Quantile Geometry: Finite-Rank Drift Ratios for Causal Discovery

Anonymous authors
Paper under double-blind review

Abstract

Unknown nonlinear transformations of an effect can obscure the structure used to distinguish cause from effect. We study the drift ratio, which measures how the effect's conditional quantiles move as the cause changes, relative to their local spacing. This quantity is unchanged by smooth increasing transformations of the effect. We introduce finite-rank drift-ratio (FRDR) models, which assume that the drift-ratio surface has finite rank, meaning that its variation is captured by finitely many shared functions of the quantile level. The framework includes additive noise models, post-nonlinear models, location-scale noise models, and post-nonlinear heteroscedastic noise models. Under suitable conditions, we prove generic identifiability of the causal direction in FRDR models: the reverse drift ratio has rank greater than any fixed finite bound, apart from exceptional distributions of the cause. We derive a quantile geometry (QG) score, QG-FRDR, that applies across these model classes, along with four specialized scores that exploit their additional structure. These scores require neither noise reconstruction nor estimation of the unknown transformation. Experiments demonstrate fast computation for QG-FRDR and strong directional accuracy for its weak formulation on structural-model simulations. The class-specific QG-ANM, QG-LSNM, and QG-PNL scores achieve high accuracy within their target classes.

1 Introduction

Bivariate causal discovery distinguishes $X \to Y$ from $Y \to X$ through assumptions that generally hold in only one direction (Spirtes et al., 2001). Additive noise models (ANMs) assume $Y = a(X) + \epsilon$, with $\epsilon \perp X$ (Hoyer et al., 2008); location–scale noise models (LSNMs) allow $Y = a(X) + b(X)\epsilon$, with $b > 0$ (Immer et al., 2023). Post-nonlinear models (PNLs) apply an increasing function g to an ANM, while post-nonlinear heteroscedastic noise models (PNL-HNMs) combine both extensions as $Y = g(a(X) + b(X)\epsilon)$ (Zhang & Hyvärinen, 2009; Chen et al., 2026; Kamimura & Pham, 2026).

Standard approaches fit a structural model and test whether the reconstructed noise is independent of the candidate cause (Hoyer et al., 2008; Zhang & Hyvärinen, 2009); alternatives use likelihood or distributional assumptions (Immer et al., 2023; Klippert & Marx, 2026). Fitting becomes harder when both a varying noise scale and an unknown transformation must be estimated. Such transformations arise from nonlinear measurement processes or a choice of response scale (Zhang & Hyvärinen, 2009; Box & Cox, 1964; Hothorn et al., 2014). Methods that directly model conditional distributions also face this issue, since the transformation changes the object being fitted (Xi et al., 2025; Chen et al., 2026; Kamimura & Pham, 2026).

We instead model the drift ratio $R(x,u) = \partial_x Q(x,u)/\partial_u Q(x,u)$, where $Q(x,u) = Q_{Y|X=x}(u)$. It measures how conditional quantiles move with the cause relative to their local spacing across quantile levels. For $\widetilde{Y} = h(Y)$ with $h' > 0$, both derivatives acquire the factor $h'(Q(x,u))$, which cancels from R. This invariant was introduced in the identifiability analysis of the low-rank quantile surface (LRQS) model (Kamimura & Pham, 2026); here, we make it the modeling object. We propose finite-rank drift-ratio (FRDR) models:

$$R(x,u) = \sum_{k=1}^{K} \alpha_k(x)\phi_k(u). \tag{1}$$

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

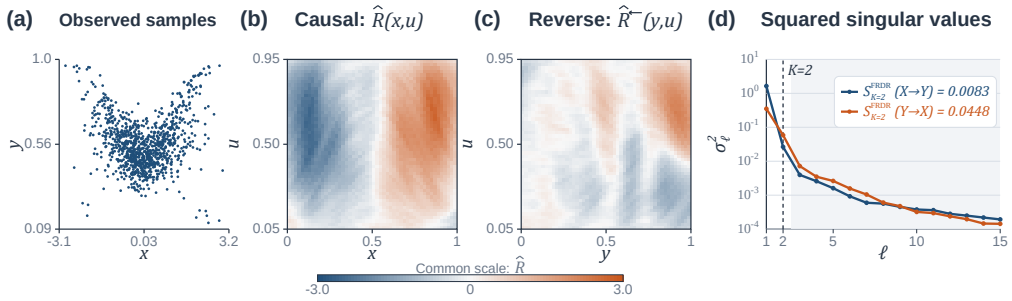

Figure 1: **Identifiability logic.** (a) Observed samples from a PNL-HNM $Y = g\big(0.4X^2 + (0.4 + X^4/15)\varepsilon\big)$, where g = sigmoid, and $X, \varepsilon \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0,1)$. (b) Estimated drift ratio $R(x,u)$ in the causal direction. The population drift ratio has rank 2. The drift ratio is invariant of the effect transformation g; another g will give the same population field. (c) Estimated reverse drift ratio $R^{\leftarrow}(y,u)$. The population reverse drift ratio generically has rank greater than 2. (d) The score $S^{\text{FRDR}}_{K=2}$, based on the sum of squared singular values beyond the first two, measures departure from surfaces of rank at most 2. Its population value is zero in the causal direction and generically positive in the reverse direction.

All curves $R(x,\cdot)$ therefore lie in the span of the same K functions of the quantile level. Unlike LRQS, FRDR imposes rank directly on R, without estimating a transformation. The rank-two class contains ANMs, LSNMs, PNLs, and PNL-HNMs.

Contributions. Our main result is a basis-free identifiability theorem for FRDR (Theorem 1): with the forward conditional distribution fixed, any reverse rank bound L restricts $\log p_X$ locally to a finite union of finite-parameter families, for every L and independently of the forward rank K. Unlike the fixed-basis results for QPE and multi-component LRQS, the theorem does not prescribe the reverse basis and therefore covers learned-basis implementations.

We further show that ANM, LSNM, PNL, and PNL-HNM are rank-two FRDR models with sharper population identities that hold causally and fail generically in reverse (Proposition 2). These results yield QG-FRDR, a weighted-SVD score; QG-FRDR-W, a weak-form rank-constrained fit that avoids differentiating the estimated CDF; and four class-specific scores. None requires residual reconstruction, independence testing, or estimation of an output transformation. We also show that zero population scores characterize model membership: on the working rectangle for ANM and LSNM, and locally under additional conditions for PNL and PNL-HNM (Appendix D.3).

On 48 synthetic settings (Section 8), QG-FRDR-W reaches 93.2% overall accuracy, second to LRQS est-g at 95.0% and similar cost, while QG-ANM, QG-LSNM, and QG-PNL reach 99.3–100% on their target families in about 0.01 s per pair. QG-FRDR is the fastest general-purpose method at 0.013 s per pair, but performs poorly on several mechanisms with small reverse rank-two gaps.

2 Related work

QPE (Chen et al., 2026) and LRQS (Kamimura & Pham, 2026) are the closest conditional-quantile approaches. QPE constrains a velocity field ψ in observed (x, y) coordinates. Its relation to the drift ratio is

$$R(x,u) = f\{Q(x,u) \mid x\}\, \psi\{x, Q(x,u)\}.$$

Both the density factor and the change to quantile coordinates can alter functional rank, so the QPE and FRDR rank constraints concern different fields. LRQS instead constrains the quantile surface after an increasing output transformation; its est-g implementation estimates that transformation jointly with the surface. FRDR imposes rank directly on the transformation-invariant field R. An LRQS representation need not yield a finite-rank drift ratio; Appendix A gives an example.

The reverse basis need not be specified. The finite-span identifiability result for QPE and the multi-component result for LRQS assume prescribed reverse basis functions for their respective modeling objects (Chen et al., 2026; Kamimura & Pham, 2026). These fixed-basis results do not establish identifiability for learned-span variants such as QPE-f lowrank and multi-component LRQS

fitting. In contrast, Theorem 1 generically excludes all reverse drift-ratio spans of dimension at most L, for each fixed L, without prescribing their basis functions. Appendix A gives a broader review.

3 Models

3.1 Notation and regularity

Let (X, Y) have continuous, strictly increasing marginal distribution functions. For the candidate direction $X \to Y$, write $F(y \mid x)$ and $f(y \mid x)$ for the conditional distribution and density. Subscripts denote partial derivatives. For the reverse direction, write $Q^{\leftarrow}(y, u) = Q_{X|Y=y}(u)$, $R^{\leftarrow}(y, u) = Q_y^{\leftarrow}(y, u)/Q_u^{\leftarrow}(y, u)$. For the reverse-rank theorem, we use z for the value of Y and write $\kappa(x, z) = \log p_{Y|X=x}(z)$ and $\nu(z) = \log p_Y(z)$. For each candidate direction, quantile-based population quantities are considered on a compact rectangle $I \times J_0$. Here, I is a compact interval of positive length contained in the interior of the candidate cause's support: that of X for $X \to Y$, and that of Y for $Y \to X$. The interval $J_0 = [u_{\text{lo}}, u_{\text{hi}}]$, with $0 < u_{\text{lo}} < u_{\text{hi}} < 1$, contains the quantile levels. Write μ_I and μ_{J_0} for the normalized Lebesgue measures on I and J_0, respectively.

Assumption 1. *On each compact interior region considered, the forward and reverse conditional quantile surfaces are sufficiently smooth, strictly increasing in the quantile level, and induce positive continuous conditional densities. All fields entering the scores are square integrable. The precise L-dependent conditions required by Theorem 1 are stated in Appendix C.1; marginal densities are assumed positive whenever copula-coordinate derivatives are used.*

The rank of a field: $\operatorname{rank} R = \dim \operatorname{span}\{R(x, \cdot) : x \in I\}$. Rank at most K means that every row is a linear combination of K common functions of u.

Genericity. We hold the forward conditional distribution fixed and vary the smooth cause density. For each fixed reverse rank bound, satisfying that bound restricts $\log p_X$ locally to a finite union of finite-parameter families. Generic identifiability below means identifiability outside the corresponding exceptional set of cause densities.

3.2 Finite-rank drift-ratio models

For a conditional quantile surface Q with $Q_u > 0$, write $R_Q(x, u) = Q_x(x, u)/Q_u(x, u)$, or simply R when Q is clear. For each integer $K \geq 1$, define the FRDR model $\mathcal{M}_K^R$ by $\mathcal{M}_K^R = \{Q : \operatorname{rank} R_Q \leq K\}$. Equivalently, $Q \in \mathcal{M}_K^R$ iff $R_Q(x, u) = \sum_{k=1}^K \alpha_k(x)\phi_k(u)$ for continuous coefficients α_k and continuously differentiable functions ϕ_k. The functions α_k and ϕ_k are not fixed in advance.

The model $\mathcal{M}_2^R$ contains all four structural models studied below. The PNL-HNM representation is

$$Y = g\{a(X) + b(X)\varepsilon\}, \qquad X \perp\!\!\!\perp \varepsilon, \qquad g' > 0, \quad b > 0. \tag{2}$$

If q is the quantile function of ε, then $Q(x, u) = g\{a(x) + b(x)q(u)\}$, and differentiation gives

$$R_Q(x, u) = \alpha(x)\frac{1}{q'(u)} + \beta(x)\frac{q(u)}{q'(u)}, \qquad \alpha = \frac{a'}{b}, \quad \beta = \frac{b'}{b}. \tag{3}$$

Thus PNL-HNM belongs to $\mathcal{M}_2^R$. Its named submodels are PNL with $b \equiv 1$, LSNM with $g = \text{id}$, and ANM with $g = \text{id}$, $b \equiv 1$. Setting $b \equiv 1$ in Eq. (3) removes the second term, so PNL and ANM have rank at most one; LSNM and PNL-HNM have rank at most two. Thus all four lie in $\mathcal{M}_2^R$. The converse does not hold: a general rank-two drift ratio need not have the shared-q form in Eq. (3).

The drift ratio can also be written directly in terms of the conditional distribution. Differentiating $F(Q(x, u) \mid x) = u$ with respect to x and u gives $F_x(Q \mid x) + f(Q \mid x)Q_x = 0$ and $f(Q \mid x)Q_u = 1$, hence

$$R_Q(x, u) = Q_x(x, u) f\{Q(x, u) \mid x\} = -F_x(y \mid x)\big|_{y=Q(x,u)}. \tag{4}$$

Equation (4) plays a central role in both the calculations of QG-FRDR and QG-FRDR-W.

4 Identifiability of FRDR

Fixing $p(y \mid x)$ fixes the forward drift ratio but leaves the cause density p_X unrestricted. In reverse, Bayes' rule gives $p(x \mid y) \propto p(y \mid x)p_X(x)$, so the drift ratio also depends on p_X. The theorem

below converts a reverse rank bound into differential constraints on $\log p_X$, without prescribing the reverse basis.

Theorem 1. *Fix integers $K, L \geq 1$ and a positive smooth forward conditional density whose conditional quantile surface belongs to $\mathcal{M}_K^R$. Suppose Assumption 1 and the L-dependent regularity conditions in Appendix C.1 hold. Assume that, for some z_0 in the interior of the reverse conditioning interval, $\kappa_{xz}(\cdot, z_0)$ is not identically zero on $\{Q^{\leftarrow}(z_0, u) : u \in J_0\}$. If $\operatorname{rank} R^{\leftarrow} \leq L$, then the restriction of $\log p_X$ to a smaller interior interval belongs to a finite union of finite-parameter families determined by systems of ordinary differential equations. Consequently, $\operatorname{rank} R^{\leftarrow} \geq L + 1$ generically. No relation between K and L is required.*

Proof sketch. If $\operatorname{rank} R^{\leftarrow} \leq L$, the rows and their first L derivatives in z lie in a space of dimension at most L. Nondegeneracy gives $R^{\leftarrow}(z_0, \cdot) \neq 0$. Hence, for some $1 \leq m \leq L$ and constants A_j,

$$\partial_z^m R^{\leftarrow}(z_0, \cdot) = \sum_{j=0}^{m-1} A_j \partial_z^j R^{\leftarrow}(z_0, \cdot).$$

To express this relation in terms of the cause density, let $H(x, z) = P(X \leq x \mid Y = z)$ and set $S = -H_z$, $w = 1/p_{X|Y=z}(x)$, and $c = \nu' - \kappa_z$. Differentiating $H\{Q^{\leftarrow}(z, u), z\} = u$ gives $Q_z^{\leftarrow} = Sw$ and $R^{\leftarrow}(z, u) = S\{Q^{\leftarrow}(z, u), z\}$, with Sw evaluated at $x = Q^{\leftarrow}(z, u)$. For any smooth $\Psi(x, z)$, the chain rule therefore gives

$$\partial_z \Psi\{Q^{\leftarrow}(z, u), z\} = (\Psi_z + Sw\Psi_x)\{Q^{\leftarrow}(z, u), z\}.$$

Repeatedly applying this identity computes the derivatives in the linear relation. Substituting $u = H(x, z_0)$ then writes that relation in x.

Bayes' rule gives $S_x = c/w$ and $w_z = cw$. For $m \geq 3$, the highest-derivative term in the resulting relation is $c_x S^m w^{m-2} \partial_x^{m-2} w$. Since $c_x = -\kappa_{xz}$, nondegeneracy provides an interval with $c_x S \neq 0$, where the relation determines $\partial_x^{m-2} w$. For $m = 1, 2$, it determines w algebraically, after one x-derivative when $m = 1$.

The remaining unknown functions are lower x-derivatives of w and $\partial_z^j S$, which satisfy $\partial_x(\partial_z^j S) = \partial_z^j(c/w)$ for $0 \leq j \leq m$. Using $w_z = cw$ removes the z-derivatives of w, leaving a finite system of ordinary differential equations in x at $z = z_0$. Its solutions depend on finitely many coefficients, initial values, and derivatives of ν at z_0. Finally,

$$\log p_X(x) = -\log w(x, z_0) - \kappa(x, z_0) + \nu(z_0).$$

The possibilities $m = 1, \ldots, L$ therefore give the claimed finite union of finite-parameter families. Appendix C gives the details. □

Remark 1 (Scope of the assumptions). Since $\kappa_{xz}(x, z) = \partial_x \partial_z \log p_{X,Y}(x, z)$, the nondegeneracy condition in Theorem 1 is the same as Assumption 2 of Kamimura & Pham (2026), with the nonzero point required to lie in the working quantile range. The exclusions derived in their Appendix C.4 also apply here: uniform-noise location–scale models, constant-scale exponential and Laplace noise models, and pure-scale power-law noise models, including their smooth increasing output transformations. These mechanisms fall outside the theorem, but need not be nonidentifiable. Our L-dependent regularity requirements are given in Appendix C.1.

5 Causal scores for FRDR

5.1 QG-FRDR

Let $\sigma_1 \geq \sigma_2 \geq \cdots$ be the singular values of R under μ_I and μ_{J_0}, and write $s_R^2 = \|R\|^2_{L^2(\mu_I \otimes \mu_{J_0})} = \sum_{\ell \geq 1} \sigma_\ell^2$. For $s_R > 0$, define

$$S_K^{\mathrm{FRDR}}(R) = \frac{\sum_{\ell > K} \sigma_\ell^2}{s_R^2}. \tag{5}$$

Set the score to zero when $s_R = 0$. The numerator is the squared L^2 distance to the set of fields of rank at most K. Thus $S_K^{\mathrm{FRDR}} \in [0, 1]$ measures the fraction of squared L^2 norm not captured by a rank-K approximation. We select the direction with the smaller score.

Proposition 1. *Under the assumptions of Theorem 1 with $L = K$, $S_K^{\mathrm{FRDR}}(R) = 0$ and $S_K^{\mathrm{FRDR}}(R^{\leftarrow}) > 0$ generically.*

Appendix C.3 gives the proof.

5.2 QG-FRDR-W

QG-FRDR estimates R and measures its departure from the rank bound. Noise in an unconstrained estimate can obscure the low-rank structure. QG-FRDR-W instead fits a rank-constrained field through integral identities, following the weak-form approach of Messenger & Bortz (2021). For a candidate direction $d = (A, B)$, use copula coordinates $T = F_A(A)$ and $V = F_B(B)$, and write $F_d(v \mid t)$ and $R_d(t, u)$ for the conditional CDF and drift ratio in these coordinates. Equation (4) gives $\partial_t F_d(v \mid t) = -R_d\{t, F_d(v \mid t)\}$. For a smooth test function w_m that vanishes at the endpoints of I, integration by parts gives

$$\int_I w_m'(t) F_d(v \mid t)\, dt = -\int_I w_m(t) \partial_t F_d(v \mid t)\, dt = \int_I w_m(t) R_d\{t, F_d(v \mid t)\}\, dt.$$

Only the known test function is differentiated; the conditional CDF is merely evaluated.

We fit $R_M(t, u) = \psi(t)^\top M \chi(u)$, where ψ and χ are spline bases and rank $M \le K$. Let $\mathcal{E}_d(M)$ denote the weighted mean squared discrepancy between the two integrals over test functions w_m and response thresholds v_ℓ, replacing F_d by its estimate $\widehat{F}_d$ on both sides and R_d by R_M on the right. The tests are constructed from localized functions of several widths to assess fit over shorter and longer intervals. After approximately minimizing $\mathcal{E}_d(M)$ under the rank constraint, we score the fitted matrix $\widehat{M}_d$ by

$$\widehat{S}_K^{\mathrm{FRDR,weak}}(d) = \frac{\mathcal{E}_d(\widehat{M}_d)}{\max\{\mathcal{E}_d(0), \varepsilon_{\mathrm{den}}\}}, \qquad \varepsilon_{\mathrm{den}} = 10^{-10}. \tag{6}$$

The denominator normalizes by the error of the zero field; the direction with the smaller score is selected. This finite-test criterion measures fit to the integral identities, not exact FRDR membership. Appendix G.1 specifies the tests, weights, and fitting procedure.

6 Class-specific causal scores

Each of the four structural classes above has drift-ratio rank at most two, so QG-FRDR and its variant apply to all of them. Their more specific functional forms, however, imply stronger geometric identities than the rank bound alone. We now turn these identities into class-specific causal scores; see Table 1 for a summary. Derivations and proofs are given in Appendix D. See Fig. 2 for a visualization for the cases of LSNM and PNL.

Model	Forward geometric identity	Score
FRDR	The drift ratio has rank at most K.	QG-FRDR in (5), QG-FRDR-W in (6)
ANM	$Q(x, u) - Q(x, v)$ is independent of x.	QG-ANM, (7)
LSNM	Interquantile-spacings are proportional.	QG-LSNM, (8)
PNL	Columns $R(\cdot, u)$ are positively proportional.	QG-PNL, (9)
PNL-HNM	G_{123} and W_{12} are linearly dependent.	QG-PNL-HNM, (10)

Table 1: Population invariants in the candidate direction $X \to Y$.

For functions in the relevant L^2 space, write $\mathcal{A}(f, g) := \|f\|^2\|g\|^2 - \langle f, g\rangle^2$. This Cauchy–Schwarz defect vanishes exactly when f and g are linearly dependent.

6.1 Scores from the quantile surface

ANM. Under $Q(x, u) = a(x) + q(u)$, every interquantile difference $Q(x, u) - Q(x, v) = q(u) - q(v)$ is independent of x; locally, this is equivalent to $Q_{xu} = 0$. We therefore define

$$S_{X\to Y}^{\mathrm{ANM}} := \int_{J_0^2} \mathrm{Var}_{x\sim\mu_I}[Q(x, u) - Q(x, v)]\ \mathrm{d}\mu_{J_0}^{\otimes 2}(u, v). \tag{7}$$

LSNM. For $u, v \in J_0$, let $\Delta_{uv} := Q(\cdot, v) - Q(\cdot, u) \in L^2(I)$. Under $Q(x, u) = a(x) + b(x)q(u)$, $\Delta_{uv} = \{q(v) - q(u)\}b$, so all spacing functions with a common reference level are proportional. Equivalently, $(\log Q_u)_{xu} = 0$. Define

$$S^{\mathrm{LSNM}}_{X\to Y} := \int_{J_0^3} \mathcal{A}\big(\Delta_{u_0u_1}, \Delta_{u_0u_2}\big)\, \mathrm{d}\mu^{\otimes 3}_{J_0}(u_0, u_1, u_2). \tag{8}$$

6.2 Scores from the drift ratio

PNL. Under $Q(x, u) = g\{a(x) + q(u)\}$,

$$R(x, u) = \frac{a'(x)}{q'(u)}.$$

Thus every nonzero column $R(\cdot, u)$ is a positive multiple of every other nonzero column. For nonzero $f, g \in L^2(I)$, let

$$\mathrm{pcol}(f, g) := 1 - \frac{\langle f, g\rangle_+^2}{\|f\|^2\|g\|^2}, \qquad \langle f, g\rangle_+ := \max\{\langle f, g\rangle, 0\},$$

and set $\mathrm{pcol}(f, g) = 0$ if either function is zero. The QG-PNL score is

$$S^{\mathrm{PNL}}_{X\to Y} := \int_{J_0^2} \mathrm{pcol}\{R(\cdot, u), R(\cdot, v)\}\, \mathrm{d}\mu^{\otimes 2}_{J_0}(u, v). \tag{9}$$

PNL-HNM. For $x_i \in I$, write $R_i = R(x_i, \cdot)$ and $W_{12} = R_1'R_2 - R_1R_2'$, where primes denote derivatives with respect to u. Substituting (3) gives

$$W_{12}(u) = \frac{\beta(x_1)\alpha(x_2) - \alpha(x_1)\beta(x_2)}{q'(u)}, \qquad R_i'(u) + \frac{q''(u)}{q'(u)}R_i(u) = \beta(x_i).$$

Hence $W_{12}' = -(q''/q')W_{12}$, and eliminating q''/q' yields $G_{123} := W_{12}R_3' - W_{12}'R_3 = \beta(x_3)W_{12}$. The QG-PNL-HNM score is:

$$S^{\text{PNL-HNM}}_{X\to Y} := \frac{1}{s_R^{10}} \int_{I^3} \mathcal{A}(G_{123}, W_{12})\, \mathrm{d}\mu^{\otimes 3}_I(x_1, x_2, x_3). \tag{10}$$

When $s_R = 0$, set $S^{\text{PNL-HNM}}_{X\to Y} := 0$. The reason we divide by s_R^{10} is to remove the global scaling of R, explained in Section 7.

Proposition 2 (Model-specific identification)**.** *If the forward mechanism is ANM, LSNM, PNL, or PNL-HNM, then its corresponding score in (7)–(10) is zero. Under the conditions of Theorem 1, the corresponding reverse score is generically positive, using $L = 1$ for ANM and PNL and $L = 2$ for LSNM and PNL-HNM.*

QG-ANM and QG-LSNM are exact membership scores on the working rectangle. QG-PNL and QG-PNL-HNM admit local converses on nondegenerate subrectangles where, respectively, the column proportionality factor is strictly positive or a pairwise Wronskian is nonvanishing; see Appendix D.3.

7 Practical computation of the scores

For a candidate direction $A \to B$, define $A^C = F_A(A)$ and $B^C = F_B(B)$. These are the coordinates T and V used by QG-FRDR-W. For the Q-based scores, also define $B^{\mathrm{std}} = \frac{B - \mathbb{E}[B]}{\sqrt{\mathrm{Var}(B)}}$, assuming $0 < \mathrm{Var}(B) < \infty$. Table 2 summarizes the estimation coordinates for each score.

For QG-ANM and QG-LSNM, an increasing reparameterization of A only relabels the conditional quantile curves, so it preserves the ANM and LSNM identities; using A^C fixes this otherwise arbitrary conditioning scale. The candidate effect is different: a nonlinear transformation of B generally destroys these Q-based identities, so B is only affinely standardized.

Such a transformation cancels from the drift ratio, allowing the R-based scores to use B^C:

Table 2: Estimation coordinates and fitted quantities for a candidate direction $A \to B$.

Score(s)	Data coordinates	Estimated or fitted quantity
QG-ANM, QG-LSNM	A^C, B^{std}	Conditional quantile surface $\widehat{Q}_{B^{\text{std}}\mid A^C}$
QG-PNL, QG-PNL-HNM, QG-FRDR	A^C, B^C	Drift-ratio field $\widehat{R}_{B^C\mid A^C}$
QG-FRDR-W	A^C, B^C	Conditional CDF $\widehat{F}_{B^C\mid A^C}$ and rank-constrained fit R_M

Lemma 1 (Marginal reparameterization). *For a candidate direction $A \to B$, let $\widetilde{A} = m(A)$ and $\widetilde{B} = h(B)$, where m and h are C^1 diffeomorphisms with strictly positive derivatives. Writing $a = m^{-1}(t)$, $R_{\widetilde{B}\mid\widetilde{A}}(t,u) = \frac{R_{B\mid A}(a,u)}{m'(a)}$. Hence a transformation of B leaves the drift ratio unchanged, while a transformation of A only relabels and positively rescales its rows. Consequently, marginal reparameterization preserves the rank of the drift-ratio field and the zero sets of the PNL and PNL-HNM identities.*

See Appendix E for a proof.

For the R-based scores, applying the lemma with $m = F_A$ and $h = F_B$ gives $R_{B^C\mid A^C}(t,u) = \frac{R_{B\mid A}(a,u)}{f_A(a)}$, with $a = F_A^{-1}(t)$. The positive row factor preserves the rank, PNL, and PNL-HNM zero conditions, but can change the size of a nonzero discrepancy. The copula coordinate A^C therefore puts both candidate causes on the same uniform scale for directional comparison. Empirically, we use $A_i^C = \text{rank}(A_i)/(n+1)$ and $B_i^C = \text{rank}(B_i)/(n+1)$, with average ranks for ties. These coordinates are unchanged by strictly increasing marginal transformations.

Finally, the unnormalized integral in the numerator of Eq. (10) is homogeneous of degree ten in R. Since $s_{cR} = |c|s_R$, division by s_R^{10} makes QG-PNL-HNM invariant to a global rescaling $R \mapsto cR$. QG-PNL and QG-FRDR are already scale-free and require no additional normalization.

Estimation and comparison. For the direct scores, we estimate the conditional CDF by Gaussian-kernel averaging and obtain conditional quantiles by inversion. We compute $\widehat{R}$ using Eq. (4), differentiating the kernel weights. Score integrals are approximated by weighted grid sums. For QG-FRDR, multiplying each row and column of the drift-ratio matrix by the square root of its quadrature weight gives the matrix used for singular value decomposition. The direction with the smaller score is selected. Algorithm 1 and Appendix G give the direct-score procedure; Appendix G.1 specifies the weak fit.

8 Experiments

8.1 Experimental setup

Synthetic datasets. We evaluate 48 settings grouped by their generating construction: two FRDR, 20 PNL-HNM, 11 LSNM, 12 PNL, and three ANM settings. Each contains 100 independently generated cause–effect pairs with $n = 1000$ observations per pair.

The two FRDR settings specify $R(x,u) = \sum_{k=1}^{K_0} \alpha_k(x)\phi_k(u)$, with Gaussian-process coefficient functions and $K_0 \in \{2,3\}$, and numerically construct Q from $Q_x = RQ_u$. The other settings combine Gaussian-process or analytic location–scale mechanisms with Gaussian, uniform, or standardized beta noise and identity, cubic, sigmoid, exponential, or hyperbolic-tangent output transformations. Uniform-noise settings are included even though they fall outside the nondegeneracy condition of Theorem 1. Appendix F gives the structural equations and generation procedure.

Cause–effect benchmarks. We use 19 groups from the evaluation suite distributed with QPE (Chen et al., 2026), comprising the SIM benchmarks and Tübingen cause–effect pairs (Mooij et al., 2016), the collections of Guyon et al. (2019), causal-velocity and constrained-QPE simulations (Xi et al., 2025; Chen et al., 2026), and DREAM4 in-silico gene-network benchmarks (Marbach et al., 2009).

We omit AN, AN-c, LS, LS-c, and MNU because their structural families are already represented in our synthetic evaluation. The retained groups contain 3,359 evaluable pairs. We compute accuracy within each group and average the 19 accuracies with equal weight. Appendix F.3 lists the groups and their results.

Compared methods. The general-purpose comparison includes QG-FRDR, its multiscale weak-fit version QG-FRDR-W, CVEL (Xi et al., 2025), and QPE-k and QPE-f lowrank (Chen et al., 2026). We also include the concurrent LRQS method (Kamimura & Pham, 2026), with an estimated output transformation (est-g) and with the transformation fixed to the identity (fix-g), because it provides a closely related comparison based on low-rank conditional quantile surfaces. These methods are evaluated on every setting in Tables 3 and 4.

We separately evaluate the model-specific QG-ANM, QG-LSNM, QG-PNL, and QG-PNL-HNM scores only on their corresponding synthetic families. On the ANM family, we also include the ANM baseline and DIVOT (Hoyer et al., 2008; Tu et al., 2022).

Implementation, metrics, and runtime. QG-FRDR uses a 50×50 conditioning–quantile grid and $u \in [0.01, 0.99]$. We set the drift-ratio rank bound to $K = 2$ for QG-FRDR and QG-FRDR-W on every dataset. For LRQS est-g and LRQS fix-g, we set the implementation's rank parameter to 2; this constrains a different object. The $K_0 = 3$ FRDR setting therefore tests deliberate misspecification of the rank bound used by the two FRDR scores.

The QG scores form their copula coordinates internally, whereas the baseline methods receive affine-standardized inputs when required by their implementations. For the QG scores, smaller values indicate a better fit. Accuracy records whether the inferred direction matches the known generating direction. Runtime is reported as mean wall-clock computation time per pair. QPE-f lowrank uses GPU acceleration, whereas all other methods use CPU computation. Further preprocessing details are given in Appendix F, and implementation and computational-environment details are provided in Appendix G.

8.2 Main synthetic results

Table 3 shows that QG-FRDR-W reaches 93.2% overall accuracy, compared with 75.2% for QG-FRDR, 72.3% for CVEL, and 72.4% for QPE-k; LRQS est-g reaches 95.0%. QG-FRDR-W has the highest accuracy on PNL and ANM, whereas LRQS est-g leads on PNL-HNM and LSNM.

The direct and weak formulations have complementary strengths. On the directly generated FRDR settings, QG-FRDR achieves 82.0% and 65.0% for $K_0 = 2$ and 3, versus 71.0% and 49.0% for QG-FRDR-W. Conversely, on the five analytic linear-location mechanisms and their transformed counterparts, QG-FRDR achieves only 0–4%, while QG-FRDR-W reaches 68–88% (Appendix F.2). For their Gaussian-noise versions, Appendix F.7 finds $S_2^{\mathrm{FRDR}}(R^{\leftarrow}) < 10^{-2}$. Together with the finite-sample diagnostics, these results indicate that direction-dependent estimation error can obscure the small population gaps. Theorem 1 establishes generic positivity, not a quantitative lower bound.

QG-FRDR is the fastest general-purpose method at 0.013 s per pair, versus 2.663 s for QG-FRDR-W. The model-specific QG-ANM, QG-LSNM, and QG-PNL scores achieve 99.3–100% accuracy on their target families in 0.008–0.011 s per pair; QG-LSNM reaches 99.5% on LSNM, compared with 98.7% for LRQS est-g. QG-PNL-HNM reaches only 72.0%, versus 93.4% for QG-FRDR-W on the same PNL-HNM settings, so its stronger population identity does not improve finite-sample accuracy here.

Sensitivity analysis. Appendix F.6.1 varies sample size, quantile range, grid resolution, and fitted rank on independently generated samples from the same 48 settings. Larger samples improve overall accuracy, while setting $K = 1$ causes the largest decline. The tested coarser grids and wider quantile range have little overall effect, although the family-level effects vary (Figure 3 and Table 12).

8.3 Cause–effect benchmark results

The QG-FRDR variants are less accurate on the heterogeneous benchmarks than on the synthetic families. QPE-k obtains the highest dataset-averaged accuracy, 69.1%, followed closely by LRQS

Table 3: Causal-direction accuracy (%) and mean wall-clock time per pair on 48 synthetic settings, with 100 pairs per setting. Panel A evaluates general-purpose methods on all settings; Panel B evaluates model-specific methods on their target families. Family accuracies average equally over settings within each family, and Overall averages equally over all 48 settings. Runtime averages over the pairs evaluated by each method. Both FRDR scores use drift-ratio rank bound $K = 2$; both LRQS implementations use rank parameter 2. Boldface marks the best value in each column of Panel A and within the ANM comparison in Panel B.

Panel A: General-purpose methods

Method	FRDR	PNL-HNM	LSNM	PNL	ANM	Overall	Time (s/pair)
QG-FRDR	**73.5**	74.0	54.7	91.8	92.7	75.2	**0.013**
QG-FRDR-W	60.0	93.4	90.3	**99.4**	**100.0**	93.2	2.663
LRQS est-g	55.0	**97.2**	**98.7**	95.3	93.0	**95.0**	2.425
LRQS fix-g	56.5	93.1	97.7	81.4	95.0	89.8	0.061
CVEL	49.0	77.8	64.5	84.3	32.3	72.3	2.168
QPE-k	51.0	69.7	97.9	51.0	96.7	72.4	0.061
QPE-f lowrank	35.0	64.7	60.9	49.3	62.7	58.6	26.658

Panel B: Model-specific methods

Method	Target family	Accuracy (%)	Time (s/pair)
QG-ANM	ANM	**100.0**	**0.008**
ANM	ANM	**100.0**	6.805
DIVOT	ANM	98.7	1.092
QG-LSNM	LSNM	99.5	0.011
QG-PNL	PNL	99.3	0.010
QG-PNL-HNM	PNL-HNM	72.0	0.103

Table 4: Dataset-averaged causal-direction accuracy and mean computation time per pair on 19 cause–effect benchmark groups. Accuracy is computed within each group and then averaged with equal group weights. Runtime is averaged over all 3,359 evaluated pairs.

Method	Accuracy (%)	Time (s/pair)
QG-FRDR	50.8	**0.011**
QG-FRDR-W	54.8	2.801
LRQS est-g	60.4	3.185
LRQS fix-g	68.7	0.047
CVEL	68.3	2.606
QPE-k	**69.1**	0.058
QPE-f lowrank	65.8	3.849

fix-g and CVEL. QG-FRDR and QG-FRDR-W reach 50.8% and 54.8%, respectively. Nevertheless, QG-FRDR remains the fastest method at 0.011 seconds per pair. The dataset-level results in Appendix F.3 show substantial variation across benchmark groups rather than a uniform loss.

9 Limitations and conclusion

Limitations. We consider bivariate models without hidden variables. The theory does not quantify score gaps or provide finite-sample guarantees. Accuracy depends on conditional-distribution estimation and the rank K; the weak score does not certify FRDR membership. **Conclusion.** We establish basis-free generic identifiability for FRDR, encompassing ANM, LSNM, PNL, and PNL-HNM, and derive causal scores without reconstructing noise or estimating an output transformation. These include QG-FRDR, its weak formulation, and four class-specific scores with exact or local model-membership characterizations. Experiments highlight the speed of direct scoring and the higher overall synthetic accuracy of the weak formulation. Quantifying score gaps and controlling direction-dependent estimation error remain important directions for future work.

AI USE STATEMENT

In this work, we used generative AI tools for helping develop theoretical models or conceptual frameworks, formulating mathematical claims, assisting in the writing of proofs, implementing methods, and assisting with translation. We have not used generative AI tools for generating synthetic data sets, providing critical ingredients for proving mathematical claims, proposing or refining hypotheses, designing or providing feedback on research methodology or experiments, cleaning and reformating dataset, supporting qualitative and thematic data analysis, or interpreting results. Additionally, we used generative AI tools for drafting parts of the paper, summarizing or analyzing existing literature, editing the paper to improve readability, and identifying relevant literature. We have reviewed all AI-assisted work. LLM-generated code was verified and tested for correctness. We take responsibility for the final content of this work, including text, claims or artifacts produced with the aid of generative AI.

REPRODUCIBILITY STATEMENT

Upon acceptance, we will release the complete implementation of our methods and all experimental code in a public repository. The repository will include the environment specifications, experimental configurations, and instructions needed to reproduce the reported results.

REFERENCES

David M. Bortz, Daniel A. Messenger, and Vanja Dukic. Direct estimation of parameters in ODE models using WENDy: Weak-form estimation of nonlinear dynamics. *Bulletin of Mathematical Biology*, 85:110, 2023. doi: 10.1007/s11538-023-01208-6.

G. E. P. Box and D. R. Cox. An analysis of transformations. *Journal of the Royal Statistical Society: Series B (Methodological)*, 26(2):211–243, 07 1964. ISSN 0035-9246. doi: 10.1111/j.2517-6161.1964.tb00553.x. URL https://doi.org/10.1111/j.2517-6161.1964.tb00553.x.

Yikang Chen, Xingzhe Sun, and Dehui Du. Causal discovery via quantile partial effect. In *The Fourteenth International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=80vdaC5DsD.

Isabelle Guyon, Alexander Statnikov, and Berna Bakir Batu. *Cause effect pairs in machine learning*. Springer, 2019.

Torsten Hothorn, Thomas Kneib, and Peter Bühlmann. Conditional transformation models. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 76(1):3–27, 01 2014. ISSN 1369-7412. doi: 10.1111/rssb.12017. URL https://doi.org/10.1111/rssb.12017.

Patrik O. Hoyer, Dominik Janzing, Joris Mooij, Jonas Peters, and Bernhard Schölkopf. Nonlinear causal discovery with additive noise models. In *Advances in Neural Information Processing Systems 21*, pp. 689–696. Curran Associates Inc., 2008.

Alexander Immer, Christoph Schultheiss, Julia E Vogt, Bernhard Schölkopf, Peter Bühlmann, and Alexander Marx. On the identifiability and estimation of causal location-scale noise models. In *International Conference on Machine Learning*, pp. 14316–14332, 2023.

Ryo Kamimura and Thong Pham. Causal discovery via transformed low-rank quantile surfaces. *arXiv preprint*, September 2026. doi: 10.48550/arXiv.2609.16931.

Daniel Klippert and Alexander Marx. Skewness-robust causal discovery in location-scale noise models. In *Forty-third International Conference on Machine Learning*, 2026. URL https://openreview.net/forum?id=mXn4ueFKnQ.

Daniel Marbach, Thomas Schaffter, Claudio Mattiussi, and Dario Floreano. Generating realistic in silico gene networks for performance assessment of reverse engineering methods. *Journal of computational biology*, 16(2):229–239, 2009.

Daniel A. Messenger and David M. Bortz. Weak SINDy: Galerkin-based data-driven model selection. *Multiscale Modeling & Simulation*, 19(3):1474–1497, 2021.

Joris M Mooij, Jonas Peters, Dominik Janzing, Jakob Zscheischler, and Bernhard Schölkopf. Distinguishing cause from effect using observational data: methods and benchmarks. *Journal of Machine Learning Research*, 17(32):1–102, 2016.

Jonas Peters, Joris M Mooij, Dominik Janzing, and Bernhard Schölkopf. Causal discovery with continuous additive noise models. *Journal of Machine Learning Research*, 15:2009–2053, 2014.

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(Oct):2003–2030, 2006.

Shohei Shimizu, Takanori Inazumi, Yasuhiro Sogawa, Aapo Hyvärinen, Yoshinobu Kawahara, Takashi Washio, Patrik O Hoyer, and Kenneth Bollen. DirectLiNGAM: A direct method for learning a linear non-Gaussian structural equation model. *Journal of Machine Learning Research*, 12:1225–1248, 2011.

P. Spirtes, C. Glymour, and R. Scheines. *Causation, Prediction, and Search*. MIT press, 2nd edition, 2001.

Natasa Tagasovska, Valérie Chavez-Demoulin, and Thibault Vatter. Distinguishing cause from effect using quantiles: bivariate quantile causal discovery. In *Proceedings of the 37th International Conference on Machine Learning*, ICML'20. JMLR.org, 2020.

Ruibo Tu, Kun Zhang, Hedvig Kjellstrom, and Cheng Zhang. Optimal transport for causal discovery. In *International Conference on Learning Representations*, 2022.

Johnny Xi, Hugh Dance, Peter Orbanz, and Benjamin Bloem-Reddy. Distinguishing cause from effect with causal velocity models. In *Proceedings of the 42nd International Conference on Machine Learning*, ICML'25. JMLR.org, 2025.

K. Zhang and A. Hyvärinen. On the identifiability of the post-nonlinear causal model. In *Proc. 25th Conference on Uncertainty in Artificial Intelligence (UAI2009)*, pp. 647–655, 2009.

A Detailed related work

LiNGAM identifies a linear causal direction through non-Gaussianity and is commonly estimated by independent-component or regression-based procedures (Shimizu et al., 2006; 2011). ANM and PNL identifiability results compare the existence of independent-noise representations in the two directions (Hoyer et al., 2008; Zhang & Hyvärinen, 2009). Their reverse-direction analyses use Bayes' rule to convert a reverse structural constraint into differential restrictions on the cause density. LSNM extends the additive model by allowing the conditional scale to depend on the cause (Immer et al., 2023). Regression followed by an independence test, including RESIT, gives a standard estimation strategy for ANM (Peters et al., 2014). Our scores differ at the estimation stage: they are computed from conditional quantile geometry and do not require a fitted residual followed by a separate independence statistic.

Several causal direction methods use conditional quantiles or other distributional summaries. The bQCD criterion compares predictive behavior across conditional quantiles (Tagasovska et al., 2020). The optimal-transport formulation of Tu et al. (2022) yields, in the ANM case, the same conditional-quantile translation relation $Q(x,u) - Q(x,v) = q(u) - q(v)$. QPE constrains a conditional-quantile velocity field in observed (x, y) coordinates; its finite-span identifiability result assumes a prescribed basis, while its flexible implementation learns the velocity and basis functions (Chen et al., 2026). In contrast, the rank-K FRDR class $\mathcal{M}_K^R$ constrains the output-invariant field R in (x,u) coordinates and uses a direct spectral projection when the basis is unspecified. If ψ denotes the QPE velocity, the two fields satisfy

$$R(x,u) = f\{Q(x,u) \mid x\}\, \psi\{x, Q(x,u)\}.$$

The density factor and the change to quantile coordinates can alter functional rank.

The LRQS model of Kamimura & Pham (2026) assumes an increasing function h such that $h\{Q(x,u)\} = a(x) + \sum_{k=1}^{K} b_k(x) q_k(u)$. Its drift ratio is

$$R(x,u) = \frac{a'(x) + \sum_k b_k'(x) q_k(u)}{\sum_k b_k(x) q_k'(u)}.$$

This ratio need not have finite rank. For example, $Q(x,u) = u + xu^2$ on $x, u \in (0,1)$ is an LRQS surface with

$$R(x,u) = \frac{u^2}{1+2xu}.$$

Rows at distinct positive values of x are linearly independent: after clearing denominators, evaluation at their distinct poles forces every coefficient in a finite linear relation to vanish. Thus LRQS does not generally imply a finite-rank drift ratio. LRQS est-g jointly estimates an output transformation and a low-rank quantile surface, whereas QG-FRDR applies a weighted singular-value decomposition directly to R.

B Visualization of quantile geometry for specific models

We visualize the forward and reverse scores for the LSNM and PNL models in Figure 2.

C Proof of the basis-free FRDR theorem

C.1 Additional regularity

For a reverse rank bound L, the proof uses derivatives of the reverse conditional distribution through order $L+1$ in z, the function $w = 1/p_{X|Y}$ and, when $L \geq 3$, its x-derivatives through order $L-2$, together with the corresponding mixed derivatives of the log-densities. We assume that these derivatives exist and are continuous and that differentiation under the integral defining the reverse conditional distribution is valid. We also assume that $z \mapsto R^{\leftarrow}(z,\cdot)$ is a C^L map into $L^2(J_0)$.

C.2 Proof of Theorem 1

Proof. The forward rank bound follows from $Q \in \mathcal{M}_K^R$. We show that reverse rank at most L restricts the cause density to the stated local families.

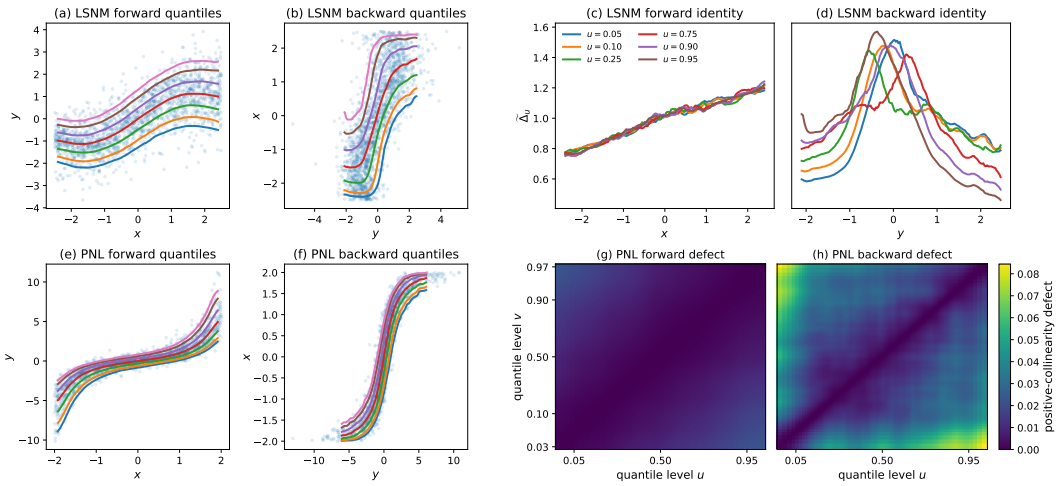

Figure 2: Forward and reverse quantile geometry for LSNM (top) and PNL (bottom). Columns 1 and 2 show estimated conditional quantile curves for context. For LSNM, columns 3 and 4 show normalized interquantile-spacing functions, which collapse in the causal direction but not in reverse. For PNL, they show the positive-collinearity defect between drift-ratio columns $R(\cdot, u)$ and $R(\cdot, v)$; the forward heatmap is nearly zero, whereas the reverse heatmap exhibits more violations.

Express the reverse distribution using Bayes' rule. Write

$$\eta(x) = \log p_X(x), \qquad \nu(z) = \log p_Y(z), \qquad \kappa(x,z) = \log p_{Y|X=x}(z).$$

Bayes' rule gives

$$h(x,z) := p_{X|Y=z}(x) = \exp\{\eta(x) + \kappa(x,z) - \nu(z)\}. \tag{11}$$

Let $H(x,z) = P(X \le x \mid Y = z)$ and define

$$S = -H_z, \qquad w = \frac{1}{h}, \qquad c = \nu' - \kappa_z.$$

Differentiating (11) gives $h_z = -ch$ and $w_z = cw$. Since $H_x = h$, we also obtain

$$S_x = -H_{zx} = -h_z = ch = \frac{c}{w}. \tag{12}$$

We will need the same equation for successive z-derivatives of S. Set $S_j = \partial_z^j S$, so $S_0 = S$, and define

$$P_0 = c, \qquad P_{j+1} = \partial_z P_j - cP_j.$$

Indeed, if $(S_j)_x = P_j/w$, then

$$(S_{j+1})_x = \partial_z\left(\frac{P_j}{w}\right) = \frac{\partial_z P_j}{w} - \frac{P_j w_z}{w^2} = \frac{\partial_z P_j - cP_j}{w}.$$

Consequently,

$$(S_j)_x = \frac{P_j}{w}, \qquad L \ge j \ge 0. \tag{13}$$

At $z = z_0$, the functions $P_0, \dots, P_j$ are determined by the fixed function κ and the finitely many constants $\nu'(z_0), \dots, \nu^{(j+1)}(z_0)$.

Differentiate the reverse drift ratio. The identity $H\{Q^{\leftarrow}(z,u), z\} = u$ gives

$$h\{Q^{\leftarrow}(z,u), z\}\, Q^{\leftarrow}_z + H_z\{Q^{\leftarrow}(z,u), z\} = 0, \qquad h\{Q^{\leftarrow}(z,u), z\}\, Q^{\leftarrow}_u = 1.$$

Thus $Q^{\leftarrow}_z = Sw$ and $Q^{\leftarrow}_u = w$, with the right-hand sides evaluated at $x = Q^{\leftarrow}(z,u)$. It follows that

$$R^{\leftarrow}(z,u) = S\{Q^{\leftarrow}(z,u), z\}. \tag{14}$$

For any smooth function $\Psi(x, z)$, the chain rule gives

$$\begin{aligned}\partial_z \Psi\{Q^{\leftarrow}(z,u), z\} &= \Psi_z\{Q^{\leftarrow}(z,u), z\} + \Psi_x\{Q^{\leftarrow}(z,u), z\} Q_z^{\leftarrow}(z,u) \\ &= \big(\Psi_z + Sw\Psi_x\big)\{Q^{\leftarrow}(z,u), z\}.\end{aligned}$$

To shorten the formulas, define

$$\mathcal{D}\Psi := \Psi_z + Sw\Psi_x. \tag{15}$$

Set $\mathcal{J}_0 = S$ and $\mathcal{J}_{j+1} = \mathcal{D}\mathcal{J}_j$. Repeated application of the displayed chain rule gives, with $\rho(z) = R^{\leftarrow}(z, \cdot)$,

$$\rho^{(j)}(z,u) = \mathcal{J}_j\{Q^{\leftarrow}(z,u), z\}. \tag{16}$$

Use the reverse rank bound. Let $\mathcal{W}$ be the span of the functions $\rho(z)$ over the reverse conditioning interval. The assumption rank $R^{\leftarrow} \le L$ means that $\dim \mathcal{W} \le L$.

Derivatives of ρ also belong to $\mathcal{W}$. For example,

$$\rho'(z_0) = \lim_{t \to 0} \frac{\rho(z_0 + t) - \rho(z_0)}{t}.$$

Every difference quotient belongs to $\mathcal{W}$, and its limit remains there because $\mathcal{W}$ is finite-dimensional. Repeating the argument gives $\rho^{(j)}(z_0) \in \mathcal{W}$ for $0 \le j \le L$.

Moreover, $\rho(z_0) \neq 0$. Otherwise, (14) would give $S(x, z_0) = 0$ throughout $\{Q^{\leftarrow}(z_0, u) : u \in J_0\}$. Equation (12) would then imply $c = 0$ there, hence $c_x = -\kappa_{xz} = 0$, contradicting nondegeneracy.

There is therefore a first $m \in \{1, \dots, L\}$ for which $\rho(z_0), \dots, \rho^{(m)}(z_0)$ are linearly dependent. The preceding functions are linearly independent, so

$$\rho^{(m)}(z_0) = \sum_{j=0}^{m-1} A_j \rho^{(j)}(z_0) \tag{17}$$

for constants $A_0, \dots, A_{m-1}$.

We now write this equality in terms of x. For $x \in \{Q^{\leftarrow}(z_0, u) : u \in J_0\}$, set $u = H(x, z_0)$. Since

$$Q^{\leftarrow}\{z_0, H(x, z_0)\} = x,$$

equation (16) gives

$$\rho^{(j)}\{z_0, H(x, z_0)\} = \mathcal{J}_j(x, z_0).$$

Substituting these expressions into (17) yields

$$\mathcal{J}_m(x, z_0) = \sum_{j=0}^{m-1} A_j \mathcal{J}_j(x, z_0). \tag{18}$$

The L^2 equality holds pointwise because the functions involved are continuous.

Identify the highest derivative of w. We need to determine which derivative of w can be solved for in (18).

First,

$$\mathcal{J}_1 = \mathcal{D}S = S_1 + SwS_x = S_1 + cS.$$

More generally, (13) gives

$$\mathcal{D}S_j = S_{j+1} + Sw(S_j)_x = S_{j+1} + SP_j.$$

Using $P_1 = c_z - c^2$ and $\mathcal{D}c = c_z + Swc_x$, we obtain

$$\begin{aligned}\mathcal{J}_2 &= \mathcal{D}(S_1 + cS) \\ &= S_2 + S(c_z - c^2) + (c_z + Swc_x)S + c(S_1 + cS) \\ &= S_2 + cS_1 + 2c_z S + c_x S^2 w.\end{aligned}$$

Thus

$$\mathcal{J}_2 = S_2 + cS_1 + 2c_z S + c_x S^2 w. \tag{19}$$

For every $j \geq 2$, we claim that

$$\mathcal{J}_j = c_x S^j w^{j-2} \partial_x^{j-2} w + G_j, \tag{20}$$

where G_j involves $S_0, \ldots, S_j$, derivatives of c, and derivatives of w of order at most $j-3$. For $j = 2$, $G_2 = S_2 + cS_1 + 2c_z S$ does not involve w.

To verify the claim, note that $w_z = cw$ implies

$$\mathcal{D}(\partial_x^r w) = \partial_x^r (cw) + Sw\, \partial_x^{r+1} w.$$

The first term contains derivatives of w of order at most r; the second contains the derivative of order $r+1$.

Suppose (20) holds. When $\mathcal{D}$ is applied to its displayed term, differentiating $\partial_x^{j-2} w$ contributes

$$c_x S^j w^{j-2} \cdot Sw\, \partial_x^{j-1} w = c_x S^{j+1} w^{j-1} \partial_x^{j-1} w.$$

Differentiating its coefficient introduces only lower derivatives of w. Applying $\mathcal{D}$ to G_j introduces derivatives of w of order at most $j-2$; derivatives of the S_i are replaced using $\mathcal{D}S_i = S_{i+1} + SP_i$. This proves (20) at order $j+1$.

Obtain a finite system of differential equations. From now on, evaluate every function at $z = z_0$.

We first select an interval on which division by the relevant coefficients is valid. Since $c_x = -\kappa_{xz}$, there is an interval on which $c_x \neq 0$. The function S cannot vanish throughout that interval: otherwise $S_x = c/w$ would imply $c = 0$ there, hence $c_x = 0$. By continuity, we can therefore choose a smaller interval on which $c_x S \neq 0$.

If $m \geq 3$, substitute (20) into (18) and solve for $\partial_x^{m-2} w$:

$$\partial_x^{m-2} w = \frac{\sum_{j=0}^{m-1} A_j \mathcal{J}_j - G_m}{c_x S^m w^{m-2}}. \tag{21}$$

The right-hand side uses only $S_0, \ldots, S_m$ and derivatives of w of order at most $m-3$. Together with

$$(S_j)_x = \frac{P_j}{w}, \qquad 0 \leq j \leq m,$$

and the equations relating successive derivatives of w, this gives a first-order system for

$$S_0, \ldots, S_m, \quad w, \partial_x w, \ldots, \partial_x^{m-3} w.$$

Thus no additional unknown functions are needed.

If $m = 2$, the relation is $\mathcal{J}_2 = A_1 \mathcal{J}_1 + A_0 S$. Substituting the formulas for $\mathcal{J}_1$ and $\mathcal{J}_2$ gives

$$S_2 + cS_1 + 2c_z S + c_x S^2 w = A_1(S_1 + cS) + A_0 S.$$

Hence

$$w = \frac{A_1(S_1 + cS) + A_0 S - S_2 - cS_1 - 2c_z S}{c_x S^2}. \tag{22}$$

Substituting this expression for w into (13) for $j = 0, 1, 2$ gives a first-order system for S_0, S_1, S_2.

If $m = 1$, the relation is $S_1 + cS = A_0 S$. Differentiate with respect to x:

$$(S_1)_x + c_x S + cS_x = A_0 S_x.$$

Using $(S_1)_x = (c_z - c^2)/w$ and $S_x = c/w$ gives

$$\frac{c_z - c^2}{w} + c_x S + \frac{c^2}{w} = \frac{A_0 c}{w}.$$

Therefore

$$w = \frac{A_0 c - c_z}{c_x S}, \qquad S_x = \frac{c}{w}. \tag{23}$$

Substituting the expression for w into the second equation gives a first-order equation for S alone.

In each case, the denominators are nonzero along the solution on the selected interval, and the right-hand sides are locally Lipschitz in the unknown quantities. The solution is therefore determined by the coefficients and initial values.

Recover the cause density. Bayes' rule gives

$$\eta(x) = -\log w(x, z_0) - \kappa(x, z_0) + \nu(z_0). \tag{24}$$

The function κ is fixed. Although ν depends on the cause density, the preceding system uses only the finitely many values $\nu(z_0), \dots, \nu^{(m+1)}(z_0)$.

At fixed initial points, the parameters consist of the m coefficients A_j, these $m + 2$ values of ν and its derivatives, and at most $2m - 1$ initial values. Their total is at most $4m + 1$. Allowing the initial points in x and z to vary adds at most two parameters. Normalization and compatibility with the full joint distribution can only restrict these families further.

Thus, for each m, the restriction of $\log p_X$ to the selected interval belongs to a finite-parameter family. Since $m \in \{1, \dots, L\}$, reverse rank at most L is contained locally in a finite union of such families. Outside these families, rank $R^{\leftarrow} \geq L + 1$. □

C.3 Proof of Proposition 1

Proof. If $s_R = 0$, the forward score is zero by definition. Otherwise, rank $R \leq K$ implies $\sigma_\ell = 0$ for $\ell > K$, so (5) gives $S_K^{\mathrm{FRDR}}(R) = 0$. By Theorem 1 with $L = K$, the reverse field generically has rank at least $K + 1$. Its $(K + 1)$st singular value is therefore positive, and $S_K^{\mathrm{FRDR}}(R^{\leftarrow}) > 0$. □

D Details for the class-specific scores

D.1 Proof of Proposition 2

Proof. For ANM, LSNM, and PNL-HNM, the score integrands are nonnegative and continuous. A zero integral therefore makes the corresponding defect vanish pointwise. For PNL, the analogous almost-everywhere statement is sufficient.

ANM. If $Q(x, u) = a(x) + q(u)$, then $Q(x, u) - Q(x, v) = q(u) - q(v)$, so $S_{X\to Y}^{\mathrm{ANM}} = 0$.

Conversely, suppose $S_{Y\to X}^{\mathrm{ANM}} = 0$. Fix $v_0 \in J_0$ and a reverse conditioning value y_0. The zero-variance identity gives

$$Q^{\leftarrow}(y, u) - Q^{\leftarrow}(y, v_0) = Q^{\leftarrow}(y_0, u) - Q^{\leftarrow}(y_0, v_0).$$

Thus, with

$$\widetilde{a}(y) := Q^{\leftarrow}(y, v_0), \qquad \widetilde{q}(u) := Q^{\leftarrow}(y_0, u) - Q^{\leftarrow}(y_0, v_0),$$

we have $Q^{\leftarrow}(y, u) = \widetilde{a}(y) + \widetilde{q}(u)$. Consequently,

$$R^{\leftarrow}(y, u) = \frac{\widetilde{a}'(y)}{\widetilde{q}'(u)}$$

has rank at most one.

LSNM. For any $f, g \in L^2(I)$,

$$\int_{I^2} \{f(x_1)g(x_2) - g(x_1)f(x_2)\}^2 \, \mathrm{d}\mu_I^{\otimes 2}(x_1, x_2) = 2\mathcal{A}(f, g). \tag{25}$$

Hence (8) is exactly the integrated squared spacing-determinant score.

Under $Q(x, u) = a(x) + b(x)q(u)$, $\Delta_{uv} = \{q(v) - q(u)\}b$, so every defect in (8) is zero.

Conversely, suppose $S_{Y\to X}^{\mathrm{LSNM}} = 0$. Fix $u_0 < u_1$ and a reverse conditioning value y_0. Define

$$\widetilde{b}(y) := \frac{Q^{\leftarrow}(y, u_1) - Q^{\leftarrow}(y, u_0)}{Q^{\leftarrow}(y_0, u_1) - Q^{\leftarrow}(y_0, u_0)}, \qquad \widetilde{a}(y) := Q^{\leftarrow}(y, u_0),$$

and

$$\widetilde{q}(u) := Q^{\leftarrow}(y_0, u) - Q^{\leftarrow}(y_0, u_0).$$

Strict monotonicity in u gives $\widetilde{b}(y) > 0$. Since $\mathcal{A}(\Delta_{u_0u_1}^{\leftarrow}, \Delta_{u_0u}^{\leftarrow}) = 0$, evaluation at y_0 yields

$$Q^{\leftarrow}(y, u) - Q^{\leftarrow}(y, u_0) = \widetilde{b}(y)\{Q^{\leftarrow}(y_0, u) - Q^{\leftarrow}(y_0, u_0)\}.$$

Therefore

$$Q^{\leftarrow}(y,u) = \widetilde{a}(y) + \widetilde{b}(y)\widetilde{q}(u),$$

and

$$R^{\leftarrow}(y,u) = \frac{\widetilde{a}'(y)}{\widetilde{b}(y)}\frac{1}{\widetilde{q}'(u)} + \frac{\widetilde{b}'(y)}{\widetilde{b}(y)}\frac{\widetilde{q}(u)}{\widetilde{q}'(u)}$$

has rank at most two.

PNL. Under PNL, $R(x,u) = a'(x)/q'(u)$, with $q'(u) > 0$. Thus all nonzero columns are positively proportional and $S^{\mathrm{PNL}}_{X\to Y} = 0$.

Suppose $S^{\mathrm{PNL}}_{Y\to X} = 0$, and write $C_u := R^{\leftarrow}(\cdot,u) \in L^2$. If every C_u is zero, the reverse field has rank zero. Otherwise, by Fubini's theorem and continuity, one may choose a nonzero reference column C_{u_0} such that $\mathrm{pcol}(C_u, C_{u_0}) = 0$ for almost every u. Equality in Cauchy–Schwarz then gives $C_u = c(u)C_{u_0}$ with $c(u) \geq 0$. Hence the reverse field has rank at most one.

PNL-HNM. Write

$$R(x,u) = \alpha(x)\phi_1(u) + \beta(x)\phi_2(u), \qquad \phi_1 = \frac{1}{q'}, \quad \phi_2 = \frac{q}{q'}.$$

A direct calculation gives

$$\phi_1\phi_2' - \phi_1'\phi_2 = \frac{1}{q'}.$$

Consequently,

$$W_{12}(u) = \frac{\beta(x_1)\alpha(x_2) - \alpha(x_1)\beta(x_2)}{q'(u)}. \tag{26}$$

Writing $\tau = q''/q'$, we also have

$$R_i' + \tau R_i = \beta(x_i), \qquad W_{12}' = -\tau W_{12}.$$

It follows that

$$G_{123} = W_{12}R_3' - W_{12}'R_3 = \beta(x_3)W_{12},$$

which proves $S^{\text{PNL-HNM}}_{X\to Y} = 0$. Moreover, under $R \mapsto cR$, $W_{12} \mapsto c^2W_{12}$ and $G_{123} \mapsto c^3G_{123}$; hence the numerator of (10) is homogeneous of degree 10, justifying the normalization by s_R^{10}.

It remains to consider the reverse implication for PNL-HNM. Write $r_y = R^{\leftarrow}(y,\cdot)$ and $W_{y_1y_2} = r_{y_1}'r_{y_2} - r_{y_1}r_{y_2}'$. Suppose the reverse score is zero.

We first retain nondegeneracy when restricting the quantile interval. From (14) and (12),

$$r_{z_0}''(u) = -\kappa_{xz}\{Q^{\leftarrow}(z_0,u), z_0\}Q_u^{\leftarrow}(z_0,u).$$

Nondegeneracy therefore gives an interval on which $r_{z_0}'' \neq 0$. Since r_{z_0} cannot vanish throughout that interval, choose a compact subinterval $J_1 \subset \mathrm{int}(J_0)$ with nonempty interior on which $r_{z_0}r_{z_0}'' \neq 0$.

Suppose some $W_{y_1y_2}$ is nonzero at a point of J_1. Choose a compact subinterval $J_0' \subset J_1$ with nonempty interior on which $W := W_{y_1y_2}$ is nowhere zero. The zero score and equality in Cauchy–Schwarz give

$$Wr_y' - W'r_y = \lambda(y)W$$

for every y, where $\lambda(y)$ is independent of u. Thus, for any $u_* \in J_0'$,

$$r_y(u) = c(y)W(u) + \lambda(y)W(u)\int_{u_*}^{u}\frac{\mathrm{d}t}{W(t)}. \tag{27}$$

The restricted rows therefore lie in a common space of dimension at most two.

Otherwise, every pairwise Wronskian vanishes on J_1. Since r_{z_0} is nowhere zero there,

$$\left(\frac{r_y}{r_{z_0}}\right)' = \frac{W_{yz_0}}{r_{z_0}^2} = 0,$$

so the restricted rows have rank at most one. In both cases, the reverse rank is at most two on an interval where nondegeneracy still holds at z_0.

Combining the four cases, Theorem 1 applies with $L = 1$ for ANM and PNL and $L = 2$ for LSNM and PNL-HNM, proving the generic reverse statements. □

D.2 Why the PNL-HNM identity is stronger than rank two

Consider

$$R(x,u) = \gamma(x)\cos u + \delta(x)\sin u.$$

This field has rank at most two. For any x_1, x_2,

$$W_{12} = \delta(x_1)\gamma(x_2) - \gamma(x_1)\delta(x_2),$$

which is constant in u. If this constant is nonzero, then $W'_{12} = 0$, and the identity $G_{123} = \beta(x_3)W_{12}$ underlying the score (10) would require R'_3 to be constant in u for every x_3. This fails for generic γ and δ. Thus the PNL-HNM identity restricts the relation between the two basis functions.

D.3 Local converses

The following elementary observation converts drift-ratio identities into local structural representations.

Lemma 2 (Output equivalence of the drift ratio)**.** *Let $Q_1, Q_2 \in C^1$ have positive u-derivatives and the same drift-ratio field on an open domain. Near every interior point, there is a C^1 function h with $h' > 0$ such that $Q_2 = h \circ Q_1$.*

Proof. Use the local coordinates $(x, s) = (x, Q_1(x,u))$, which are valid because $Q_{1,u} > 0$, and write the inverse as $u = v(x,s)$. After restricting to a rectangle in these coordinates, define $G(x,s) = Q_2(x, v(x,s))$. Differentiating $Q_1(x, v(x,s)) = s$ gives

$$G_x = \left(Q_{2,x} - Q_{2,u}\frac{Q_{1,x}}{Q_{1,u}}\right)(x, v(x,s)) = 0.$$

Thus $G(x,s) = h(s)$, and

$$h'(s) = \frac{Q_{2,u}(x, v(x,s))}{Q_{1,u}(x, v(x,s))} > 0.$$

Hence $Q_2 = h \circ Q_1$ locally. □

Proposition 3 (Local PNL converse)**.** *Suppose, on a rectangle,*

$$R(x,u) = A(x)\phi(u), \qquad \phi(u) > 0.$$

Then

$$Q(x,u) = g\{a(x) + q(u)\}$$

locally for some $g \in C^1$ with $g' > 0$.

Proof. Choose a and q such that $a' = A$ and $q' = 1/\phi$. Then $q' > 0$, and $Q_0(x,u) := a(x) + q(u)$ satisfies

$$\frac{(Q_0)_x}{(Q_0)_u} = A(x)\phi(u) = R(x,u).$$

Lemma 2 gives $Q = g \circ Q_0$. □

Proposition 4 (Local PNL-HNM converse)**.** *Suppose that W_{12} is nowhere zero for some x_1, x_2, and that for every x,*

$$W_{12}R_u(x,\cdot) - W'_{12}R(x,\cdot) = \lambda(x)W_{12}$$

for some scalar $\lambda(x)$. Then

$$Q(x,u) = g\{a(x) + b(x)q(u)\}$$

locally, where $b > 0$ and $g \in C^1$ satisfies $g' > 0$.

Proof. Let $W = W_{12}$ and define

$$T(u) := -\frac{W'(u)}{W(u)}, \qquad q'(u) := \exp\left\{\int_{u_*}^{u} T(v)\,\mathrm{d}v\right\} > 0.$$

The assumed identity gives $R_u + TR = \lambda$. Since $q''/q' = T$,

$$(q'R)_u = \lambda(x)q'.$$

Table 5: Composition of the synthetic evaluation. Each setting contains 100 independently generated pairs, and each pair contains $n = 1000$ observations.

Family	Construction	Settings	Pairs
FRDR	finite-rank drift ratio, GP coefficients	2	200
PNL-HNM	12 GP-based + 8 analytic mechanisms	20	2,000
LSNM	3 GP-based + 8 analytic mechanisms	11	1,100
PNL	3 noises × 4 transformations	12	1,200
ANM	3 noise distributions	3	300
Total		48	4,800

After choosing an antiderivative q, integration in u yields

$$R(x,u) = \frac{\alpha(x)}{q'(u)} + \lambda(x)\frac{q(u)}{q'(u)}$$

for some function α. Choose local solutions $b > 0$ and a of

$$\frac{b'}{b} = \lambda, \qquad a' = \alpha b.$$

Then $Q_0(x,u) := a(x) + b(x)q(u)$ has drift ratio R. Lemma 2 therefore gives $Q = g \circ Q_0$. □

A zero QG-PNL score gives the representation in Proposition 3 on subrectangles where the proportionality factors are strictly positive. For QG-PNL-HNM, a zero score and a nowhere-vanishing Wronskian give the representation in Proposition 4.

E Proof of Lemma 1

Proof. Increasing maps commute with conditional quantiles: $Q_{\widetilde{B}|\widetilde{A}=t}(u) = h\{Q_{B|A=a}(u)\}$. Differentiating in t and u cancels the factor h' and contributes $\mathrm{d}a/\mathrm{d}t = 1/m'(a)$, giving the stated formula. The multiplier is positive and independent of u, so it preserves the row span, positive proportionality of columns, and the PNL-HNM identity. □

F Additional experimental details

F.1 Synthetic datasets and protocol

Table 5 summarizes the 48 synthetic settings. Each setting contains 100 independently generated pairs, and each pair contains $n = 1000$ observations.

GP-based ANM, PNL, LSNM, and PNL-HNM settings. The base construction is

$$X \sim \mathcal{N}(0,2), \qquad Z = f(X) + s(X)\varepsilon, \qquad Y = h(Z).$$

The function f and, for the location–scale settings, an independent function g are sampled from a zero-mean Gaussian process with RBF covariance $k(x,x') = \exp\{-(x-x')^2/2\}$. We set $s(x) = 1$ for the additive-noise settings and $s(x) = |g(x)|$ for the location–scale settings. The three zero-mean, unit-variance noise distributions are standard Gaussian, $\mathrm{Unif}(-\sqrt{3},\sqrt{3})$, and $\sqrt{20}\{B - 1/2\}$ with $B \sim \mathrm{Beta}(2,2)$. The output transformation h is the identity, cube, logistic sigmoid, exponential, or hyperbolic tangent. The identity settings form the three ANM and three GP-based LSNM settings; the four nonidentity transformations form the 12 PNL and 12 GP-based PNL-HNM settings.

Analytic LSNM and PNL-HNM settings. To make the location and scale functions explicit, we additionally generate

$$X \sim \mathcal{N}(0,1), \qquad Z = a(X) + b(X)\varepsilon.$$

The LSNM response is $Y = Z$. Its paired PNL-HNM response is $Y = h(Z)$, where h is balanced across the cube, logistic sigmoid, exponential, and hyperbolic-tangent transformations over the 100

Table 6: Analytic location–scale mechanisms. The rank column gives the rank of the resulting drift-ratio representation.

ID	$a(x)$	$b(x)$	Rank
LQ	$0.4x$	$0.4 + 0.2x^2$	2
L4	$0.4x$	$0.4 + x^4/15$	2
LB	$0.4x$	$0.4 + 0.6x^2/(1 + x^2)$	2
LH	$0.4x$	$0.4 + 0.3\{\cosh(x) - 1\}$	2
LP	$0.4x$	$0.4 + 0.5\sin^2(x)$	2
QQ	$0.4x^2$	$0.4 + 0.2x^2$	1
Q4	$0.4x^2$	$0.4 + x^4/15$	2
PB	$1.2x^2/(1 + x^2)$	$0.4 + 0.6x^2/(1 + x^2)$	1

pairs. The three noise distributions above are also balanced over pairs. Table 6 defines all eight mechanisms. For each mechanism and pair, the LSNM and PNL-HNM versions share X, ε, Z, and the observed column order.

FRDR settings. The two FRDR settings directly specify

$$R(x, u) = \frac{Q_x(x, u)}{Q_u(x, u)} = \sum_{k=1}^{K_0} \alpha_k(x)\phi_k(u),$$

where

$$\phi_1(u) = u(1 - u), \quad \phi_2(u) = u(1 - u)(2u - 1), \quad \phi_3(u) = u(1 - u)(6u^2 - 6u + 1).$$

The rank-two setting uses the first two terms and the rank-three setting uses all three. Independently for each pair, the coefficient functions are drawn from an RBF Gaussian process with unit length scale on a 257-point grid over $[-5, 5]$ and are scaled by 1.6, 0.4, and 0.4, respectively. The conditional quantile surface is recovered from $Q_x = RQ_u$ by a 160-step RK4 characteristic solver, using the standardized logistic quantile at $x = 0$. Finally, $X \sim \mathcal{N}(0, 1)$, $U \sim \mathrm{Unif}(10^{-5}, 1 - 10^{-5})$, and $Y = Q(X, U)$. All rank-sensitive methods are fitted with $K = 2$ on both FRDR settings, so `frdr-k3-gp` evaluates rank misspecification.

The QG-FRDR and QG-FRDR-W implementations construct copula coordinates from the unstandardized observations. LRQS, CVEL, QPE-k, and QPE-f lowrank receive the affine-standardized inputs required by their implementations. Pair orientations are randomized, and accuracy is the percentage for which the known generating direction is selected.

F.2 Setting-level synthetic results

Tables 7 and 8 report the setting-level results for the seven general-purpose methods.

Tables 9 and 10 give the corresponding setting-level results for the model-specific scores in Panel B of Table 3.

F.3 Dataset-level benchmark results

Table 11 gives the dataset-level accuracies underlying Table 4.

F.4 Sensitivity details

The sensitivity samples are generated at $n = 5000$, with the first 500 and 1000 observations used for the smaller-sample conditions. This prefix design isolates sample-size effects within each pair. For the 30 GP-based ANM/PNL/LSNM/PNL-HNM settings, the same RBF Gaussian processes are generated on a fixed grid and linearly interpolated, avoiding repeated factorization of a 5000×5000 covariance matrix. The Table 3 samples are not modified.

All conditions except $K = 1$ fit $K = 2$, including the rank-three FRDR setting; the $K = 1$ condition alone changes the fitted rank.

Table 7: Setting-level causal-direction accuracy (%) for general-purpose methods: FRDR and PNL-HNM settings.

Setting	QG-FRDR	QG-FRDR-W	LRQS est-g	LRQS fix-g	CVEL	QPE-k	QPE-f lowrank
FRDR							
frdr-k2-gp	**82.0**	71.0	58.0	63.0	44.0	56.0	35.0
frdr-k3-gp	**65.0**	49.0	52.0	50.0	54.0	46.0	35.0
pnl-hnm-gaussian-cube	99.0	97.0	**100.0**	97.0	99.0	48.0	73.0
pnl-hnm-gaussian-sigmoid	**100.0**	98.0	99.0	95.0	20.0	79.0	56.0
pnl-hnm-gaussian-exp	98.0	**100.0**	**100.0**	**100.0**	94.0	32.0	85.0
pnl-hnm-gaussian-tanh	**97.0**	**97.0**	96.0	78.0	66.0	45.0	54.0
pnl-hnm-uniform-cube	98.0	97.0	**100.0**	95.0	99.0	76.0	75.0
pnl-hnm-uniform-sigmoid	98.0	99.0	**100.0**	97.0	39.0	94.0	38.0
pnl-hnm-uniform-exp	95.0	94.0	98.0	**100.0**	95.0	65.0	76.0
pnl-hnm-uniform-tanh	97.0	97.0	**99.0**	83.0	74.0	60.0	46.0
pnl-hnm-beta-cube	**99.0**	98.0	98.0	88.0	98.0	63.0	82.0
pnl-hnm-beta-sigmoid	98.0	96.0	**99.0**	97.0	22.0	91.0	37.0
pnl-hnm-beta-exp	96.0	99.0	**100.0**	**100.0**	92.0	39.0	75.0
pnl-hnm-beta-tanh	**97.0**	**97.0**	**97.0**	77.0	76.0	53.0	52.0
pnl-hnm-extended-lq	0.0	83.0	94.0	96.0	96.0	93.0	**99.0**
pnl-hnm-extended-l4	2.0	72.0	98.0	**100.0**	97.0	91.0	85.0
pnl-hnm-extended-lb	4.0	**88.0**	86.0	84.0	87.0	80.0	**88.0**
pnl-hnm-extended-lh	0.0	68.0	95.0	**98.0**	95.0	90.0	93.0
pnl-hnm-extended-lp	2.0	88.0	91.0	90.0	89.0	75.0	**94.0**
pnl-hnm-extended-qq	**100.0**	**100.0**	96.0	92.0	73.0	86.0	32.0
pnl-hnm-extended-q4	**100.0**	**100.0**	99.0	**100.0**	85.0	83.0	40.0
pnl-hnm-extended-pb	**100.0**	**100.0**	98.0	95.0	60.0	51.0	13.0

Table 8: Setting-level causal-direction accuracy (%) for general-purpose methods: LSNM, PNL, and ANM settings.

Setting	QG-FRDR	QG-FRDR-W	LRQS est-g	LRQS fix-g	CVEL	QPE-k	QPE-f lowrank
LSNM							
pnl-hnm-gaussian-identity	98.0	99.0	**100.0**	**100.0**	18.0	94.0	59.0
pnl-hnm-uniform-identity	96.0	97.0	**100.0**	**100.0**	24.0	**100.0**	36.0
pnl-hnm-beta-identity	**100.0**	98.0	99.0	99.0	17.0	98.0	41.0
lsnm-extended-lq	0.0	82.0	**99.0**	96.0	93.0	98.0	96.0
lsnm-extended-l4	2.0	72.0	**100.0**	96.0	**100.0**	98.0	89.0
lsnm-extended-lb	4.0	88.0	95.0	96.0	74.0	**100.0**	99.0
lsnm-extended-lh	0.0	69.0	**99.0**	95.0	93.0	98.0	93.0
lsnm-extended-lp	2.0	88.0	97.0	97.0	67.0	**98.0**	96.0
lsnm-extended-qq	**100.0**	**100.0**	99.0	98.0	**100.0**	**100.0**	19.0
lsnm-extended-q4	**100.0**	**100.0**	**100.0**	**100.0**	**100.0**	99.0	24.0
lsnm-extended-pb	**100.0**	**100.0**	98.0	98.0	24.0	94.0	18.0
PNL							
pnl-an-gaussian-cube	**100.0**	**100.0**	99.0	82.0	**100.0**	34.0	43.0
pnl-an-gaussian-sigmoid	99.0	**100.0**	95.0	96.0	4.0	75.0	48.0
pnl-an-gaussian-exp	98.0	98.0	**100.0**	**100.0**	**100.0**	16.0	61.0
pnl-an-gaussian-tanh	**98.0**	97.0	91.0	59.0	70.0	10.0	31.0
pnl-an-uniform-cube	82.0	**100.0**	95.0	56.0	99.0	74.0	22.0
pnl-an-uniform-sigmoid	82.0	**100.0**	94.0	91.0	87.0	**100.0**	64.0
pnl-an-uniform-exp	86.0	99.0	98.0	97.0	**100.0**	49.0	72.0
pnl-an-uniform-tanh	84.0	**100.0**	93.0	69.0	**100.0**	43.0	51.0
pnl-an-beta-cube	93.0	**100.0**	96.0	58.0	**100.0**	60.0	29.0
pnl-an-beta-sigmoid	92.0	**100.0**	95.0	95.0	52.0	97.0	50.0
pnl-an-beta-exp	92.0	**100.0**	98.0	**100.0**	**100.0**	29.0	69.0
pnl-an-beta-tanh	96.0	**99.0**	89.0	74.0	**99.0**	25.0	52.0
ANM							
pnl-an-gaussian-identity	99.0	**100.0**	90.0	94.0	2.0	90.0	77.0
pnl-an-uniform-identity	86.0	**100.0**	93.0	98.0	76.0	**100.0**	60.0
pnl-an-beta-identity	93.0	**100.0**	96.0	93.0	19.0	**100.0**	51.0

F.5 Score-gap and convergence details

Table 13 first computes the mean normalized gap and near-tie rate within each dataset and then averages those quantities over datasets. The median column is the median of the dataset mean gaps. This gives each benchmark dataset equal weight despite unequal pair counts. The Overall rows combine the 48 synthetic and 19 benchmark datasets with the same equal-dataset weighting used in panel (a) of Figure 4.

Table 9: Setting-level causal-direction accuracy (%) for the model-specific QG-PNL-HNM and QG-LSNM scores.

PNL-HNM setting	QG-PNL-HNM	LSNM setting	QG-LSNM
pnl-hnm-gaussian-cube	95.0	pnl-hnm-gaussian-identity	100.0
pnl-hnm-gaussian-sigmoid	98.0	pnl-hnm-uniform-identity	99.0
pnl-hnm-gaussian-exp	97.0	pnl-hnm-beta-identity	99.0
pnl-hnm-gaussian-tanh	93.0	lsnm-extended-lq	100.0
pnl-hnm-uniform-cube	94.0	lsnm-extended-l4	100.0
pnl-hnm-uniform-sigmoid	95.0	lsnm-extended-lb	97.0
pnl-hnm-uniform-exp	93.0	lsnm-extended-lh	99.0
pnl-hnm-uniform-tanh	97.0	lsnm-extended-lp	100.0
pnl-hnm-beta-cube	95.0	lsnm-extended-qq	100.0
pnl-hnm-beta-sigmoid	97.0	lsnm-extended-q4	100.0
pnl-hnm-beta-exp	93.0	lsnm-extended-pb	100.0
pnl-hnm-beta-tanh	92.0		
pnl-hnm-extended-lq	0.0		
pnl-hnm-extended-l4	0.0		
pnl-hnm-extended-lb	0.0		
pnl-hnm-extended-lh	0.0		
pnl-hnm-extended-lp	0.0		
pnl-hnm-extended-qq	100.0		
pnl-hnm-extended-q4	100.0		
pnl-hnm-extended-pb	100.0		

Table 10: Setting-level causal-direction accuracy (%) for the model-specific PNL and ANM comparisons.

PNL setting	QG-PNL	ANM setting	QG-ANM	ANM	DIVOT
pnl-an-gaussian-cube	100.0	pnl-an-gaussian-identity	**100.0**	**100.0**	96.0
pnl-an-gaussian-sigmoid	100.0	pnl-an-uniform-identity	**100.0**	**100.0**	**100.0**
pnl-an-gaussian-exp	100.0	pnl-an-beta-identity	**100.0**	**100.0**	**100.0**
pnl-an-gaussian-tanh	99.0				
pnl-an-uniform-cube	100.0				
pnl-an-uniform-sigmoid	98.0				
pnl-an-uniform-exp	100.0				
pnl-an-uniform-tanh	98.0				
pnl-an-beta-cube	98.0				
pnl-an-beta-sigmoid	99.0				
pnl-an-beta-exp	100.0				
pnl-an-beta-tanh	99.0				

No exact score ties occur in these evaluations. All QG-FRDR-W fits converge on the 4,800 synthetic pairs. On the benchmarks, one true-direction fit does not report convergence (Multi, pair 97), while its reverse-direction fit converges. Both scores are finite and are retained; removing this pair does not change any reported accuracy to one decimal place.

F.6 Robustness and score separation

F.6.1 Sensitivity analysis

We vary one component of QG-FRDR at a time: sample size, quantile range, grid resolution, or fitted rank. To make the $n = 500$, 1000, and 5000 conditions nested, this analysis uses a separate set of samples generated from the same 48 data-generating settings as Table 3. The reference condition is therefore not a re-evaluation of the exact main-experiment pairs. It uses $n = 1000$, $u \in [0.01, 0.99]$, a 50×50 grid, and $K = 2$.

Figure 3 shows that the two coarser grids and the wider quantile range produce only small changes in the overall averages from the reference condition. Restricting the quantile range to $[0.20, 0.80]$ lowers accuracy from 75.6% to 74.0% and the mean normalized gap from 0.334 to 0.304. Increasing the sample size to 5000 raises these values to 77.8% and 0.426, respectively. The largest deterioration occurs when the fitted rank is reduced to $K = 1$, for which accuracy falls to 69.4% and the mean gap to 0.224. Family-level accuracies are reported in Appendix F.4.

Table 11: Dataset-level causal-direction accuracy (%) on the 19 cause–effect benchmarks.

Dataset	QG-FRDR	QG-FRDR-W	LRQS est-g	LRQS fix-g	CVEL	QPE-k	QPE-f lowrank
SIM	62.0	71.0	61.0	69.0	64.0	**83.0**	74.0
SIM-c	73.0	67.0	52.0	77.0	62.0	**79.0**	66.0
SIM-g	68.0	68.0	71.0	72.0	80.0	**83.0**	55.0
SIM-ln	69.0	**78.0**	56.0	64.0	65.0	68.0	72.0
Cha	66.3	69.3	60.7	**70.3**	67.0	59.7	55.3
Multi	62.7	72.0	84.3	85.7	87.0	**88.3**	69.7
Net	78.0	76.0	70.7	85.3	51.0	**89.3**	63.0
Per	11.0	54.0	64.0	74.0	**100.0**	77.0	83.0
Sig	28.0	37.0	54.0	52.0	69.0	**89.0**	73.0
Vex	8.0	11.0	78.0	70.0	91.0	63.0	**97.0**
Qd-V	41.0	41.0	57.0	73.0	**83.0**	42.0	73.0
Sig-V	62.0	57.0	60.0	79.0	**90.0**	67.0	67.0
RbF-V	45.0	51.0	44.0	65.0	**90.0**	68.0	69.0
NN-V	44.0	43.0	61.0	79.0	**88.0**	53.0	77.0
Tue	63.6	48.5	68.7	**77.8**	34.3	53.5	48.5
D4-s1	33.3	50.0	50.0	41.7	41.7	**58.3**	33.3
D4-s2a	48.9	51.7	53.1	64.2	46.6	**67.0**	65.6
D4-s2b	52.2	49.0	51.4	53.8	41.4	**61.0**	53.0
D4-s2c	49.2	46.6	50.5	53.4	47.2	**64.2**	55.4
Dataset average	50.8	54.8	60.4	68.7	68.3	**69.1**	65.8

Table 12: Sensitivity accuracy (%) by synthetic family. Ref. denotes the $n = 1000$, $u \in [0.01, 0.99]$, 50×50, $K = 2$ reference condition on the sensitivity-specific samples.

Family	Ref.	$n = 500$	$n = 5000$	$u = .20$–$.80$	$u = .005$–$.995$	10^2 grid	20^2 grid	$K = 1$
FRDR	72.5	69.5	78.0	69.0	72.5	60.5	70.0	70.0
PNL-HNM	74.0	73.5	75.9	73.0	73.9	74.6	73.9	60.2
LSNM	54.8	54.9	57.5	53.9	54.6	55.9	54.7	48.8
PNL	93.1	92.0	95.1	90.3	93.1	93.5	93.0	96.4
ANM	94.7	93.7	96.3	92.3	94.3	95.0	94.0	97.3
Overall	75.6	75.0	77.8	74.0	75.5	75.7	75.4	69.4

F.6.2 Normalized score-gap analysis

For each pair, let S_{true} and S_{rev} be the QG score in the generating and reverse directions. We define

$$G = \frac{S_{\text{rev}} - S_{\text{true}}}{S_{\text{rev}} + S_{\text{true}}},$$

so that $G > 0$ supports the generating direction and values near zero indicate weak separation. Figure 4 first averages G within each dataset, so every dataset contributes one point regardless of its number of pairs. Panel (a) provides an overall view of the 48 synthetic and 19 benchmark datasets, and panel (b) displays the same dataset-level values by family.

Across all 67 datasets in panel (a), the median dataset mean gap is 0.386 for QG-FRDR and 0.466 for QG-FRDR-W. Panel (b) shows that this overall distribution combines predominantly positive separation on the synthetic families with much weaker and more heterogeneous separation on the benchmarks. On the 48 synthetic datasets, the dataset-averaged mean gap is 0.332 for QG-FRDR and 0.411 for QG-FRDR-W. On the 19 benchmarks, these values fall to −0.007 and 0.037, respectively, and the dataset-averaged proportions of near ties, defined by $|G| < 0.05$, rise to 17.1% and 18.2%. Thus, the lower benchmark accuracy is accompanied by substantially weaker directional score separation. Detailed overall and family summaries, together with convergence checks, appear in Appendix F.5.

F.7 Population score gaps for the analytic mechanisms

For the eight mechanisms of Table 6 with Gaussian cause and Gaussian noise, the forward drift ratio is available in closed form from (3), and the reverse drift ratio follows from (14) by evaluating $H(x, z) = P(X \le x \mid Y = z)$, $S = -H_z$ and $Q^{\leftarrow}(z, u)$ numerically on a fine grid in x. Both fields are expressed in copula coordinates as in Section 7 and scored on the 50×50 grid with the trapezoidal weights of Appendix G; no sampling is involved, so the entries of Table 14 are population values of the scores as implemented. The table also reports a linear Gaussian calibration, for which both directions

Table 13: Dataset-balanced normalized score-gap summaries. Near ties satisfy $|G| < 0.05$.

Method	Family	Mean gap	Median gap	Near tie (%)
QG-FRDR	FRDR	0.140	0.140	14.0
	PNL-HNM	0.347	0.572	1.6
	LSNM	0.155	0.457	2.2
	PNL	0.466	0.473	3.4
	ANM	0.468	0.449	4.7
	Benchmark	−0.007	0.017	17.1
	Overall (67)	0.236	0.386	6.9
QG-FRDR-W	FRDR	0.061	0.061	22.5
	PNL-HNM	0.400	0.485	5.1
	LSNM	0.327	0.449	8.3
	PNL	0.531	0.546	1.2
	ANM	0.539	0.540	1.0
	Benchmark	0.037	0.010	18.2
	Overall (67)	0.305	0.466	9.0

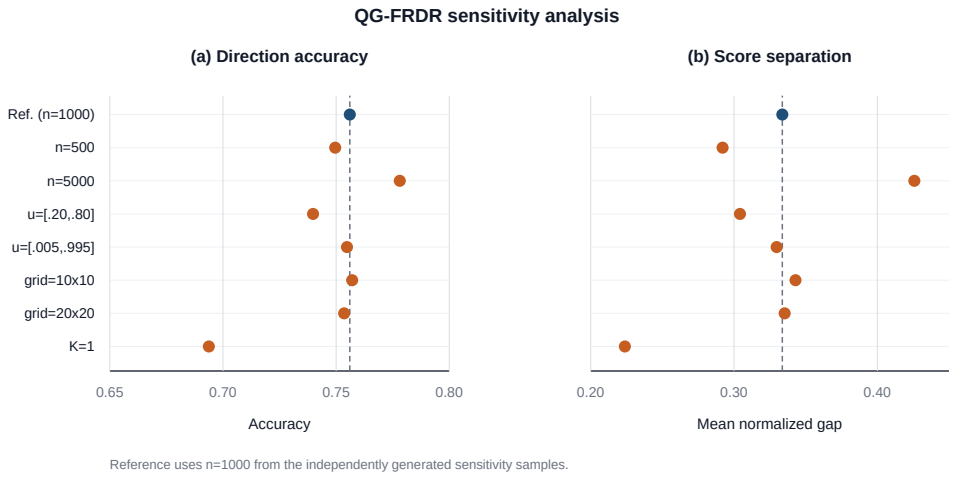

Figure 3: Sensitivity of QG-FRDR across 48 synthetic data-generating settings. Panel (a) reports direction accuracy and panel (b) the mean normalized gap. The blue point denotes the reference condition, while the orange points vary one component at a time. The vertical dashed line marks the reference value in each panel. The reference uses $n = 1000$ observations from independently generated sensitivity samples that follow the settings in Table 3 and permit nested comparisons across $n = 500$, 1000, and 5000.

have rank one. Because QG-FRDR operates on copula coordinates, the LSNM and PNL-HNM versions of each analytic pair have identical inputs and identical scores, which is why their accuracies coincide in Tables 7 and 8.

For the eight nonlinear mechanisms, the forward rank-two score vanishes to numerical precision, whereas the reverse score is positive but small. These numerical results are consistent with Theorem 1, which establishes generic positivity without a quantitative lower bound.

For the five linear-location mechanisms, the rank-one population score is smaller in reverse than in the causal direction. This illustrates how an underspecified rank can favor the reverse direction even in population. With $K = 2$, finite-sample estimation can instead obscure the small population gap. In the linear-Gaussian calibration, the mean estimated scores at $n = 1000$ are 0.029 and 0.028, although both population scores are zero. For the five linear-location mechanisms under the noise-balanced protocol, the mean estimated scores are 0.026–0.033 causally and 0.014–0.018 in reverse. The corresponding ranges for QQ, Q4, and PB are 0.014–0.021 and 0.039–0.073.

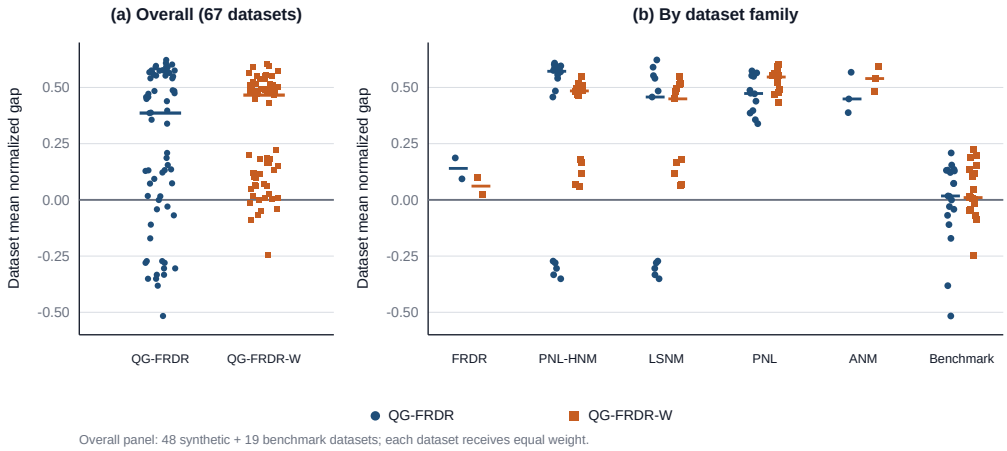

Figure 4: Dataset-level normalized score gaps for QG-FRDR and QG-FRDR-W. Each point is the mean gap within one dataset. Panel (a) pools the 48 synthetic datasets from Table 3 and the 19 benchmarks from Table 4, giving each dataset equal weight. Panel (b) groups the same values by dataset family. The short horizontal segment marks the median for each method within the displayed group. Positive values favor the generating direction; the horizontal zero line indicates no average separation.

Table 14: Population scores for the analytic mechanisms of Table 6, with independent standard Gaussian cause and noise. Both the copula conditioning grid and the conditional quantile grid contain 50 equally spaced points on $[0.01, 0.99]$, with normalized trapezoidal weights. $S_K = S_K^{\mathrm{FRDR}}$, and σ_3/σ_1 is the ratio of the third to the first singular value of the weighted reverse drift-ratio matrix. Zeros denote analytically vanishing quantities. The linear Gaussian calibration uses $a(x) = 0.4x$ and $b(x) = 0.6$. The last column repeats the QG-FRDR accuracies from the noise-balanced experiments in Tables 7 and 8, not Gaussian-only evaluations.

Mechanism	rank R	$S_1(R)$	$S_2(R)$	$S_1(R^{\leftarrow})$	$S_2(R^{\leftarrow})$	σ_3/σ_1	Acc. (%)
LQ	2	0.410839	0	0.008749	0.00198826	0.033900	0
L4	2	0.271760	0	0.052354	0.00693861	0.074371	2
LB	2	0.068950	0	0.024789	0.00730781	0.069414	4
LH	2	0.454188	0	0.009403	0.00069790	0.022314	0
LP	2	0.275808	0	0.115084	0.00881207	0.070126	2
QQ	1	0	0	0.004938	0.00008882	0.009046	100
Q4	2	0.052910	0	0.030888	0.00054592	0.021419	100
PB	1	0	0	0.008029	0.00032074	0.017547	100
Linear Gaussian	1	0	0	0	0	0	–

These diagnostics indicate that directional performance depends on both the population gap and direction-dependent estimation error. The Gaussian-noise population calculations should be distinguished from the noise-balanced finite-sample evaluation. QG-FRDR-W performs better on the five linear-location mechanisms, as shown in Tables 7 and 8.

G Implementation details

Grid evaluation. All score integrals are evaluated as weighted averages over the conditioning and quantile grids, using positive weights w_j^x and w_k^u that reflect grid spacing and sum to one. For QG-FRDR, let $\widehat{\boldsymbol{R}}$ contain the values $\widehat{R}(t_j, u_k)$; scale row j by $\sqrt{w_j^x}$ and column k by $\sqrt{w_k^u}$, and denote the singular values of the resulting matrix by $\widehat{\sigma}_1 \geq \widehat{\sigma}_2 \geq \cdots$. We define $\widehat{S}_K^{\mathrm{FRDR}} = \sum_{\ell>K} \widehat{\sigma}_\ell^2 / \sum_{\ell\geq 1} \widehat{\sigma}_\ell^2$,

with value zero when the denominator is zero. This is the finite-grid counterpart of the population score in Eq. (5).

Algorithm 1 accepts any target rank $K < \min(n_x, n_u)$, where n_x and n_u are the numbers of conditioning and quantile grid points. The main experiments use $K = 2$; Appendix F.6.1 also evaluates $K = 1$.

Algorithm 1 Computing the QG scores and selecting a direction

Require: Paired observations $\{(X_i, Y_i)\}_{i=1}^n$ and target rank K (default $K = 2$).
1: $\mathcal{M} \leftarrow \{\text{ANM}, \text{LSNM}, \text{PNL}, \text{PNL-HNM}, \text{FRDR}\}$.
2: **for** each candidate direction $(A, B) \in \{(X, Y), (Y, X)\}$ **do**
3: Replace A by its normalized empirical rank and B by its affine standardization, estimate $\widehat{Q}_{B^{\text{std}}|A^C}$, and compute QG-ANM and QG-LSNM.
4: Replace B by its normalized empirical ranks, estimate $\widehat{F}_{B^C|A^C}$ and $\widehat{Q}_{B^C|A^C}$, and compute $\widehat{R}$ using (4).
5: **if** $\widehat{R}$ is zero on the grid **then**
6: Set QG-PNL, QG-PNL-HNM, and QG-FRDR to zero.
7: **else**
8: Compute QG-PNL and QG-PNL-HNM using (9) and (10), respectively.
9: Compute QG-FRDR at rank K.
10: **end if**
11: Store the five scores as $\widehat{s}_{A\to B}$.
12: **end for**
13: **for** each criterion $m \in \mathcal{M}$ **do**
14: Estimate $\widehat{d}_m$ as the direction with the smaller score, using the tie convention in Appendix G.
15: **end for**
16: **return** $\{\widehat{s}_{X\to Y}, \widehat{s}_{Y\to X}, (\widehat{d}_m)_{m\in\mathcal{M}}\}$.

Preprocessing Empirical copula coordinates use average ranks for ties, divided by $n + 1$. Affine standardization uses the sample mean and the sample standard deviation with denominator $n - 1$.

Because the candidate cause is represented by its empirical copula coordinate for every score, both the Q-based and R-based conditioning grids contain $n_x = 50$ equally spaced points on $[0.01, 0.99]$. The quantile grid contains $n_u = 50$ equally spaced points on the same interval.

For a grid $t_1 < \cdots < t_m$, the default normalized trapezoidal weights are

$$\widetilde{w}_1 = \frac{t_2 - t_1}{2}, \qquad \widetilde{w}_j = \frac{t_{j+1} - t_{j-1}}{2}, \quad \widetilde{w}_m = \frac{t_m - t_{m-1}}{2},$$
$$w_j = \frac{\widetilde{w}_j}{\sum_{\ell=1}^m \widetilde{w}_\ell}.$$

for $2 \le j \le m - 1$. User-supplied weights must be strictly positive and are normalized to sum to one. QG-PNL-HNM uses a separate set of trapezoidal weights on the common interior grid $u_2, \ldots, u_{m-1}$.

Weighted spectral statistics. For normalized quadrature weights, the matrix representing the gridded field in normalized indicator bases is

$$\widehat{\boldsymbol{R}}_w = \text{diag}(\sqrt{w_1^x}, \ldots, \sqrt{w_{n_x}^x})\,\widehat{\boldsymbol{R}}\,\text{diag}(\sqrt{w_1^u}, \ldots, \sqrt{w_{n_u}^u}). \tag{28}$$

Its squared Frobenius norm is the quadrature approximation of $\|R\|_{L^2}^2$, and its singular values give (5). The ratio is defined as zero when the denominator is zero. A nontrivial rank-K statistic requires $K < \min(n_x, n_u)$.

Kernel estimator. At a query point x_0, the Gaussian-kernel weights are

$$w_i(x_0) = \exp\left\{-\frac{1}{2}\left(\frac{x_i - x_0}{h}\right)^2\right\}.$$

Writing

$$A(y, x_0) = \sum_i w_i(x_0)\mathbf{1}\{y_i \le y\}, \qquad B(x_0) = \sum_i w_i(x_0),$$

the conditional distribution estimate is $\widehat{F}(y \mid x_0) = A(y, x_0)/B(x_0)$. The estimated quantile is the smallest ordered response whose cumulative kernel weight is at least $uB(x_0)$. If no bandwidth is supplied, the implementation uses

$$h = \max\{0.9\widehat{\sigma}n^{-1/5}, \varepsilon_{\text{num}}\}, \qquad \widehat{\sigma} = \begin{cases} \min\{s_x, \text{IQR}(x)/1.349\}, & \text{IQR}(x) > 0, \\ s_x, & \text{IQR}(x) = 0, \end{cases}$$

where s_x is the sample standard deviation on the scoring scale. A supplied bandwidth must be positive and finite.

For the default drift-ratio estimator,

$$w_i'(x_0) = w_i(x_0)\frac{x_i - x_0}{h^2}, \qquad \partial_{x_0}\widehat{F}(y \mid x_0) = \frac{A'(y, x_0)B(x_0) - A(y, x_0)B'(x_0)}{B(x_0)^2 + \varepsilon_{\text{num}}},$$

and $\widehat{R}(x, u) = -\partial_x \widehat{F}(y \mid x)|_{y=\widehat{Q}(x,u)}$. If the total kernel weight is numerically zero, unit weight is assigned to the observation nearest to x_0.

Discrete score evaluation. All multidimensional integrals use tensor products of the normalized one-dimensional weights. For QG-LSNM, the implementation avoids constructing the full (x_1, x_2) tensor by using, for $D_v(x) = Q(x, u_0) - Q(x, v)$,

$$\iint \{D_v(x_1)D_w(x_2) - D_w(x_1)D_v(x_2)\}^2 \, d\mu(x_1)d\mu(x_2) = 2\|D_v\|^2\|D_w\|^2 - 2\langle D_v, D_w\rangle^2.$$

For QG-PNL, weighted column inner products are computed through a Gram matrix. A column with squared norm at most ε_{num} times the largest squared column norm is treated as a zero column.

For a possibly nonuniform quantile grid, set $h_k^- = u_k - u_{k-1}$ and $h_k^+ = u_{k+1} - u_k$. On the common interior grid, QG-PNL-HNM uses

$$\widehat{R}_u(u_k) = -\frac{h_k^+}{h_k^-(h_k^- + h_k^+)}\widehat{R}(u_{k-1}) + \frac{h_k^+ - h_k^-}{h_k^- h_k^+}\widehat{R}(u_k) + \frac{h_k^-}{h_k^+(h_k^- + h_k^+)}\widehat{R}(u_{k+1}),$$

$$\widehat{R}_{uu}(u_k) = \frac{2}{h_k^- + h_k^+}\left\{\frac{\widehat{R}(u_{k+1}) - \widehat{R}(u_k)}{h_k^+} - \frac{\widehat{R}(u_k) - \widehat{R}(u_{k-1})}{h_k^-}\right\}.$$

Differentiating the Wronskian and cancelling the mixed terms gives

$$W'_{x_1,x_2} = R_{uu}(x_1, \cdot)R(x_2, \cdot) - R(x_1, \cdot)R_{uu}(x_2, \cdot). \tag{29}$$

If $\widehat{s}_R = 0$, the normalized QG-PNL-HNM and singular-value statistics are set to zero. Otherwise, QG-PNL-HNM is evaluated by first replacing $\widehat{R}$ with $\widehat{R}/\widehat{s}_R$ and then computing the unnormalized discrete statistic; this is algebraically equivalent to division by $\widehat{s}_R^{10}$. Quantities that are nonnegative in exact arithmetic, including Gram determinants and the reduced QG-LSNM expression above, are clipped at zero after floating-point evaluation. The default numerical floor is $\varepsilon_{\text{num}} = 10^{-12}$.

Handling ties. An optional relative tie tolerance $\tau \geq 0$ reports a tie when

$$\left|\widehat{S}_{X\to Y} - \widehat{S}_{Y\to X}\right| \leq \tau \max\{1, |\widehat{S}_{X\to Y}|, |\widehat{S}_{Y\to X}|\}.$$

The default is $\tau = 0$, so only exact numerical equality is reported as a tie; otherwise the smaller score determines the direction.

G.1 Multiscale weak estimation

We fit the drift ratio through integral equations, following the weak-form approach of Messenger & Bortz (2021). For a candidate direction $d = (A, B)$, let $T = F_A(A)$ and $V = F_B(B)$, and write $F_d(v \mid t)$ and $R_d(t, u)$ for the conditional CDF and drift ratio in these coordinates. Since $\partial_t F_d(v \mid t) = -R_d\{t, F_d(v \mid t)\}$, integration by parts gives

$$\int_I w_m'(t)F_d(v \mid t)\, dt = \int_I w_m(t)R_d\{t, F_d(v \mid t)\}\, dt \tag{30}$$

for tests w_m that vanish at the endpoints of I. The integrals use ordinary Lebesgue measure. The identity requires the FRDR representation on the conditional probability range visited by $F_d(v \mid t)$.

Conditional CDF estimation. We replace the observations by $T_i = \mathrm{rank}(A_i)/(n+1)$ and $V_i = \mathrm{rank}(B_i)/(n+1)$, using average ranks for ties. Let $z_i = \Phi^{-1}(T_i)$, where Φ is the standard Gaussian CDF. The estimate is

$$\widehat{F}_d(v \mid t) = \frac{\sum_{i=1}^n k_i(t)\mathbf{1}\{V_i \le v\}}{\sum_{i=1}^n k_i(t)}, \qquad k_i(t) = \exp\left[-\frac{\{z_i - \Phi^{-1}(t)\}^2}{2h_d^2}\right].$$

We use $h_d = \max\{0.9\min(s_z, \mathrm{IQR}(z)/1.349)n^{-1/5}, 0.015\}$, where s_z is the sample standard deviation, and clip the estimated CDF to $[10^{-6}, 1-10^{-6}]$. The normal-score transformation determines kernel distances; the weak integrals remain in the uniform coordinate t.

Multiscale tests. Motivated by the multiscale construction in Bortz et al. (2023), we form compactly supported bumps $b_j(t) = \eta\{2(t-c_j)/\rho_j\}$, where $\eta(s) = \exp\{-s^2/(1-s^2)\}$ for $|s| < 1$ and is zero otherwise. At each width $\rho_j \in \{0.08, 0.16, 0.32, 0.64\}$, we place 32 equally spaced centers so that every support lies within $I = [0.02, 0.98]$. The different widths measure discrepancies over different portions of the cause interval.

Let ω_g be the trapezoidal weights on the cause grid t_g. Normalize each bump so that $\sum_g \omega_g b_j'(t_g)^2 = 1$. For $G_{jg} = \sqrt{\omega_g}\, b_j'(t_g)$, compute $G = U_b \Sigma V_b^\top$ and retain the leading $n_w = 48$ components. The tests are $\mathbf{w}(t) = \Sigma_{48}^{-1} U_{b,48}^\top \mathbf{b}(t)$, where $\mathbf{b}$ collects the normalized bumps. This gives $\sum_g \omega_g \mathbf{w}'(t_g)\mathbf{w}'(t_g)^\top = I_{n_w}$. The construction is independent of the observations and is shared by both directions.

Rank-constrained fitting. Let ψ and $\boldsymbol{\chi}$ be cubic B-spline bases on I and $[0,1]$, respectively. We represent the field as $R_M(t,u) = \psi(t)^\top M \boldsymbol{\chi}(u)$ with $\mathrm{rank}(M) \le K$. For response-rank thresholds $v_1, \dots, v_{n_v}$, define

$$\begin{aligned} \widehat{L}_{d,\ell} &= \int_I \mathbf{w}'(t)\widehat{F}_d(v_\ell \mid t)\, dt, \\ \widehat{D}_{d,\ell}(M) &= \int_I \mathbf{w}(t)\, \psi(t)^\top M \boldsymbol{\chi}\{\widehat{F}_d(v_\ell \mid t)\}\, dt, \\ r_{d,\ell}(M) &= \widehat{L}_{d,\ell} - \widehat{D}_{d,\ell}(M). \end{aligned} \tag{31}$$

Both sides use the same estimated CDF, without differentiating it. A single matrix M is fitted to all retained tests.

For weighting, set $\widehat{\gamma}_d(t) = n_v^{-1}\sum_\ell \widehat{F}_d(v_\ell \mid t)\{1 - \widehat{F}_d(v_\ell \mid t)\}$ and $\widehat{C}_d = \int_I \mathbf{w}'(t)\mathbf{w}'(t)^\top \widehat{\gamma}_d(t)\, dt$. This is a covariance proxy. Pool the directions as $\widehat{C} = (\widehat{C}_{X\to Y} + \widehat{C}_{Y\to X})/2$. With $\tau = \mathrm{tr}(\widehat{C})/n_w$, define

$$C_\delta = (1-\delta)\widehat{C} + \delta\tau I_{n_w}, \qquad \Omega = \frac{n_w C_\delta^{-1}}{\mathrm{tr}(C_\delta^{-1})}, \qquad \delta = \tfrac{1}{2}.$$

If $\widehat{C} = 0$, use $\Omega = I_{n_w}$.

We approximately minimize

$$\mathcal{E}_d(M) = \frac{1}{n_w n_v}\sum_{\ell=1}^{n_v} r_{d,\ell}(M)^\top \Omega r_{d,\ell}(M), \qquad \mathrm{rank}(M) \le K, \tag{32}$$

without a coefficient penalty. We enforce the rank constraint by writing $M = UV^\top$, where each factor has K columns. We first reduce the least-squares system using QR factorization, then optimize U and V with the Levenberg–Marquardt algorithm. The derivatives of the residuals with respect to both factors are computed analytically.

Directional score. For the fitted matrix $\widehat{M}_d$, we use Eq. (6) and select the direction with the smaller score.

We use $K = 2$, 30 splines per axis with uniformly spaced interior knots, 800 equally spaced cause points on I, and 35 equally spaced response-rank thresholds on $[0.01, 0.99]$. All integrals are evaluated by trapezoidal quadrature. The same observations are used for fitting and scoring.

G.2 Computational environments

The experiments were conducted in two computing environments according to the hardware requirements of the evaluated methods.

Local CPU environment. All methods except QPE-f lowrank were evaluated in the local CPU environment, including QG-FRDR, QG-FRDR-W, the model-specific QG scores, LRQS, CVEL, QPE-k, ANM, and DIVOT. The environment used Windows 11 with an 11th Gen Intel(R) Core(TM) i7-1165G7 processor at 2.80 GHz (4 cores and 8 threads) and 8.0 GB of memory operating at 3200 MT/s.

The proposed QG-FRDR methods were implemented in Python 3.11.9 using NumPy 2.4.6 and SciPy 1.17.1.

GPU environment. QPE-f lowrank was evaluated in a separate GPU environment equipped with a single NVIDIA RTX PRO 6000 Blackwell Max-Q Workstation Edition GPU with 97,887 MiB (approximately 95.6 GiB) of GPU memory. The GPU environment used Python 3.10.12 and PyTorch `2.2.0a0+81ea7a4`.