^{hor} or P^{ver} from Step 9. First, every irreducible component D of $\overline{E}_m$ satisfies $D \subseteq \overline{E}_m \subsetneq X$, so $D \subsetneq X$ and the two-term chain $D \subsetneq X$ of irreducible closed sets exists; this is the required chain of length $2 = 1 + 1$ from D to X . Second, P vanishes on W_m and not on X , so the prime ideal of W_m contains the ideal of X together with P and, by Step 7 and Step 8, no prime of the coordinate ring lies strictly between them: the ideal of W_m is a minimal prime over the ideal of X enlarged by P . Equivalently, no irreducible closed Y satisfies $W_m \subsetneq Y \subsetneq X$, so the component W_m admits no chain of length 3 from W_m to X . The two clauses are exactly codimension 1, so

$$\text{codim}(\overline{E}_m, C_{m,K}^{\text{right}}) = 1, \quad \text{codim}(\overline{E}_m, C_{m,K}^{\text{left}}) = 1.$$

12. **Exclusion of codimension m for $m \geq 2$.** Suppose $\overline{E}_m$ had codimension m in X with $m \geq 2$. Applied to the component W_m , the first clause of that statement supplies a strictly increasing chain $Y_0 \subsetneq Y_1 \subsetneq \dots \subsetneq Y_m$ of irreducible closed sets with $Y_0 = W_m$ and $Y_m = X$. Keeping only the indices $0, 1, m$ leaves a strictly increasing chain of three irreducible closed sets running from W_m to X , which the second clause of Step 11 forbids. Hence

$$\text{codim}(\overline{E}_m, C_{m,K}^{\text{right}}) \neq m, \quad \text{codim}(\overline{E}_m, C_{m,K}^{\text{left}}) \neq m \quad (m \geq 2).$$

13. **Forward preimage.** By Definition 13, $\theta \in H_m^{\text{right}}$ means $\theta \in \Theta_{m,K}^{\text{right},\circ}$ and $\Phi_{m,K}^{\text{right}}(\theta) \in E_m$. Write $t = \Phi_{m,K}^{\text{right}}(\theta)$, with the fibers of Definition 8. If θ is generic then $\theta \in R_{m,K}^{\text{right}}(t) \cap \Theta_{m,K}^{\text{right},\circ}$, so the first alternative in the definition of E_m reduces to $R_{m,K}^{\text{left}}(t) \neq \emptyset$; and the second alternative also forces $R_{m,K}^{\text{left}}(t) \neq \emptyset$, since it asserts in particular that the reverse fiber over t contains a generic point. Conversely, if θ is generic and $R_{m,K}^{\text{left}}(t) \neq \emptyset$ then the first alternative holds outright. Therefore

$$\theta \in H_m^{\text{right}} \iff \theta \in \Theta_{m,K}^{\text{right},\circ} \text{ and } R_{m,K}^{\text{left}}(\Phi_{m,K}^{\text{right}}(\theta)) \neq \emptyset.$$

Coordinates outside the range $2 \leq r \leq K$ vanish under the truncation convention, so they impose no further condition on either side.

14. **Reverse preimage.** The same argument with the two arrows interchanged gives

$$\eta \in H_m^{\text{left}} \iff \eta \in \Theta_{m,K}^{\text{left},\circ} \text{ and } R_{m,K}^{\text{right}}(\Phi_{m,K}^{\text{left}}(\eta)) \neq \emptyset.$$

15. **Low-dimensional specializations.** Let $m \in \{1, 2\}$. A pair $p = (p^{\text{right}}, p^{\text{left}})$ solves the incidence system A_m of [Definition 14](#) precisely when both members lie in the retained coordinate space, their retained cumulants agree,

$$\Phi_{m,K}^{\text{right}}(p^{\text{right}}) = \Phi_{m,K}^{\text{left}}(p^{\text{left}}),$$

and at least one member lies in its generic locus. If $t \in E_m$, then by definition one arrow has a generic fiber point over t and the other a fiber point over t ; calling them p^{right} and p^{left} , their retained cumulants both equal t , so the displayed equality holds, the generic-locus disjunction holds, and $t = \Phi_{m,K}^{\text{right}}(p^{\text{right}})$; thus p solves A_m and projects to t . Conversely, if p solves A_m and $t = \Phi_{m,K}^{\text{right}}(p^{\text{right}})$, then $p^{\text{right}} \in R_{m,K}^{\text{right}}(t)$ and, by the displayed equality, $p^{\text{left}} \in R_{m,K}^{\text{left}}(t)$; whichever member is generic supplies the corresponding alternative in the definition of E_m , so $t \in E_m$. Coordinates beyond the truncation order vanish on both sides by convention. Taking $m = 1$ (with $K = 4$) and $m = 2$ (with $K = 6$) gives

$$E_1 = \{t : \exists p \text{ satisfying } A_1, \Phi_{1,4}^{\text{right}}(p^{\text{right}}) = t\}, \quad E_2 = \{t : \exists p \text{ satisfying } A_2, \Phi_{2,6}^{\text{right}}(p^{\text{right}}) = t\}.$$

□

Proof of [Theorem 4](#). Fix $m \geq 3$ and put $L = 2m + 1$ and $n = m + 2$. As in the apolar setting write $\ell_u(x, y) = u_1x + u_2y$, write $q(u)$ for the evaluation of a binary form at a direction u , and, for a loading family $(w_j)_{j=0}^{m+1}$, write its support annihilator in the normalization of [Definition 12](#) as

$$Q_D = \prod_{j=0}^{m+1} (w_{j2}x - w_{j1}y),$$

a nonzero homogeneous form of degree n vanishing exactly on the n loading directions. Expanding the binomial in $\ell_{u_j}^r$ and comparing with [Definition 4](#) gives, for $2 \leq r \leq L$,

$$f_r(x, y) = \sum_{a=0}^r \binom{r}{a} [\Phi_{m,L}^{\text{right}}(\theta)]_{r,a} x^{r-a} y^a = \sum_{j=0}^{m+1} c_{jr} \ell_{u_j}(x, y)^r,$$

and likewise on the reverse side; and for q homogeneous of degree n ,

$$q(\partial) f_{n+k} = \frac{(n+k)!}{k!} \sum_{j=0}^{m+1} c_{j,n+k} q(u_j) \ell_{u_j}^k, \quad 0 \leq k \leq m-1,$$

the range of k being limited by $n+k \leq L$. This is one block fewer than at order $2m+2$.

1. **The lower-order rank determinant.** Out of the m available contraction equations select $m+2$ scalar rows: the single equation at $k=0$, the coefficient of x in the equation at $k=1$, and all m coefficients of the degree- $(m-1)$ form at $k=m-1$. Explicitly, define the $(m+2) \times (m+2)$ matrix $N^{\text{right}}(\theta)$, with columns indexed by the source labels $j = 0, \dots, m+1$, by

$$N_{0,j}^{\text{right}} = c_{j,m+2}, \quad N_{1,j}^{\text{right}} = c_{j,m+3} u_{j1}, \quad N_{i+2,j}^{\text{right}} = \binom{m-1}{i} c_{j,2m+1} (u_{j1})^{m-1-i} (u_{j2})^i \quad (0 \leq i \leq m-1),$$

and let $\Pi_r(\theta) = \det N^{\text{right}}(\theta)$. The three row blocks come from the three distinct available orders $m+2$, $m+3$ and $2m+1$, which are pairwise distinct precisely because $m \geq 3$. By construction, $N^{\text{right}}(\theta)e$ collects the selected rows of the contraction equations applied to a vector $e = (e_j)_j$; hence

$$\Pi_r(\theta) \neq 0 \implies \left[\sum_j c_{j,n+k} e_j \ell_{u_j}^k = 0 \text{ for all } 0 \leq k \leq m-1 \right] \implies e = 0.$$

The determinant is a nonzero polynomial: at the forward parameter

$$\gamma = 1, \quad \rho_i = i+1, \quad c_{j,m+2} = \mathbb{K}_{\{j=m+1\}}, \quad c_{j,m+3} = \mathbb{K}_{\{j=0\}}, \quad c_{j,2m+1} = \mathbb{K}_{\{1 \leq j \leq m\}}, \quad c_{jr} = 0 \text{ otherwise,}$$

the first row is supported on the column $j = m+1$, the second on the column $j = 0$, and the last m rows on the columns $j = 1, \dots, m$, where they form the matrix $\binom{m-1}{i} \rho_j^i$; the matrix is therefore block triangular after a permutation of columns, with a Vandermonde determinant in the m distinct nodes $\rho_1, \dots, \rho_m$ as its only nonconstant factor, so it is nonzero. All weights of this witness vanish outside the orders $2 \leq r \leq L$. Define N^{left} and Π_l by the coordinate-reversed formulas, with the reverse loadings v_j and weights d_{jr} , and the mirror witness gives $\Pi_l \neq 0$ likewise.

2. **The exceptional sets.** On the real forward parameter space put

$$\Delta_m^{\text{right}} = \left(\prod_{i=1}^m \rho_i \right) \Pi_r, \quad E^{\text{right}} = \{\theta : \Delta_m^{\text{right}}(\theta) = 0\},$$

and on the real reverse parameter space

$$\Delta_m^{\text{left}} = \left(\prod_{i=1}^m \sigma_i \right) \Pi_l, \quad E^{\text{left}} = \{\eta : \Delta_m^{\text{left}}(\eta) = 0\}.$$

Both Δ_m^{right} and Δ_m^{left} are nonzero real polynomials by the determinant witnesses just constructed, so E^{right} and E^{left} are proper real algebraic subsets of their parameter spaces. Their complements meet the feasible regions: the truncated moment problem at order L supplies a real cumulant vector c^0 and a radius $\varepsilon > 0$ such that

$$|c'_r - c_r^0| < \varepsilon \ (2 \leq r \leq L) \implies \exists \text{ a centered non-Gaussian law } \nu, \mathbb{E}|\cdot|^L < \infty, \text{Cum}_r(\nu) = c'_r \ (2 \leq r \leq L).$$

Consider the box in which every slope coordinate ranges over $\mathbb{R}$ and every retained weight coordinate c_{jr} ranges over $(c_r^0 - \varepsilon, c_r^0 + \varepsilon)$. Each factor is infinite, and the product of Δ_m^{right} with the genericity polynomial of [Definition 7](#) is a nonzero polynomial in the retained coordinates, hence cannot vanish at every point of the box. A point where this product does not vanish has $\Delta_m^{\text{right}} \neq 0$; the same nonvanishing also gives $\gamma \neq 0$, pairwise distinct slopes, and retained source cumulants realized by centered non-Gaussian real source laws with finite L -th moment, so the point lies in $F_{m,L}^{\text{right}}$ by [Definition 10](#) and outside E^{right} . The coordinate-reversed argument gives the reverse feasible point. Therefore

$$F_{m,L}^{\text{right}} \setminus E^{\text{right}} \neq \emptyset, \quad F_{m,L}^{\text{left}} \setminus E^{\text{left}} \neq \emptyset.$$

3. **The lower-order kernel identity.** Let $\theta \in F_{m,L}^{\text{right}} \setminus E^{\text{right}}$ and read θ as a complex parameter; the cumulant maps have real polynomial coordinates, so complexification changes no value. From $\Delta_m^{\text{right}}(\theta) \neq 0$,

$$\rho_i \neq 0 \ (1 \leq i \leq m), \quad \Pi_r(\theta) \neq 0,$$

and feasibility gives $\gamma \neq 0$ together with pairwise distinct $\gamma, \rho_1, \dots, \rho_m$; so the $m+1$ finite forward loading directions $u_0 = (1, \gamma)$, $u_j = (1, \rho_j)$ have distinct nonzero second coordinates, and together with $u_{m+1} = (0, 1)$ the n directions are pairwise non-proportional. Now let q be homogeneous of degree n with

$$q(\partial)f_{n+k} = 0 \quad \text{for every } 0 \leq k \leq m-1.$$

Dividing the contraction display by the nonzero factorial and applying the selected rows defining N^{right} , the vector $e = (q(u_j))_{j=0}^{m+1}$ is annihilated by $N^{\text{right}}(\theta)$, whose determinant is nonzero; hence $q(u_j) = 0$ for every j . A binary form of degree n vanishing at n pairwise non-proportional directions is divisible by the corresponding product of linear forms, so, comparing degrees, $q = c Q_{\text{right}}$ for some $c \in \mathbb{C}$, where

$$Q_{\text{right}} = x(\gamma x - y) \prod_{j=1}^m (\rho_j x - y)$$

is the forward support annihilator. Conversely every multiple of Q_{right} vanishes at all directions and hence annihilates every block. Thus

$$[q(\partial)f_{n+k} = 0 \text{ for every } 0 \leq k \leq m-1] \iff \exists c \in \mathbb{C} : q = c Q_{\text{right}}.$$

4. **Forward generic separation.** Suppose some $\eta = (\delta, \sigma, d) \in F_{m,L}^{\text{left}}$ satisfied

$$\Phi_{m,L}^{\text{right}}(\theta) = \Phi_{m,L}^{\text{left}}(\eta),$$

so that all divided-power blocks of the two parameters coincide. Let

$$Q_{\text{left}} = -y \prod_{j=1}^m (x - \sigma_j y) (x - \delta y)$$

be the reverse support annihilator of η , homogeneous of degree n and nonzero. It vanishes at every reverse direction, so it annihilates every reverse block of orders $n, \dots, L$, hence, by the displayed equality, every forward block f_{n+k} , $0 \leq k \leq m-1$. The lower-order kernel identity for θ gives $c \in \mathbb{C}$ with

$$Q_{\text{left}} = c Q_{\text{right}}.$$

If $c = 0$ then $Q_{\text{left}} = 0$, which is false. If $c \neq 0$, then $y \mid Q_{\text{left}}$ – the reverse family contains $v_0 = (1, 0)$, contributing the factor $-y$ – forces $y \mid Q_{\text{right}}$; but no factor of Q_{right} is divisible by y , since $\gamma \neq 0$ and every $\rho_j \neq 0$. Both cases are impossible, so no such η exists. This proves the forward generic-separation clause with the theorem-local exceptional set E^{right} .

5. **Reverse generic separation.** Let $\eta \in F_{m,L}^{\text{left}} \setminus E^{\text{left}}$. The coordinate-reversed exceptional-set conditions give $\sigma_i \neq 0$ for all i and $\Pi_1(\eta) \neq 0$, while feasibility gives $\delta \neq 0$ and pairwise distinct $\delta, \sigma_1, \dots, \sigma_m$; hence the $m+1$ finite reverse directions $v_j = (\sigma_j, 1)$, $v_{m+1} = (\delta, 1)$ have distinct nonzero first coordinates and, with $v_0 = (1, 0)$, the n reverse directions are pairwise non-proportional. Applying the coordinate-reversed version of the kernel argument gives

$$[q(\partial)f_{n+k} = 0 \text{ for every } 0 \leq k \leq m-1] \iff \exists c \in \mathbb{C} : q = c Q_{\text{left}}, \quad t = \Phi_{m,L}^{\text{left}}(\eta).$$

If some $\theta \in F_{m,L}^{\text{right}}$ satisfied $\Phi_{m,L}^{\text{left}}(\eta) = \Phi_{m,L}^{\text{right}}(\theta)$, its forward support annihilator Q_{right} would annihilate the reverse blocks and hence equal $c Q_{\text{left}}$; $c = 0$ is impossible since $Q_{\text{right}} \neq 0$, and $c \neq 0$ transfers the divisibility $x \mid Q_{\text{right}}$ (the forward family contains $u_{m+1} = (0, 1)$, contributing the factor x) to $x \mid Q_{\text{left}}$, contradicting $\delta \neq 0$ and $\sigma_j \neq 0$. Hence no such θ exists.

6. **The information order.** The exceptional-set construction and the two separation arguments above establish, for each arrow, a proper real algebraic exceptional subset whose complement meets the corresponding feasible region and off which the opposite arrow admits no representation with the same truncated cumulants; that is exactly the generic real separation criterion of [Definition 11](#) at truncation order $L = 2m+1$. Since $m \geq 3$ we have $2 \leq 2m+1$, so $2m+1$ belongs to the set of admissible orders whose infimum defines $K^*(m)$, whence

$$K^*(m) \leq 2m+1.$$

If $K^*(m)$ were equal to $2m+2$, the displayed inequality would read $2m+2 \leq 2m+1$, which is false; hence

$$K^*(m) \neq 2m+2.$$

□

Proof of [Theorem 1](#). Put $K = 2m+2$ and $n = m+2$. For a direction $u = (u_1, u_2) \in \mathbb{C}^2$ write

$$\ell_u(x, y) = u_1 x + u_2 y, \quad q(u) = q(u_1, u_2) \text{ the evaluation of a binary form } q \text{ at } u,$$

and, for a loading family $(w_j)_{j=0}^{m+1}$, write its support annihilator in the normalization of [Definition 12](#) as the degree- n form

$$Q_D = \prod_{j=0}^{m+1} (w_{j2} x - w_{j1} y),$$

one linear factor per loading direction, each vanishing exactly on that direction. Throughout, $\Theta_{m,K}^{\text{right}}$ and $\Theta_{m,K}^{\text{left}}$ are the finite retained-coordinate parameter spaces of [Definition 3](#), in which every source weight outside the orders $2 \leq r \leq K$ is zero.

1. **Contraction of a divided-power block.** For a forward parameter $\theta = (\gamma, \rho, c)$ and $2 \leq r \leq K$, expanding the binomial in $\ell_{u_j}^r$ and comparing with [Definition 4](#) gives

$$f_r(x, y) = \sum_{a=0}^r \binom{r}{a} [\Phi_{m,K}^{\text{right}}(\theta)]_{r,a} x^{r-a} y^a = \sum_{j=0}^{m+1} c_{jr} \ell_{u_j}(x, y)^r,$$

and the same identity with d_{jr} , v_j , and $\Phi_{m,K}^{\text{left}}$ follows from [Definition 5](#). If q is a homogeneous binary form of degree n , then differentiating a power of a linear form gives, for every $k \geq 0$,

$$q(\partial) \ell_u^{n+k} = \frac{(n+k)!}{k!} q(u) \ell_u^k.$$

Consequently, for $0 \leq k \leq m$,

$$q(\partial) f_{n+k} = \frac{(n+k)!}{k!} \sum_{j=0}^{m+1} c_{j,n+k} q(u_j) \ell_{u_j}^k,$$

and the scalar $(n+k)!/k!$ is nonzero.

2. **The contraction minor and the loci.** For a forward parameter define the $(m+2) \times (m+2)$ matrix $M^{\text{right}}(\theta)$ with rows indexed by $i = 0, \dots, m+1$ and columns by the source labels $j = 0, \dots, m+1$,

$$M^{\text{right}}(\theta)_{0,j} = c_{j,n}, \quad M^{\text{right}}(\theta)_{a+1,j} = \binom{m}{a} c_{j,K} (u_{j1})^{m-a} (u_{j2})^a \quad (0 \leq a \leq m),$$

and let $M_r(\theta) = \det M^{\text{right}}(\theta)$, a polynomial in the coordinates of θ . These rows are the degree-zero block and the coefficients of the degree- m block in the weighted contraction map

$$e = (e_j)_{j=0}^{m+1} \mapsto \left(\sum_{j=0}^{m+1} c_{j,n+k} e_j \ell_{u_j}^k \right)_{k=0}^m.$$

Thus nonsingularity of this selected square coefficient matrix forces the displayed weighted contraction map to be injective.

For a reverse parameter define instead

$$M^{\text{left}}(\eta)_{0,j} = d_{j,n}, \quad M^{\text{left}}(\eta)_{a+1,j} = \binom{m}{a} d_{j,K} (v_{j1})^a (v_{j2})^{m-a} \quad (0 \leq a \leq m),$$

and let $M_l(\eta) = \det M^{\text{left}}(\eta)$. The last $m+1$ rows are the coefficients of $x^a y^{m-a}$ in the degree- m reverse contraction block, so $M_l(\eta) \neq 0$ gives injectivity of the corresponding reverse weighted contraction map.

Each of M_r and M_l is a nonzero polynomial. At the forward parameter

$$\theta^*: \quad \gamma = 1, \quad \rho_i = i+1 \ (1 \leq i \leq m), \quad c_{j,n} = \mathbb{K}_{\{j=m+1\}}, \quad c_{j,K} = \mathbb{K}_{\{j \leq m\}}, \quad c_{jr} = 0 \text{ otherwise,}$$

the matrix $M^{\text{right}}(\theta^*)$ has last column $(1, 0, \dots, 0)^\top$ and, in its remaining columns and last $m+1$ rows, the entries $\binom{m}{a} (j+1)^a$; it is therefore triangular with a Vandermonde determinant in the distinct nodes $1, \dots, m+1$ as its surviving factor, so $M_r(\theta^*) \neq 0$. All weights of θ^* vanish outside the orders $2 \leq r \leq K$, so M_r remains a nonzero polynomial after every off-band weight coordinate is set to 0. The reverse witness has $\delta = m+1$, $\sigma_i = i+1$, $d_{0,n} = 1$, $d_{j,K} = 1$ for $j \neq 0$, and all other retained witness weights zero; the same Vandermonde calculation with the displayed reverse minor gives the corresponding statements for M_l .

3. **The cut-out polynomials.** Let Δ_{gen}^r and Δ_{gen}^l be the forward and reverse generic polynomials from [Definition 7](#), namely

$$\Delta_{\text{gen}}^r(\gamma, \rho, c) = \gamma \prod_{i=1}^m (\gamma - \rho_i) \prod_{1 \leq i < i' \leq m} (\rho_i - \rho_{i'}) \prod_{j=0}^{m+1} \prod_{r=2}^K c_{jr},$$

and

$$\Delta_{\text{gen}}^l(\delta, \sigma, d) = \delta \prod_{i=1}^m (\delta - \sigma_i) \prod_{1 \leq i < i' \leq m} (\sigma_i - \sigma_{i'}) \prod_{j=0}^{m+1} \prod_{r=2}^K d_{jr}.$$

Put $\Delta_\rho = \prod_{i=1}^m \rho_i$ in the forward chart and $\Delta_\sigma = \prod_{i=1}^m \sigma_i$ in the reverse chart, and define

$$R_r = \Delta_{\text{gen}}^r \Delta_\rho M_r, \quad R_l = \Delta_{\text{gen}}^l \Delta_\sigma M_l.$$

Each factor is a nonzero polynomial that stays nonzero after the off-band weight coordinates are set to 0: this is immediate for Δ_{gen}^r , Δ_{gen}^l , Δ_ρ , and Δ_σ , since they involve only slopes and retained weights, and it follows for the determinants from the witness calculations above. Hence R_r and R_l have the same two properties.

4. **The loci U_m^{right} and U_m^{left} .** Set

$$U_m^{\text{right}} = \{\theta \in \Theta_{m,K}^{\text{right}} : R_r(\theta) \neq 0\}, \quad U_m^{\text{left}} = \{\eta \in \Theta_{m,K}^{\text{left}} : R_l(\eta) \neq 0\}.$$

Inside the retained-coordinate space, the nonvanishing locus of a polynomial that is not identically zero there is the complement of a proper relatively closed set, hence relatively Zariski open, and it is relatively Zariski dense because any polynomial vanishing on it multiplies R_r , respectively R_l , to the zero polynomial and therefore vanishes identically. Since Δ_{gen}^r divides R_r , we get $U_m^{\text{right}} \subseteq \Theta_{m,K}^{\text{right},\circ}$; likewise Δ_{gen}^l divides R_l , so $U_m^{\text{left}} \subseteq \Theta_{m,K}^{\text{left},\circ}$.

5. **Open neighborhoods meeting feasibility.** Set

$$O^{\text{right}} = \{\theta \in \mathbb{R}^{m+1} \times \mathbb{R}^{(m+2)(K-1)} : R_r(\theta_{\mathbb{C}}) \neq 0\}, \quad O^{\text{left}} = \{\eta \in \mathbb{R}^{m+1} \times \mathbb{R}^{(m+2)(K-1)} : R_l(\eta_{\mathbb{C}}) \neq 0\},$$

with $\theta_{\mathbb{C}}$ and $\eta_{\mathbb{C}}$ the coordinatewise complexifications of [Definition 12](#). These are Euclidean open because $\theta \mapsto R_r(\theta_{\mathbb{C}})$ and $\eta \mapsto R_l(\eta_{\mathbb{C}})$ are continuous. The truncated moment problem supplies a real cumulant vector c^0 and a radius $\varepsilon > 0$ such that

$$|c'_r - c_r^0| < \varepsilon \quad (2 \leq r \leq K) \implies \exists \text{ a centered non-Gaussian law } \nu \text{ with } \mathbb{E}|\cdot|^K < \infty \text{ and } \text{Cum}_r(\nu) = c'_r \quad (2 \leq r \leq K).$$

Consider the box in which each slope coordinate ranges over all of $\mathbb{R}$ and each retained weight coordinate c_{jr} ranges over $(c_r^0 - \varepsilon, c_r^0 + \varepsilon)$. Each factor of the box is infinite, and R_r is a nonzero polynomial after pinning the off-band coordinates, so it cannot vanish at every point of the box; pick a point θ where it does not vanish. Then $\Delta_{\text{gen}}^r(\theta) \neq 0$ gives $\gamma \neq 0$ and pairwise distinct slopes, and each weight family c_j lies in the ε -box, hence is realized by a centered non-Gaussian real source law with finite K -th moment. Thus $\theta \in F_{m,K}^{\text{right}}$ by [Definition 10](#), and $R_r(\theta_{\mathbb{C}}) \neq 0$ puts $\theta \in O^{\text{right}}$. Hence $O^{\text{right}} \cap F_{m,K}^{\text{right}} \neq \emptyset$, and the same argument with R_l gives $O^{\text{left}} \cap F_{m,K}^{\text{left}} \neq \emptyset$. Finally, if $\theta \in O^{\text{right}} \cap F_{m,K}^{\text{right}}$, then $\theta_{\mathbb{C}}$ has all off-band weights zero because feasibility pins them, and $R_r(\theta_{\mathbb{C}}) \neq 0$, so $\theta_{\mathbb{C}} \in U_m^{\text{right}}$; likewise on the reverse side:

$$(O^{\text{right}} \cap F_{m,K}^{\text{right}})_{\mathbb{C}} \subseteq U_m^{\text{right}}, \quad (O^{\text{left}} \cap F_{m,K}^{\text{left}})_{\mathbb{C}} \subseteq U_m^{\text{left}}.$$

6. **What a point of U_m^{right} provides.** Let $\theta = (\gamma, \rho, c) \in U_m^{\text{right}}$. Factoring $R_r(\theta) \neq 0$ gives

$$\Delta_{\text{gen}}^r(\theta) \neq 0, \quad \rho_i \neq 0 \ (1 \leq i \leq m), \quad M_r(\theta) \neq 0.$$

Hence θ is generic, $\gamma \neq 0$, the second coordinates of the $m + 1$ finite loading directions,

$$u_{02} = \gamma, \quad u_{j2} = \rho_j \ (1 \leq j \leq m),$$

are pairwise distinct and all nonzero, and the forward weighted contraction map displayed above is injective.

7. **The common contraction kernel is the line $\mathbb{C}Q_D$.** Let D be the forward loading-direction support of θ and

$$Q_D = (\gamma x - y) \prod_{j=1}^m (\rho_j x - y) \cdot x,$$

the last factor coming from $u_{m+1} = (0, 1)$. If $q = cQ_D$ is homogeneous of degree n , then $q(u_j) = 0$ for every j , so every contraction $q(\partial)f_{n+k}$, $0 \leq k \leq m$, vanishes by the divided-power contraction identity. Conversely, suppose q is homogeneous of degree n and $q(\partial)f_{n+k} = 0$ for all $0 \leq k \leq m$. Dividing the contraction identity by the nonzero factorial, the vector $e = (q(u_j))_{j=0}^{m+1}$ is sent to 0 by the forward weighted contraction map. Injectivity therefore gives

$$q(u_j) = 0 \quad \text{for every } j = 0, \dots, m+1.$$

The $n = m+2$ directions $u_0, \dots, u_{m+1}$ are pairwise non-proportional, because their second coordinates $\gamma, \rho_1, \dots, \rho_m$ are distinct and nonzero and u_{m+1} is the vertical direction. A degree- n binary form vanishing at these n pairwise non-proportional directions is divisible by the corresponding product of linear forms, i.e. by Q_D ; comparing degrees gives $q = cQ_D$ for some $c \in \mathbb{C}$. Thus, for every homogeneous binary form q of degree n ,

$$[q(\partial)f_{n+k} = 0 \text{ for every } 0 \leq k \leq m] \iff \exists c \in \mathbb{C} : q = cQ_D.$$

Moreover $Q_D \neq 0$, every factor being a nonzero linear form; $x \mid Q_D$ by the displayed factorization; and $y \nmid Q_D$, since the factor x is not divisible by y , and the factors $\gamma x - y$, $\rho_j x - y$ are not because $\gamma \neq 0$ and $\rho_j \neq 0$.

8. **The reverse fiber over $\Phi_{m,K}^{\text{right}}(\theta)$ is empty.** Suppose $\eta = (\delta, \sigma, d) \in R_{m,K}^{\text{left}}(\Phi_{m,K}^{\text{right}}(\theta))$. Fiber membership in [Definition 8](#) compares the two maps on the retained coordinates, and both cumulant maps vanish off that range, so

$$\Phi_{m,K}^{\text{left}}(\eta) = \Phi_{m,K}^{\text{right}}(\theta), \quad \text{hence} \quad f_r^{\text{left}} = f_r \quad \text{for all } r.$$

Let $Q_D^{\text{left}} = \prod_j (v_{j2}x - v_{j1}y)$ be the support annihilator of the reverse loading family of η . It is homogeneous of degree n , it is nonzero, and it vanishes at every reverse direction. Therefore it annihilates all reverse blocks f_{n+k}^{left} , $0 \leq k \leq m$, by the divided-power contraction identity, and hence annihilates all forward blocks f_{n+k} . The forward kernel identity then yields $c \in \mathbb{C}$ with

$$Q_D^{\text{left}} = cQ_D.$$

If $c = 0$, this makes $Q_D^{\text{left}} = 0$, a contradiction. If $c \neq 0$, then $y \mid Q_D^{\text{left}}$, because the reverse family contains $v_0 = (1, 0)$ and therefore contributes the factor $-y$; the displayed scalar multiple then forces $y \mid Q_D$, contradicting the forward divisibility conclusion above. Hence

$$R_{m,K}^{\text{left}}(\Phi_{m,K}^{\text{right}}(\theta)) = \emptyset.$$

9. **Recovery of the slope multiset.** Put

$$Q(z) = (z - \gamma) \prod_{i=1}^m (z - \rho_i).$$

The entries $\gamma, \rho_1, \dots, \rho_m$ are pairwise distinct and nonzero, so $Q \neq 0$, Q is squarefree, $Q(0) = (-1)^{m+1} \gamma \prod_i \rho_i \neq 0$, and the multiset of roots of Q is exactly $\{\gamma, \rho_1, \dots, \rho_m\}$. Up to sign, Q is the dehomogenization $Q_D(1, z)$: the factors $\gamma x - y$ and $\rho_j x - y$ become $\gamma - z$ and $\rho_j - z$, and the factor x becomes 1. Now let θ'' be a generic forward parameter with

$$\Phi_{m,K}^{\text{right}}(\theta'') = \Phi_{m,K}^{\text{right}}(\theta),$$

and let Q_D'' be its support annihilator. Its own support annihilator annihilates its own divided-power blocks, hence the blocks of θ ; the forward kernel identity gives $Q_D'' = c Q_D$ with $c \neq 0$, because $Q_D'' \neq 0$. Dehomogenizing at $x = 1$ gives $Q_D''(1, z) = c Q_D(1, z)$, so the two dehomogenizations have the same root multiset. Therefore

$$\{\gamma'', \rho_1'', \dots, \rho_m''\} = \{\gamma, \rho_1, \dots, \rho_m\}.$$

10. **The reverse locus.** The argument is the same after exchanging the two coordinate axes, hence the roles of x and y and of the two arrows. For $\eta = (\delta, \sigma, d) \in U_m^{\text{left}}$, factoring $R_1(\eta) \neq 0$ gives that η is generic, $\delta \neq 0$, all $\sigma_i \neq 0$, the first coordinates

$$v_{j1} = \sigma_j \quad (1 \leq j \leq m), \quad v_{m+1,1} = \delta$$

of the finite reverse loading directions are pairwise distinct and nonzero, and the reverse weighted contraction map is injective. Writing Q_D^{left} for the support annihilator of the reverse loading family of η ,

$$Q_D^{\text{left}} = y \cdot \prod_{j=1}^m (x - \sigma_j y) \cdot (x - \delta y)$$

up to sign, the first factor coming from $v_0 = (1, 0)$. Hence $Q_D^{\text{left}} \neq 0$, $y \mid Q_D^{\text{left}}$, and $x \nmid Q_D^{\text{left}}$, because $\delta \neq 0$ and all $\sigma_i \neq 0$. Repeating the divided-power contraction and interpolation argument with the reverse map gives the reverse kernel identity, for $t = \Phi_{m,K}^{\text{left}}(\eta)$ and f_r built from t :

$$[q(\partial)f_{n+k} = 0 \text{ for every } 0 \leq k \leq m] \iff \exists c \in \mathbb{C} : q = c Q_D^{\text{left}}.$$

A forward parameter $\theta \in R_{m,K}^{\text{right}}(\Phi_{m,K}^{\text{left}}(\eta))$ would have its own support annihilator Q_D^{right} in that kernel, hence $Q_D^{\text{right}} = c Q_D^{\text{left}}$ with $c \neq 0$. But $x \mid Q_D^{\text{right}}$, because the forward family contains $u_{m+1} = (0, 1)$, while $x \nmid Q_D^{\text{left}}$, a contradiction. Therefore

$$R_{m,K}^{\text{right}}(\Phi_{m,K}^{\text{left}}(\eta)) = \emptyset.$$

Finally $Q(z) = (z - \delta) \prod_{i=1}^m (z - \sigma_i)$ is nonzero, squarefree, satisfies $Q(0) \neq 0$, and has root multiset $\{\delta, \sigma_1, \dots, \sigma_m\}$. Dehomogenizing the reverse kernel identity at $y = 1$ shows that every generic reverse parameter with the same cumulants has that same multiset.

□

References

- Saugata Basu, Richard Pollack, and Marie-Francoise Roy. *Algorithms in Real Algebraic Geometry*, volume 10 of *Algorithms and Computation in Mathematics*. Springer, 2 edition, 2006. doi: 10.1007/3-540-33099-2.
- Jacek Bochnak, Michel Coste, and Marie-Francoise Roy. *Real Algebraic Geometry*, volume 36 of *Ergebnisse der Mathematik und ihrer Grenzgebiete*. Springer, 1998.
- David R. Brillinger. The calculation of cumulants via conditioning. *Annals of the Institute of Statistical Mathematics*, 21:215–218, 1969.
- Jean-Francois Cardoso. High-order contrasts for independent component analysis. *Neural Computation*, 11(1):157–192, 1999. doi: 10.1162/089976699300016863.
- Wei Chen, Wanyang Gu, Linjun Peng, Ruichu Cai, Zhifeng Hao, and Kun Zhang. Horizontal and vertical federated causal structure learning via higher-order cumulants. *arXiv preprint arXiv:2507.06888*, 2025a.
- Wei Chen, Linjun Peng, Zhiyi Huang, Haoyue Dai, Zhifeng Hao, Ruichu Cai, and Kun Zhang. Identification of causal direction under an arbitrary number of latent confounders. *arXiv preprint arXiv:2510.22711*, 2025b.
- Pierre Comon. Independent component analysis, a new concept? *Signal Processing*, 36(3):287–314, 1994. doi: 10.1016/0165-1684(94)90029-9.
- Pierre Comon and Bernard Mourrain. Decomposition of quantics in sums of powers of linear forms. *Signal Processing*, 53(2–3):93–107, 1996. doi: 10.1016/S0165-1684(96)00097-9.
- Patrik O. Hoyer, Shohei Shimizu, Antti J. Kerminen, and Markus Palviainen. Estimation of causal effects using linear non-gaussian causal models with hidden variables. *International Journal of Approximate Reasoning*, 49(2):362–378, 2008. doi: 10.1016/j.ijar.2008.02.006.
- A. Hyvärinen and E. Oja. Independent component analysis: algorithms and applications. *Neural Networks*, 13(4–5):411–430, 2000. doi: 10.1016/S0893-6080(00)00026-5.
- Tamara G. Kolda and Brett W. Bader. Tensor decompositions and applications. *SIAM Review*, 51(3):455–500, 2009. doi: 10.1137/07070111X.
- Joseph M. Landsberg. *Tensors: Geometry and Applications*. American Mathematical Society, 2012. ISBN 9780821869079.
- Takashi Nicholas Maeda and Shohei Shimizu. Causal discovery of linear non-gaussian acyclic models in the presence of latent confounders. *arXiv preprint arXiv:2001.04197*, 2020.

- Peter McCullagh. *Tensor Methods in Statistics*. Chapman and Hall, 1987.
- Geert Mesters and Piotr Zwiernik. Non-independent components analysis. *arXiv preprint arXiv:2206.13668*, 2022.
- Yoshimitsu Morinishi and Shohei Shimizu. Differentiable causal discovery of linear non-gaussian acyclic models under unmeasured confounding. *arXiv preprint arXiv:2501.12854*, 2025.
- Judea Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, 2 edition, 2009. ISBN 9780521895606.
- Jonas Peters, Dominik Janzing, and Bernhard Schölkopf. *Elements of Causal Inference: Foundations and Learning Algorithms*. MIT Press, 2017. ISBN 9780262037310.
- Thomas S. Richardson and Peter Spirtes. Ancestral graph markov models. *The Annals of Statistics*, 30(4):962–1030, 2002. doi: 10.1214/aos/1031689015.
- Daniela Schkoda, Elina Robeva, and Mathias Drton. Causal discovery of linear non-gaussian causal models with unobserved confounding. *arXiv preprint arXiv:2408.04907*, 2024.
- Shohei Shimizu, Patrik O. Hoyer, Aapo Hyvärinen, and Antti Kerminen. A linear non-gaussian acyclic model for causal discovery. *Journal of Machine Learning Research*, 7:2003–2030, 2006. URL <https://www.jmlr.org/papers/v7/shimizu06a.html>.
- Peter Spirtes, Clark Glymour, and Richard Scheines. *Causation, Prediction, and Search*. MIT Press, 2 edition, 2001.
- Tatsuya Tashiro, Shohei Shimizu, Aapo Hyvärinen, and Takashi Washio. Parcelingam: A causal ordering method robust against latent confounders. *Neural Computation*, 26(1):57–83, 2014.
- Daniele Tramontano, Yaroslav Kivva, Saber Salehkaleybar, Mathias Drton, and Negar Kiyavash. Causal effect identification in lingam models with latent confounders. *arXiv preprint arXiv:2406.02049*, 2024.
- Daniele Tramontano, Yaroslav Kivva, Saber Salehkaleybar, Mathias Drton, and Negar Kiyavash. Causal effect identification in lvingam from higher-order cumulants. In *Proceedings of the 42nd International Conference on Machine Learning*, 2025.
- Savi Virolainen. Identification by non-gaussianity in structural threshold and smooth transition vector autoregressive models. *arXiv preprint arXiv:2404.19707*, 2024.
- Feng Xie, Biwei Huang, Zhengming Chen, Ruichu Cai, Clark Glymour, Zhi Geng, and Kun Zhang. Generalized independent noise condition for estimating causal structure with latent variables. *arXiv preprint arXiv:2308.06718*, 2023.