BEYOND AVERAGE EFFECTS: IDENTIFYING SIGNED CAUSAL EFFECTS UNDER LATENT CONFOUNDING VIA MULTI-ENVIRONMENT VARIANCE SHIFTS

Anonymous authors
Paper under double-blind review

ABSTRACT

A positive association between overnight commodity futures returns and equity opening returns does not establish causal transmission: both markets may react to the same unobserved news. Even a positive average causal effect can conceal positive and negative effects that partly cancel. We study what changes in volatility across a few environments can reveal about these two questions. Our model allows the causal slope to vary across observations and treatment to be confounded. Conditional on a pretreatment state, environments change only the scale of the treatment disturbance, while the joint distribution of the causal slope and latent common components remains unchanged. Under the model's assumptions, contrasts in second moments across two environments identify the mean and second moment of the causal slope. These moments yield exact joint bounds on its expected positive and negative parts: every pair in the resulting set is attained by a model matching the observed moments. With three environments having distinct, positive, and unknown scales, additional assumptions allow both signed effects to be point identified without recovering the full slope distribution. Simulations assess estimation and inference. In an application to Shanghai tin futures and Yunnan Tin equity returns, the estimated average causal effect is approximately 29 basis points per percentage-point tin innovation, after averaging over the observed state composition. Yet the moment-based bounds allow a negative effect with a magnitude as large as 55 basis points per percentage point—nearly twice the average effect. The same estimated moments thus admit both transmission that is nonnegative throughout and substantial cancellation between positive and negative effects.

1 INTRODUCTION

When tin futures rise while the stock market is closed, does a tin producer open higher the next morning because of it? In our motivating data, a one-percentage-point larger overnight Shanghai tin futures return is associated with about 40 basis points of additional opening-gap return for Yunnan Tin Co., Ltd. (000960). Adding other metal returns and lagged controls barely moves this number (Section 2, "A persistent, state-varying association"). This statistical association is strong, but it does not establish transmission by itself.

The obstacle is that both the futures and equity markets react to the same news. Tin futures trade longer hours than the equity does and play a price-discovery role (Hasbrouck, 2003), and they share exposure with tin equities to macroeconomic shocks, industry expectations, and liquidity conditions (Pástor & Stambaugh, 2003). A regression on tin therefore captures whatever moves both markets, and controlling for what we observe need not remove that common part, because the common drivers are themselves unobserved. The association leaves two things unsettled: how much of it is transmission rather than common movement, and whether a positive average comes from effects that are positive, or from positive and negative effects that partly cancel (Kilian & Park, 2009).

The data on the Shanghai tin future contain information that the raw association does not use: how much the overnight tin return moves around is not constant over the sample. Its dispersion is nearly the same in the first two calendar periods and roughly three and a half times larger in the third; the pe-

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.

riods are separated by changes in the exchange's fee rules (Section 2, "What environmental variation could reveal"). Here we refer to these distinct periods as environments. The use of heterogeneous environments to reveal causal structure is also central in recent work on causal discovery and causal representation learning from multiple domains or distributions (Huang et al., 2020; Zhang et al., 2024); in our setting the environments are not used to recover a causal graph or a latent representation, but to shift the scale of the treatment disturbance. What can a small number of environments still tell us about the size of a causal effect, and about its sign? Differences between environments then identify two numbers: the average causal effect and its dispersion. This holds even when the treatment is confounded, because the confounded part is common to every environment and cancels in the difference. Our goal is to state exactly what these two numbers imply about the sizes of positive and negative effects, and to quantify what remains unknown. The approach builds on identification through heteroskedasticity (Rigobon, 2003; Lewbel, 2012) and on random-coefficient models with correlated coefficients (Masten & Torgovitsky, 2016; Masten, 2018), but our method uses neither an instrumental variable nor a known scale: the identifying variation is the change in the scale of the treatment disturbance across environments.

Contributions. **First, identification without instrumental variables.** What identifies the slope is the change in the scale of the treatment disturbance across environments. The mean and the second moment of the slope follow from variance shifts alone. Two environments suffice for both, and a third makes the relations over-identified, which gives the model testable content. **Second, beyond the average: exact bounds on signed effects.** A positive average is consistent both with effects that are positive throughout and with large positive and negative effects that partly cancel. We describe the full set of positive and negative pairings the identified moments admit, show that every point of that set is realized by some compatible model. **Third, point identification of the signed effects from three or more environments.** Three environments with distinct and unknown scales identify the positive and negative contributions as points rather than as a set. A characteristic-function argument supplies the one quantity that the second moments leave undetermined, the mean absolute effect, and hence separates the two signed effects.

2 Empirical Motivation: Futures–Equity Associations and Variance Shifts

This section presents an empirical motivation for our identification strategy.

Pretreatment information, the treatment and outcome. Let d be an equity trading date and d^- the preceding date. We assign each futures night session to the equity trading date that follows it. Let $P^{\mathrm{SN}}_{d,t}$ be the Shanghai tin futures price on date d at time t, and let $\mathrm{Open}_{000960,d}$ and $\mathrm{Close}_{000960,d}$ be the opening and closing prices of Yunnan Tin Co., Ltd. (000960) on date d. The raw tin return, in percentage points, and the equity opening gap, in basis points, are

$$R^{\mathrm{SN}}_d = \log P^{\mathrm{SN}}_{d,\text{night close}} - \log P^{\mathrm{SN}}_{d^-,15:00}, \tag{1}$$

$$Y^{\mathrm{gap}}_d = 10^4 \left\{ \log \mathrm{Open}_{000960,d} - \log \mathrm{Close}_{000960,d^-} \right\}. \tag{2}$$

The tin night session closes at 01:00 and the equity market opens at 09:30, so the treatment window starts at the preceding equity close, 15:00 on d^-. Information that arrives after the futures close can still enter the opening gap.

Let W^{treat}_d collect the pretreatment controls that enter the adjusted treatment X_d in equation 3; all of them are available at the start of the treatment window (Appendix D). Let M_d collect the returns of the other three metal futures — copper, aluminum and zinc — measured over the same window as R^{SN}_d. We fit a rolling linear regression of R^{SN}_d on M_d and W^{treat}_d on completed dates and write $\widehat{g}_d(\cdot)$ for the fitted rule (Appendix D). The adjusted treatment is

$$X_d = 100 \left\{ R^{\mathrm{SN}}_d - \widehat{g}_d\big(M_d, W^{\mathrm{treat}}_d\big) \right\}. \tag{3}$$

What remains in X_d is the part of the overnight tin return that the other metals and prior information do not explain, so it plays the role of a tin-specific surprise: the object a producer's equity is exposed to (Tufano, 1998), and the kind of shock whose identity matters for equity responses (Kilian & Park, 2009). The value $\widehat{g}_d(M_d, W^{\mathrm{treat}}_d)$ is what the rolling fit predicts for date d from the same-window

metal returns and the pretreatment controls. Its coefficients come from completed dates, while M_d is realized in the current window, so X_d is available only after the window closes. Appendix D documents the window length, the control set, and the units.

What this section reports is an association, not a causal effect. The residual X_d is orthogonal to the variables we observe, not to the ones we do not, so a causal reading has to come from the model in Section 3 rather than from the construction of the treatment.

A persistent, state-varying association. Table 1 reports three OLS regressions of the opening gap on the adjusted treatment, $Y_d^{\text{gap}} = \alpha + \beta X_d + \gamma' C_d + \varepsilon_d$, estimated on the 615 trading dates for which both the treatment and the conditioning state are defined (Appendix D). The reported coefficient is β, the opening-gap response to a one-percentage-point change in X_d. The controls C_d enter cumulatively — none, then the same-window metal returns M_d, then the lagged and calendar controls. The treatment coefficient stays between 40 and 42 basis points per percentage point, while the added controls raise the explained variation substantially.

Table 1: Motivating regressions on 615 trading dates. The treatment is the adjusted tin return in percentage points; the outcome is the equity opening gap in basis points. Controls are added cumulatively. Standard errors are Newey–West with five lags (Newey & West, 1987).

Specification	$\widehat{\beta}$	Newey–West SE	p	R^2
No controls	42.05	9.79	< 0.001	0.112
+ same-window metals	40.32	7.33	< 0.001	0.378
+ lagged and calendar controls	40.19	7.17	< 0.001	0.407

The slope also differs across pretreatment states. Estimating the last row of Table 1 separately for the two states gives 27.34 after a negative overnight return against 47.40 after a positive one, on the same 615 dates. Two different slopes need not mean two different causal effects, and two positive slopes still say nothing about signs: the average can be positive while some dates are negative.

Calendar variation in treatment dispersion. Table 2 reports the dispersion of the treatment in three calendar periods: nearly the same in the first two, 1.16 and 1.04, and roughly three times larger in the third, 3.56. The periods are separated by two SHFE fee adjustments (Shanghai Futures Exchange, 2024; 2025), and we label them Era I, Era II, and Era III. Appendix D records their boundary dates, the products and contracts they cover, and why they are candidate rather than validated sources of variation.

Table 2: Era-wise moments of the fixed-scale treatment X (percentage points) and the opening gap Y (basis points). Entries are sample moments.

Environment	n	$\mathrm{Var}(X)$	$\mathrm{Cov}(X,Y)$	$\mathrm{Var}(Y)$
Era I: before Jun 2024	102	1.16	59.5	14,252
Era II: Jun 2024–Dec 2025	352	1.04	49.6	16,086
Era III: from Dec 2025	161	3.56	131.8	59,546
Eras I–II pooled	454	1.07	51.6	15,647

Because Eras I and II have similar dispersion, their contrast carries little identifying variation, so the primary specification pools the two and contrasts the pool with Era III. This choice follows from inspecting the sample; it is not a claim that Eras I and II share a latent mechanism.

3 Problem Setup and Model

Notation. Let $E \in \{1, \dots, K\}$ be the observed environment and W an observed pretreatment state. The observed treatment and outcome in environment e are (X_e, Y_e), and $Y(x)$ is the potential outcome under the intervention $X \leftarrow x$. All moments below condition on the same state $W = w$.

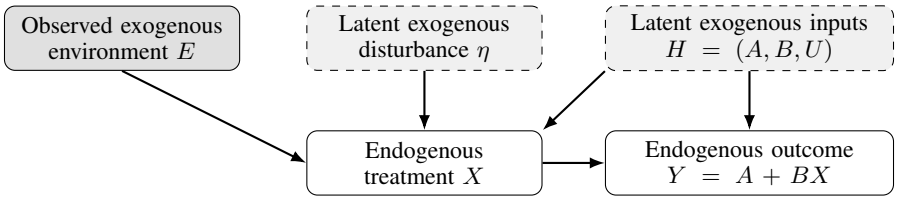


Figure 1: Structural assignments at a fixed $W = w$. Dashed boxes denote latent inputs; solid boxes denote observed variables. Arrows show functional dependence, with E entering only through the treatment-disturbance scale.

We write $\mathbb{E}_w$, Var_w, and Cov_w for means, variances, and covariances given $W = w$; and $\mathbb{E}_{e,w}$, $\mathrm{Var}_{e,w}$, and $\mathrm{Cov}_{e,w}$ for the corresponding moments given $E = e$ and $W = w$.

Structural assignments. We use the familiar correlated random-coefficient outcome equation for heterogeneous treatment effects with an endogenous regressor (Masten & Torgovitsky, 2016; Masten, 2018), but pair it with a different source of identification. The random-coefficient model in Figure 1 has the structure:

$$X_e = U + s_e(w)\eta, \qquad Y(x) = A + Bx, \qquad Y_e = Y(X_e). \tag{4}$$

Here η is a standardized noise, and $s_e(w)$ is the environment-specific scale. $A = Y(0)$ is the latent baseline outcome and $B = Y(x+1) - Y(x)$ the latent causal effect of a unit treatment increment. The latent inputs A, B and U are jointly unrestricted: a common latent factor can make the hidden component of treatment U dependent on the baseline outcome A and on the causal slope B, which is how confounding enters the model. The environment acts only through the scale $s_e(w)$ of η. The auxiliary quantity $C = A + BU = Y(U)$ is the baseline outcome that remains once the environment-specific disturbance is removed.

Relation to instrumental-variable (IV) models. The model in equation 4 is the correlated random coefficient model: the treatment is endogenous because U may be correlated with A and B, and the slope varies across observations. The model differs from the conventional IV models. Under monotonicity an instrument identifies the average effect for compliers (Imbens & Angrist, 1994); richer variation identifies average coefficients (Masten & Torgovitsky, 2016), the effect along margins of the instrument (Heckman et al., 2006), or, under rank conditions, the slope distribution itself (Masten, 2018). We use no instrument, and our environments could not serve as one: an instrument works by moving the treatment, and ours does not. Within a state, the treatment in our model has the same average in every environment, but not the same variability. Identification therefore relies on stability and exclusion across environments rather than on instrument validity and rank conditions.

Assumption 1 (Stable latent mechanism)**.** *For almost every state w and every pair of environments e, f being compared, the conditional distribution of (A, B, U, η) given $E = e, W = w$ is the same as the conditional distribution given $E = f, W = w$.*

At a fixed state, the latent inputs and the disturbance share one joint distribution across the environments being compared, and that distribution may differ from state to state.

Assumption 2 (Independent standardized disturbance)**.** *For every environment and state, we assume that $\mathbb{E}_{e,w}[\eta] = 0, \mathbb{E}_{e,w}[\eta^2] = 1$, and $\eta \perp (A, B, U) \mid (E = e, W = w)$.*

Here the dependence among A, B, U is unrestricted.

Assumption 3 (Environmental exclusion)**.** *The scales $s_e(w)$ are finite, deterministic, and nonnegative. The environment enters the structural assignments only through $s_e(w)$.*

Assumption 4 (Finite moments)**.** $\mathbb{E}_w[U^2 + C^2 + B^2] < \infty$.

These assumptions allow the treatment to be confounded: U is a component of X_e and also enters the outcome through $C = A + BU$, so dependence between U and A or B makes X_e correlated to the potential outcomes even at a fixed state. The independence of η from (A, B, U) is only a baseline convenience: it is sufficient but not necessary for the moment decomposition below, and Section 4 replaces it with Condition M, which allows η to depend on (A, B, U).

Causal targets. The first two conditional slope moments and the slope variance are denoted by

$$\boxed{\mu(w) = \mathbb{E}_w[B],\ q(w) = \mathbb{E}_w[B^2],\ \sigma_B^2(w) = q(w) - \mu(w)^2,\ t(w) := \mathbb{E}_w[|B|].}$$

To distinguish effects of opposite signs, let $x_+ = \max\{x, 0\}$ and define $P(w) = \mathbb{E}_w[B_+], N(w) = \mathbb{E}_w[(-B)_+]$. Thus $\mu(w) = P(w) - N(w)$. The quantity $t(w) := \mathbb{E}_w[|B|] = P(w) + N(w)$ is the mean absolute effect.

Observable information. At a fixed state w, each environment e gives the covariance matrix of treatment and outcome,

$$\Sigma_e(w) = \begin{pmatrix} v_e(w) & c_e(w) \\ c_e(w) & r_e(w) \end{pmatrix},$$

where $v_e(w) = \mathrm{Var}_{e,w}(X_e), c_e(w) = \mathrm{Cov}_{e,w}(X_e, Y_e), r_e(w) = \mathrm{Var}_{e,w}(Y_e)$. The three entries are the observable inputs to Theorem 1, and Appendix A uses the matrix form. Theorem 3 uses more than second moments: the fourth cumulants of X_e and Y_e, and, in the varying-amplitude branch, even their characteristic functions.

4 Identification Theory

In this section we state the identifying conditions and derive what they identify. We first show that environment contrasts identify the mean and second moment of the slope (Theorem 1), then derive the exact joint set of the signed effects (Theorem 2), and finally identify the contributions themselves under the stronger assumptions of Theorem 3.

Moment condition. We introduce one moment condition to relax the independence assumption of Section 3. Write $V = (\eta, B\eta)^\top$ formed by the standardized shock and its outcome counterpart.

Condition M.

$$\mathbb{E}_w[V] = 0, \qquad \mathrm{Var}_w(\eta) = 1, \qquad \mathrm{Cov}_w(\eta, B\eta) = \mu, \qquad \mathrm{Var}_w(B\eta) = q,$$
$$\mathrm{Cov}_w(U, \eta) = \mathrm{Cov}_w(C, B\eta) = 0, \qquad \mathrm{Cov}_w(U, B\eta) + \mathrm{Cov}_w(C, \eta) = 0.$$

The first part is a normalization: it fixes the scale of η and sets the second moments of the pair $(\eta, B\eta)$ equal to the slope moments $\mathbb{E}_w[B]$ and $\mathbb{E}_w[B^2]$. The second part is the substantive restriction: U is uncorrelated with η, C is uncorrelated with $B\eta$, and the two remaining cross-covariances cancel. Only this cancellation is required, so η may remain dependent on (A, B, U). Assumption 2 implies Condition M, so the baseline model of Section 3, in which the disturbance is independent of (A, B, U), is a special case. Condition M is also exactly what makes the observable moments depend on the environment only through the scale: the mean of (X_e, Y_e) is the same in every environment, and its variance, its covariance and the outcome variance all grow linearly in $s_e(w)^2$. Proposition A.1 states and proves that equivalence.

Theorem 1 (Moment identification)**.** *Under Assumptions 1, 3, and 4, together with Condition M, given two environments e, f satisfying if $v_e \neq v_f$ then μ, q, σ_B^2 are identified by*

$$\mu = \frac{c_e - c_f}{v_e - v_f}, \qquad q = \frac{r_e - r_f}{v_e - v_f}, \qquad \sigma_B^2 = q - \mu^2. \tag{5}$$

Here the confounding sits in the environment-invariant intercepts, the first two moments of the slope B are computed from second moments of treatment and outcome. The same values are obtained from every environment pair with a nonzero treatment-variance difference, so with three or more environments the two relations are over-identified and their agreement is a testable restriction.

Relation to existing identification results. The two equations of Theorem 1 are ratios of environment contrasts. The first, $\mu = (c_e - c_f)/(v_e - v_f)$, is the familiar heteroskedasticity contrast (Section 6); the second, $q = (r_e - r_f)/(v_e - v_f)$, is what a random slope adds. Instrumental-variable results for this model identify a different target from different information: a binary instrument gives a weighted average over compliers (Imbens & Angrist, 1994), rank conditions give average coefficients (Masten & Torgovitsky, 2016), and sufficiently many instruments give the coefficient law itself (Masten, 2018). We need neither an instrument nor an instrument-rank condition; we require instead that the environment shift the scale $s_e(w)$ while the latent mechanism stays fixed, and both moments can be identified by two environments with different treatment variances.

From the slope moments to the signed effects. Next, we study what the identified moments imply about the signed effects. Three special cases determine the signed effects. First, if $q = \mu^2$, then $B = \mu$ almost surely and $(P, N) = (\mu_+, (-\mu)_+)$. Second, if B is known to be nonnegative or nonpositive, then $t = |\mu|$ and one contribution vanishes. Third, when $\sigma_B = \sqrt{q - \mu^2} > 0$, specifying the distribution of the standardized slope $Z = (B-\mu)/\sigma_B$ determines $t = \mathbb{E}_Z[|\mu+\sigma_B Z|]$ from (μ, q). For example, under $Z \sim N(0,1)$, we have $P = \sigma_B\phi(z) + \mu\Phi(z), N = \sigma_B\phi(z) - \mu\Phi(-z)$, where $z = \frac{\mu}{\sigma_B}$, and ϕ and Φ are the standard normal density and distribution functions.

Without shape information, the identified moments μ and q capture a range for $\mathbb{E}_w[|B|]$ rather than a value, and the next result gives that range in closed form. The data enter only through Theorem 1, which delivers μ and q from the observed moments v_e, c_e, r_e; the range is then a statement about which distributions of B share those two moments, and it is exact because every such distribution can be embedded in a model that reproduces the observed moments (Theorem A.4). Throughout we assume that the observed environment moments are compatible with the model (Proposition A.2) and that at least two comparable environments have distinct treatment variances.

Theorem 2 (Joint sharp set for signed effects)**.** *Among compatible models reproducing the observed means and covariance matrices, the attainable values of $t = \mathbb{E}_w[|B|]$ form exactly the set*

$$\mathcal{T}(\mu, q) = \begin{cases} \{0\}, & q = 0, \\ [|\mu|, \sqrt{q}], & q > 0,\ \mu \neq 0, \\ (0, \sqrt{q}], & q > 0,\ \mu = 0. \end{cases} \tag{6}$$

Consequently the attainable pairs of signed effects, $P = \mathbb{E}_w[B_+]$ and $N = \mathbb{E}_w[(-B)_+]$, are exactly the points of the joint sharp set $\mathcal{I}_{PN}(w)$:

$$\mathcal{I}_{PN}(w) = \left\{ \left(\frac{t+\mu}{2}, \frac{t-\mu}{2} \right) : t \in \mathcal{T}(\mu, q) \right\}. \tag{7}$$

Every compatible model produces a pair in this set, and every pair in the set is attained by a compatible model.

The joint sharp set $\mathcal{I}_{PN}(w)$ is a segment on the line $P - N = \mu$ in the (P, N) plane, running from $\big((|\mu| + \mu)/2, (|\mu| - \mu)/2\big)$ at $t = |\mu|$ to $\big((\sqrt{q} + \mu)/2, (\sqrt{q} - \mu)/2\big)$ at $t = \sqrt{q}$ (Figure 2 in Appendix A). The upper endpoint is always attainable, but the lower endpoint is not attainable when $\mu = 0 < q$. Appendix A proves the statement and Appendix B supplies the sampling theory for the moment estimators. The next result identifies the signed effects rather than bounding them. It adds two tools to Theorem 2 — the fourth cumulants of treatment and outcome and, in the varying-amplitude case, their characteristic functions. It also needs three environments rather than two: the fourth-cumulant step interpolates a quadratic in the treatment variance, and three distinct variances determine a quadratic.

Theorem 3 (Point identification of the signed effects)**.** *Fix a state w and three environments e, f and g with positive conditional probability and observed support, finite positive scales $0 < s_e, s_f, s_g < \infty$, and pairwise distinct treatment variances. Suppose that (A1) Assumptions 1 and 3 hold; (A2) the disturbance η satisfies Assumption 2 and is symmetric about zero; (A3) $C \perp |B|$ at w; (A4) $\mathbb{E}_w[|U|^4 + |C|^4 + |B|^4 + |\eta|^4] < \infty$. Then:*

- *(C1) The slope moments μ and q are identified by equation 5.*
- *(C2) The dispersion $\xi := \mathrm{Var}_w(B^2) \geq 0$ is identified from the fourth cumulants of treatment and outcome.*
- *(C3) The signed effects are point identified in either of two cases.*
  - *(C3a)* Constant amplitude. *If $\xi = 0$, then $B^2 = q$ almost surely and $t = \mathbb{E}_w[|B|] = \sqrt{q}$.*
  - *(C3b)* Varying amplitude. *If $\xi > 0$ and the characteristic functions of X_e and Y_e vanish nowhere, then, using any two of the three environments with distinct scales,*

$$t = \frac{J_Y}{J_X}, \qquad J_Z = \int_0^\infty \frac{1 - \Re\Psi_Z(u)}{u^2}\, du,$$

  where Ψ_Z is the characteristic function of the scaled disturbance recovered from the environment ratios (Appendix A).

In either case $P = (t + \mu)/2$ and $N = (t - \mu)/2$.

5 EXPERIMENTS

5.1 MOMENT-BASED ESTIMATION AND INFERENCE

Design. We evaluate the feasible-GLS minimum-distance estimator of (μ, q) and the Wald inference based on it, following standard minimum-distance and asymptotic inference methods (Hansen, 1982; Newey & McFadden, 1994); Appendix B establishes its asymptotic normality and feasible covariance under the stated sampling conditions. The estimator matches the cross-environment contrasts in $\mathrm{Cov}(X, Y)$ and $\mathrm{Var}(Y)$ to the contrasts in $\mathrm{Var}(X)$, weighting with an estimated inverse-covariance matrix and imposing $q \geq \mu^2$.

The data-generating process is the three-environment model of Section 3, with a latent baseline correlated with the treatment level and a Gaussian slope. Total sample sizes of $n = 500$ to $20{,}000$ are allocated across the three environments in proportions $\pi = (0.145, 0.589, 0.266)$, close to the era shares of the application. The designs differ in how far apart the treatment variances are, which we call the contrast: the strong contrast has $v = (1.20, 1.05, 4.02)$, the weak one $v = (0.85, 0.90, 0.95)$, and the zero one identical variances. We report the strong contrast in Table 3; the weak and zero contrasts, the comparison with identity weighting and with the environment-ratio estimator, and the behaviour of the identified set when the contrast is weak or absent are in Appendix C, which also gives the generating mechanism, the parameters and the Monte Carlo settings.

Table 3: Moment estimation and Wald coverage under the strong-shift design; n is the total sample size, allocated across the three environments in proportions $\pi = (0.145, 0.589, 0.266)$, and "MC reps" is the number of Monte Carlo replications. Coverage is evaluated for nominal 95% marginal intervals and the joint Wald ellipse for (μ, q) using the estimated feasible-GLS covariance.

n	MC reps	RMSE($\hat{\mu}$)	RMSE($\hat{q}$)	cov(μ)	cov(q)	joint cov
500	1,200	0.196	0.360	0.949	0.917	0.884
1,000	800	0.144	0.251	0.940	0.948	0.910
2,000	500	0.098	0.190	0.948	0.936	0.924
5,000	300	0.060	0.112	0.967	0.947	0.957
20,000	150	0.029	0.055	0.967	0.960	0.960

Results. Both $\hat{\mu}$ and $\hat{q}$ become more accurate as n increases, and joint Wald coverage rises from 88.4% to 96.0% (Table 3); at the two largest sample sizes the rates are compatible with nominal coverage. Appendix C compares the three estimators — feasible-GLS minimum distance, identity-weighted minimum distance and the unweighted environment ratio — under the strong shift (Figure 3) and reports the weak- and zero-contrast results.

5.2 POINT IDENTIFICATION OF THE SIGNED EFFECTS

Design. We evaluate signed-effect estimation in designs satisfying Theorem 3. The baseline (I) has $(\mu, q, \xi, t) = (0.4225, 1.36, 1.44, 1)$ with $\xi = \mathrm{Var}(B^2)$ the dispersion of the slope's amplitude and $t = \mathbb{E}|B|$, and $(P, N) = (0.7113, 0.2887)$. Its moment-based sharp set for N is $[0, 0.372]$, exact up to rounding and attained at the lower endpoint by a slope of known positive sign, so second moments alone cannot exclude $N = 0$. Designs II–IV respectively use standardized Laplace noise, a baseline-dependent sign assignment with $\mu = 0$, and constant amplitude $|B| \equiv 1$; Appendix C gives the generating equations and assumption checks.

Estimator. When $\hat{\xi} \leq 0$ we use $\hat{t} = \sqrt{\hat{q}}$; when $\hat{\xi} > 0$ we use the characteristic-function estimator if scale recovery is feasible and some environment pair yields a usable ratio, and otherwise report $\hat{t} = \sqrt{\hat{q}}$ as a flagged *endpoint substitution*; we compute $\hat{N} = (\hat{t} - \hat{\mu})/2$. Endpoint substitution selects the upper moment bound, not evidence that $N > 0$; Appendix C gives the details.

Results. In Design I, raising n from $1{,}000$ to $16{,}000$ reduces endpoint substitution from 15% to none and lowers RMSE($\hat{N}$) from 0.063 to 0.031 (Table 4); the bias of $\hat{t}$ remains near 0.055 at the two larger sample sizes; Appendix C attributes most of this offset to truncation of the Fourier integral

Table 4: Finite-sample estimation of signed effects. Here n is the sample size per environment, so each replication uses $3n$ observations in total. Biases and RMSEs include all replications: 100 per row for Design I and 60 per row for Designs II–IV. "CF" and "endpoint" give the shares using characteristic-function estimation and flagged endpoint substitution; the remaining share uses the $\hat{\xi} \leq 0$ formula.

Design	n	bias($\hat{\mu}$)	bias($\hat{t}$)	bias($\hat{N}$)	RMSE($\hat{N}$)	CF	endpoint
I baseline	1,000	$+0.005$	$+0.090$	$+0.043$	0.063	0.85	0.15
I baseline	4,000	$+0.002$	$+0.053$	$+0.026$	0.034	0.99	0.01
I baseline	16,000	-0.001	$+0.056$	$+0.029$	0.031	1.00	0.00
II Laplace η	16,000	$+0.003$	$+0.014$	$+0.005$	0.015	1.00	0.00
III $\mu = 0$	16,000	-0.001	$+0.058$	$+0.029$	0.031	1.00	0.00
IV $\xi = 0$	16,000	-0.001	$+0.000$	$+0.001$	0.006	0.08	0.42

rather than to identification. In the zero-mean design (III) the sharp set is widest, $(0, 0.583]$, and $\hat{N}$ averages 0.529 against $N = 0.500$. In Design IV, 92% of replications use $\sqrt{\hat{q}}$; the substitution is exact there because $\xi = 0$ makes $t = \sqrt{q}$, and $t = 1$ because $|B| \equiv 1$, whereas under genuine amplitude variation it would report the largest admissible N, not evidence that $N > 0$.

In the baseline design the moment-based sharp set for N is $[0, 0.372]$, which contains $N = 0$ (Table 7), while the characteristic-function estimator returns t and N with biases $+0.056$ and $+0.029$ at $n = 16{,}000$ (Table 8); Appendix C reports the branch shares behind those numbers.

5.3 Real-Data Demonstration: Shanghai Tin Futures and Yunnan Tin Stock

Design. The treatment is the fixed-scale overnight tin innovation of Section 2 and the state is the sign of the earlier, completed overnight tin return. The primary specification pools Eras I and II into one low-variance environment and contrasts it with Era III; Appendix D documents the construction.

Results. Table 5 reports the estimated slope moments, the projections of the joint sharp set, and Gaussian benchmarks — the values the identified moments would imply for P and N if the standardized slope were normal (Section 4) — which the caption reports as shape benchmarks rather than as estimates. The application uses the moment results of Theorems 1 and 2; the point-identification route of Theorem 3, evaluated in Section 5.2, is not applied to these data, and Appendix D documents that choice by simulating the estimator at the sample sizes observed in each state.

Table 5: State-conditional effects for Shanghai tin futures on Yunnan Tin stock. Units are basis points per percentage-point tin innovation. Bounds are projections of the joint sharp set and are not confidence intervals; Gaussian columns are shape benchmarks.

State	Environments	μ	P bounds	N bounds	Gaussian P	Gaussian N
Down	Eras I, II, III	19.38	$[19.38, 87.85]$	$[0.00, 68.47]$	72.06	52.68
Up	Eras I, II, III	32.59	$[32.59, 75.34]$	$[0.00, 42.76]$	63.43	30.85
Down	Eras I–II pooled	19.51	$[19.51, 89.07]$	$[0.00, 69.57]$	73.04	53.54
Up	Eras I–II pooled	36.72	$[36.72, 78.39]$	$[0.00, 41.67]$	66.30	29.57

The estimated mean slopes are positive in both states, but the signed effects remain only partially identified: negative contributions range from zero to 69.57 bp/pp in the down state and to 41.67 bp/pp in the up state. These ranges are projections of joint segments satisfying $P - N = \mu$, so larger negative contributions require correspondingly larger positive ones, and the three-era specification preserves these conclusions. Aggregating over the observed state composition as in Corollary A.5 gives an overall average slope of 28.63 bp/pp and an aggregated negative-contribution bound of 54.78 bp/pp (Appendix D), so a positive average and the possibility of offsetting negative contributions coexist at the population level as well.

6 Related Work

Identification from environmental variation. A related causal-discovery literature uses changes across domains, environments or time as the source of identification. Ghassami et al. (2018) learn causal directions in linear systems from independent changes of the causal modules across domains, and Huang et al. (2020) develop CD-NOD for heterogeneous or nonstationary data, where distribution shifts identify changing mechanisms and orient edges. More recent work extends this perspective to causal representation learning from multiple distributions (Zhang et al., 2024). Our use of environments is narrower: we do not seek a causal graph or a latent representation, but use environment-specific shifts in the variance of the treatment disturbance to identify functionals of a random slope under latent confounding. Our mean-slope formula is an instance of identification through heteroskedasticity (Rigobon, 2003; Lewbel, 2012). Kivva et al. (2025) identify a scalar treatment effect in a linear latent-confounder model from higher-order moment changes: in their treatment-noise branch the effect is recovered uniquely from treatment and outcome moments, and that second-moment contrast coincides with the mean-slope formula used here. Their object is a single effect assumed invariant across environments, whereas here the slope is random and only its distribution is stable, so the same contrasts identify the slope's second moment and, under Theorem 3, its signed effects, while the hidden treatment component may be correlated with the slope; when $q = \mu^2$ the slope is degenerate and our target reduces to theirs. Causal Dantzig (Rothenhäusler et al., 2019) also recovers a common linear effect through invariance, while the classical nonparametric literature recovers coefficient distributions from repeated measurements or panel structure (Beran & Hall, 1992; Hoderlein et al., 2010; Arellano & Bonhomme, 2012). We instead exploit variance shifts: the environment enters the outcome equation only through the disturbance scale, so the identifying restrictions change rather than weaken and the target is two functionals of the coefficient law rather than the law itself.

Signed effects, sharpness, and inference. The closest targets are those of Kawakami & Tian (2026): their P-CACE and N-CACE are the probability-weighted positive and negative parts of the individual effect, which for a unit treatment contrast coincide with our targets, and their decomposition of the average effect is the identity $P - N = \mu$ used here. What differs is the identifying information: they work from the conditional outcome distribution under a binary treatment and assume conditional exogeneity, $Y_x \perp X \mid W$, with monotonicity over the potential outcomes upgrading their bounds to point identification, whereas we use cross-environment scale variation with a latent confounder and a random slope, show that every point is attained by a compatible model, and target the whole set (Romano & Shaikh, 2010; Chernozhukov et al., 2007).

Characteristic functions and deconvolution. Characteristic-function products separate a common signal from noise whose scale differs across groups. In the two-level heteroscedastic model of Meister et al. (2010) the relative scaling parameter is known and is unidentified in general otherwise; repeated-measurement and heteroscedastic-error designs supply replicates (Hall & Maiti, 2009; Delaigle & Meister, 2008). We use the same device without assuming the scale ratio known: linked treatment and outcome fourth cumulants recover the positive scales from the environment contrasts, and the Fourier representation of absolute moments (Pinelis, 2011; 2018) turns the reconstructed laws into both signed effects.

7 Limitations and Conclusion

Variance shifts across a few environments identify the first two moments of a confounded causal slope, and those two moments deliver the exact joint set of the slope's positive and negative effects. Under the additional assumptions of Theorem 3, three environments with distinct positive scales suffice for point identification of both effects. There are still some limitations to our conclusions. First, the identifying restrictions are substantive: environments must differ in the scale of the treatment disturbance while preserving the joint distribution of the slope and latent common components. Second, the characteristic-function route used for point identification is exact in population but may be ill-conditioned in finite samples. Third, the sampling theory assumes independent observations. Extending the analysis to dependent data is a natural next step.

AI USE STATEMENT

Generative AI tools, including GPT, were used to assist with drafting and language editing of portions of the manuscript, literature searches and comparisons, development of empirical and simulation code, interpretation of results, and development of parts of the mathematical proofs. The authors reviewed all AI-assisted content, checked the cited literature, and verified the arguments and conclusions of the AI-assisted proofs. The authors also ran and checked the empirical and simulation code and verified the reported numerical results. The authors take full responsibility for the final content of this work, including all mathematical claims, proofs, and computational results.

ETHICS STATEMENT

The study uses publicly available market data and does not involve human subjects. The empirical findings concern historical market data, and their causal interpretation depends on the stated modeling assumptions. They do not constitute investment advice.

REPRODUCIBILITY STATEMENT

Appendix D describes the empirical data construction, including the treatment and state definitions, sample restrictions, and state-weight normalization. Appendix C describes the simulation design and implementation. The companion statistical supplement provides the sampling design and asymptotic derivations summarized in Appendix B.

REFERENCES

Manuel Arellano and Stéphane Bonhomme. Identifying distributional characteristics in random coefficients panel data models. *The Review of Economic Studies*, 79(3):987–1020, 2012. doi: 10.1093/restud/rdr045.

Rudolf Beran and Peter Hall. Estimating coefficient distributions in random coefficient regressions. *The Annals of Statistics*, 20(4):1970–1984, 1992. doi: 10.1214/aos/1176348898.

Victor Chernozhukov, Han Hong, and Elie Tamer. Estimation and confidence regions for parameter sets in econometric models. *Econometrica*, 75(5):1243–1284, 2007. doi: 10.1111/j.1468-0262.2007.00794.x.

Aurore Delaigle and Alexander Meister. Density estimation with heteroscedastic error. *Bernoulli*, 14(2):562–579, 2008. doi: 10.3150/08-BEJ121.

AmirEmad Ghassami, Negar Kiyavash, Biwei Huang, and Kun Zhang. Multi-domain causal structure learning in linear systems. In *Advances in Neural Information Processing Systems*, volume 31, 2018.

Peter Hall and Tapabrata Maiti. Deconvolution methods for non-parametric inference in two-level mixed models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 71 (3):703–718, 2009. doi: 10.1111/j.1467-9868.2009.00705.x.

Lars Peter Hansen. Large sample properties of generalized method of moments estimators. *Econometrica*, 50(4):1029–1054, 1982. doi: 10.2307/1912775.

Joel Hasbrouck. Intraday price formation in U.S. equity index markets. *The Journal of Finance*, 58 (6):2375–2400, 2003. doi: 10.1046/j.1540-6261.2003.00609.x.

James J. Heckman, Sergio Urzua, and Edward Vytlacil. Understanding instrumental variables in models with essential heterogeneity. *The Review of Economics and Statistics*, 88(3):389–432, 2006. doi: 10.1162/rest.88.3.389.

Stefan Hoderlein, Jussi Klemelä, and Enno Mammen. Analyzing the random coefficient model nonparametrically. *Econometric Theory*, 26(3):804–837, 2010. doi: 10.1017/S0266466609990119.

Biwei Huang, Kun Zhang, Jiji Zhang, Joseph Ramsey, Ruben Sanchez-Romero, Clark Glymour, and Bernhard Schölkopf. Causal discovery from heterogeneous/nonstationary data. *Journal of Machine Learning Research*, 21(89):1–53, 2020. URL http://jmlr.org/papers/v21/19-232.html.

Guido W. Imbens and Joshua D. Angrist. Identification and estimation of local average treatment effects. *Econometrica*, 62(2):467–475, 1994. doi: 10.2307/2951620.

Yuta Kawakami and Jin Tian. Measures for assessing causal effect heterogeneity unexplained by covariates. *arXiv preprint arXiv:2602.08647*, 2026. doi: 10.48550/arXiv.2602.08647.

Lutz Kilian and Cheolbeom Park. The impact of oil price shocks on the U.S. stock market. *International Economic Review*, 50(4):1267–1287, 2009. doi: 10.1111/j.1468-2354.2009.00568.x.

Yaroslav Kivva, Sina Akbari, Saber Salehkaleybar, and Negar Kiyavash. Causal effect identification in heterogeneous environments from higher-order moments. In *Proceedings of the Forty-first Conference on Uncertainty in Artificial Intelligence*, volume 286 of *Proceedings of Machine Learning Research*, pp. 2223–2254. PMLR, 2025.

Arthur Lewbel. Using heteroscedasticity to identify and estimate mismeasured and endogenous regressor models. *Journal of Business & Economic Statistics*, 30(1):67–80, 2012. doi: 10.1080/07350015.2012.643126.

Matthew A. Masten. Random coefficients on endogenous variables in simultaneous equations models. *The Review of Economic Studies*, 85(2):1193–1250, 2018. doi: 10.1093/restud/rdx047.

Matthew A. Masten and Alexander Torgovitsky. Identification of instrumental variable correlated random coefficients models. *The Review of Economics and Statistics*, 98(5):1001–1005, 2016. doi: 10.1162/REST_a_00603.

Alexander Meister, Ulrich Stadtmüller, and Christian Wagner. Density deconvolution in a two-level heteroscedastic model with unknown error density. *Electronic Journal of Statistics*, 4:36–57, 2010. doi: 10.1214/09-EJS444.

Whitney K. Newey and Daniel McFadden. Large sample estimation and hypothesis testing. In Robert F. Engle and Daniel L. McFadden (eds.), *Handbook of Econometrics, Volume IV*, volume 4, pp. 2111–2245. Elsevier, 1994. doi: 10.1016/S1573-4412(05)80005-4.

Whitney K. Newey and Kenneth D. West. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55(3):703–708, 1987. doi: 10.2307/1913610.

Ľuboš Pástor and Robert F. Stambaugh. Liquidity risk and expected stock returns. *Journal of Political Economy*, 111(3):642–685, 2003. doi: 10.1086/374184.

Iosif Pinelis. Positive-part moments via the Fourier–Laplace transform. *Journal of Theoretical Probability*, 24(2):409–421, 2011. doi: 10.1007/s10959-010-0276-9.

Iosif Pinelis. Positive-part moments via characteristic functions, and more general expressions. *Journal of Theoretical Probability*, 31(1):527–555, 2018. doi: 10.1007/s10959-016-0709-1.

Roberto Rigobon. Identification through heteroskedasticity. *The Review of Economics and Statistics*, 85(4):777–792, 2003. doi: 10.1162/003465303772815727.

Joseph P. Romano and Azeem M. Shaikh. Inference for the identified set in partially identified econometric models. *Econometrica*, 78(1):169–211, 2010. doi: 10.3982/ECTA6706.

Dominik Rothenhäusler, Peter Bühlmann, and Nicolai Meinshausen. Causal dantzig: Fast inference in linear structural equation models with hidden variables under additive interventions. *The Annals of Statistics*, 47(3):1688–1722, 2019. doi: 10.1214/18-AOS1732.

Shanghai Futures Exchange. Notice on adjustments to order-submission fee schedules, 2024. URL https://www.shfe.com.cn/publicnotice/notice/202405/t20240522_801753.html. May 22. Notice No. 177 of 2024. In Chinese.

Shanghai Futures Exchange. Notice on adjustments to transaction fees for specified tin futures contracts, 2025. URL https://www.shfe.com.cn/publicnotice/notice/202512/t20251212_829775.html. December 12. Notice No. 364 of 2025. In Chinese.

Peter Tufano. The determinants of stock price exposure: Financial engineering and the gold mining industry. *The Journal of Finance*, 53(3):1015–1052, 1998. doi: 10.1111/0022-1082.00042.

A. W. van der Vaart. *Asymptotic Statistics*. Cambridge University Press, 1998. doi: 10.1017/CBO9780511802256.

Kun Zhang, Shaoan Xie, Ignavier Ng, and Yujia Zheng. Causal representation learning from multiple distributions: A general setting. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 60057–60075. PMLR, 2024.

A Proofs of the Identification Results

This appendix collects the proofs. Section 4 states Theorem 1, Theorem 2 and Theorem 3; Proposition A.1, Proposition A.2 and Corollary A.5 are stated below, next to their proofs. The proofs appear in the order in which the results are used, and the appendix closes with the discussion of the assumptions those results require.

A.1 Second moments and the signed effects

The scale family. Write $L = (U, C)^\top$, $V = (\eta, B\eta)^\top$, $M_B = \begin{pmatrix} 1 & \mu \\ \mu & q \end{pmatrix}$, $\Sigma_0 = \mathrm{Cov}_w(L)$ and $Z_s = L + sV$, so that Z_s is the pair (X_e, Y_e) that an environment with scale $s_e = s$ would deliver.

Proposition A.1 (Exact moment conditions for the scale family)**.** *Under Assumption 4, Condition M is equivalent to*

$$\mathbb{E}_w[Z_s] = \mathbb{E}_w[L], \qquad \mathrm{Cov}_w(Z_s) = \Sigma_0 + s^2 M_B \quad \textit{for all } s \geq 0.$$

Proof of Proposition A.1. Write $K_{LV} = \mathrm{Cov}_w(L, V)$.

$$\mathbb{E}_w[Z_s] = \mathbb{E}_w[L] + s\mathbb{E}_w[V],$$

and

$$\mathrm{Cov}_w(Z_s) = \Sigma_0 + s(K_{LV} + K_{LV}^\top) + s^2 \mathrm{Cov}_w(V).$$

The required decomposition holds for every finite $s \geq 0$ if and only if the mean polynomial vanishes identically, the coefficient of s in the covariance polynomial vanishes, and the coefficient of s^2 equals M_B. Comparing coefficients gives $\mathbb{E}_w[V] = 0$, $K_{LV} + K_{LV}^\top = 0$, and $\mathrm{Cov}_w(V) = M_B$, which is Condition M. The converse substitutes those restrictions into the expansion. □

Proof of Theorem 1. By Proposition A.1,

$$\mathrm{Cov}_w(Z_s) = \Sigma_0 + s^2 M_B.$$

Environment e corresponds to $s = s_e$, so

$$\Sigma_e - \Sigma_f = (s_e^2 - s_f^2) M_B.$$

The upper-left, off-diagonal, and lower-right entries give $v_e - v_f = s_e^2 - s_f^2$, $c_e - c_f = (s_e^2 - s_f^2)\mu$, and $r_e - r_f = (s_e^2 - s_f^2)q$. Dividing by $v_e - v_f \neq 0$ yields equation 5. □

Why differencing works. The environment contrast cancels the latent baseline: combining $\Sigma_e = \Sigma_0 + s_e^2 M_B$ with $v_e - v_f = s_e^2 - s_f^2$ gives

$$\Sigma_e - \Sigma_f = (v_e - v_f) M_B = (v_e - v_f) \begin{pmatrix} 1 & \mu \\ \mu & q \end{pmatrix}, \tag{8}$$

so a single contrast in treatment dispersion reveals both the mean slope μ and its second moment q; equivalently, the matrices Σ_e are affine in v_e, which is condition (iii) of Proposition A.2. For the embedding, take the least dispersed environment as the baseline ($\Sigma_0 = \Sigma_{e_*}$ and $s_e^2 = v_e - v_{e_*}$), attach any slope distribution with the identified first two moments, and reconstruct every environment; the construction is the one used in the proof of Theorem A.4.

The remaining results in this subsection carry the identified moments through to the signed effects: the compatibility characterization, the attainability of slope distributions, the joint sharp set, and its aggregation to a target population.

Proposition A.2 (Observable moment compatibility)**.** *For each environment e, let $m_e = \mathbb{E}_{e,w}[(X_e, Y_e)^\top]$ and $\Sigma_e = \mathrm{Cov}_{e,w}((X_e, Y_e))$ be the observed mean and covariance matrix of (X_e, Y_e). There exists a model satisfying Assumptions 1, 3 and 4 together with Condition M that reproduces $\{(m_e, \Sigma_e)\}_e$ if and only if*

1. *there is a single mean vector, independent of e, with $m_e = m$ for every e;*
2. *each Σ_e is positive semidefinite;*
3. *there are constants $\alpha_c, \alpha_y, \mu, q$, independent of e, with*

$$c_e = \alpha_c + \mu v_e, \qquad r_e = \alpha_y + q v_e, \qquad q \geq \mu^2.$$

Proof of Proposition A.2. *Necessity.* Under the model, $\mathbb{E}_w[X_e]$ and $\mathbb{E}_w[Y_e]$ are environment-invariant, giving common means. Covariance matrices are positive semidefinite by definition. With $u_0 = \mathrm{Var}_w(U)$, $s_e^2 = v_e - u_0$, and the covariance decomposition gives

$$c_e = \mathrm{Cov}_w(U, C) + \mu(v_e - u_0), \qquad r_e = \mathrm{Var}_w(C) + q(v_e - u_0).$$

Setting $\alpha_c = \mathrm{Cov}_w(U, C) - \mu u_0$ and $\alpha_y = \mathrm{Var}_w(C) - q u_0$ gives the affine relations, and $q \geq \mu^2$ follows from $\mathrm{Var}_w(B) \geq 0$.

Sufficiency. Let $v_* = \min_e v_e$ and choose an environment attaining it. The affine relations imply that all minimizers have the same covariance matrix Σ_*. Construct (U, C) with mean m and covariance Σ_* (for instance using a positive semidefinite square root and independent Rademacher coordinates). Choose any distribution of B with mean μ and second moment q, independent of (U, C), and an independent standardized Rademacher η. Set $A = C - BU$ and $s_e = \sqrt{v_e - v_*}$. Then $X_e = U + s_e \eta$ and $Y_e = C + s_e B \eta$ have the prescribed means and covariance matrices.

Uniqueness and the degenerate branch. If $v_e \neq v_f$, subtracting the affine relations for e and f gives $c_e - c_f = \mu(v_e - v_f)$ and $r_e - r_f = q(v_e - v_f)$, so μ and q are determined; no other pair can reproduce the same moments. If all v_e coincide, the affine relations force $c_e \equiv c$ and $r_e \equiv r$, and for any μ, q with $q \geq \mu^2$ the choices $\alpha_c = c - \mu v$ and $\alpha_y = r - q v$ satisfy condition (iii); the construction above then generates the moments, so every such pair is compatible. □

The sufficiency construction satisfies Assumption 2 as well, so the same characterization applies to the independent-disturbance baseline of Section 3. The constants α_c and α_y are affine intercepts with respect to v_e, not standalone baseline covariance entries. When these conditions hold, a nonzero variance contrast uniquely determines the compatible μ and q, which are the identified slope moments. If all v_e coincide, the covariance matrices also coincide, and every pair (μ, q) with $q \geq \mu^2$ is compatible.

Lemma A.3 (compatibility including zero variance shifts)**.** *For a candidate raw-moment vector z and parameter θ, write $\mathcal{A}(z, \theta)$ for the following conditions at every state: the environmental means agree; each covariance matrix is positive semidefinite; and there exist α_c, α_y with*

$$c_e = \alpha_c + \mu v_e, \qquad r_e = \alpha_y + q v_e, \qquad q \geq \mu^2,$$

with no requirement that the v_e differ. Then (z, θ) is generated by the specified independent-shock SCM if and only if $\mathcal{A}(z, \theta)$ holds. A sign-magnitude triple (P, N, t) is attainable with these moments if and only if

$$t \geq |\mu|, \qquad t^2 \leq q, \qquad [t > 0 \text{ or } q = 0].$$

Proof of Lemma A.3. Necessity of the compatibility conditions follows from the covariance decomposition, including when two scales are equal. For sufficiency, choose $v_{*h} = \min_e v_{eh}$ and an

environment attaining it. The affine relations force all minimizers to share a covariance matrix Σ_{*h}. Set $\tau_{eh} = v_{eh} - v_{*h} \geq 0$ and choose a bounded (U_h, C_h) with mean (a_h, b_h) and covariance Σ_{*h}, using a positive semidefinite square root times independent Rademacher coordinates. Choose B_h independently with the required two moments and an independent standardized Rademacher η_h, and set $A_h = C_h - B_h U_h$ and $s_{eh} = \sqrt{\tau_{eh}}$. This constructs the SCM and preserves the candidate moments; if all variances are equal, all covariances are forced equal and every scale can be set to zero.

For the sign–magnitude triple, necessity follows from the triangle and Cauchy–Schwarz inequalities and from $\mathbb{E}|B_h| = 0$ implying $B_h = 0$ almost surely. If $q_h = 0$, feasibility forces $\mu_h = t_h = 0$, so take $B_h = 0$. Otherwise let $\rho_h = q_h/t_h > 0$ and place masses

$$\frac{t_h(t_h - \mu_h)}{2q_h}, \qquad 1 - \frac{t_h^2}{q_h}, \qquad \frac{t_h(t_h + \mu_h)}{2q_h}$$

at $-\rho_h, 0, \rho_h$, respectively. The masses are nonnegative, sum to one, and reproduce (μ_h, q_h, t_h); the preceding embedding preserves all environmental moments. □

Compatibility constrains only the first two moments of B.

Theorem A.4 (Attainable slope distributions)**.** *Fix w and suppose that the compatibility conditions of Proposition A.2 hold together with a nonzero variance contrast, so that μ and q are the identified slope moments. Consider every model that satisfies the assumptions of that theorem and reproduces the observed conditional moments $\{(m_e, \Sigma_e)\}_e$. The conditional law of B given $W = w$ under such a model ranges over exactly the class of distributions on $\mathbb{R}$ with mean μ and second moment q. This class is nonempty exactly when $q \geq \mu^2$, and the construction in the proof of Proposition A.2 attains every one of its members.*

Proof of Theorem A.4. Proposition A.2 shows that every compatible model must have $\mathbb{E}_w[B] = \mu$ and $\mathbb{E}_w[B^2] = q$. Conversely, the sufficiency construction can choose any distribution of B with those two moments while keeping (U, C), the means, and the covariance matrices unchanged. Hence the attainable marginal distributions of B are exactly those with moments (μ, q). □

By Theorem A.4, every distribution of B with mean μ and second moment q is attainable while preserving the observed environmental moments. Determining the joint sharp set for (P, N) therefore reduces to finding the attainable values of $t = \mathbb{E}_w|B|$ under these two moment constraints.

Proof of Theorem 2. *Necessity.* Since $\mu = P - N$ and $t = P + N$, every attainable pair has the form

$$(P, N) = \left(\frac{t + \mu}{2}, \frac{t - \mu}{2}\right).$$

The triangle and Cauchy–Schwarz inequalities give

$$|\mu| \leq \mathbb{E}_w|B| \leq \sqrt{q}.$$

If $q = 0$, then $B = 0$ almost surely, so $\mu = t = 0$. If $q > 0$, then $t > 0$, since $\mathbb{E}_w|B| = 0$ would force $B = 0$ almost surely and hence $q = 0$. These observations yield the three cases in equation 6.

Attainability. Fix $t \in \mathcal{T}(\mu, q)$. For $q = 0$, the distribution concentrated at $B = 0$ attains $t = 0$. For $q > 0$, we have $t > 0$. Set $d = q/t$ and define

$$\Pr_w(B = d) = \frac{t(t + \mu)}{2q}, \qquad \Pr_w(B = -d) = \frac{t(t - \mu)}{2q}, \qquad \Pr_w(B = 0) = 1 - \frac{t^2}{q}.$$

The inequalities $|\mu| \leq t \leq \sqrt{q}$ ensure that these probabilities are nonnegative and sum to one. Direct calculation gives

$$\mathbb{E}_w[B] = \mu, \qquad \mathbb{E}_w[B^2] = q, \qquad \mathbb{E}_w|B| = t.$$

By Theorem A.4, this distribution can be embedded in a structural model preserving all observed environmental means and covariance matrices. Thus every stated value of t, and hence every pair in $\mathcal{I}_{PN}(w)$, is attainable.

The construction attains $t = \sqrt{q}$ in every case and $t = |\mu|$ whenever $\mu \neq 0$; the case $q = 0$ is immediate. When $\mu = 0 < q$, it attains every $t \in (0, \sqrt{q}]$, so zero is the infimum but is not attainable. Finally, $q = \mu^2$ implies $\mathrm{Var}_w(B) = 0$, giving the constant-slope conclusion. □

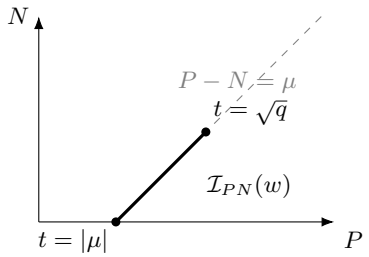


Figure 2: The sharp set of the signed effects, drawn for $\mu = 0.6$ and $\sqrt{q} = 2$. Both axes are nonnegative, since $P = \mathbb{E}_w[B_+]$ and $N = \mathbb{E}_w[(-B)_+]$. Because $P - N = \mu$ and $P + N = t$, every attainable pair lies on the dashed line $N = P - \mu$, and the attainable pairs are exactly the points of the thick segment $\mathcal{I}_{PN}(w)$, traced as t runs from $|\mu|$ to $\sqrt{q}$. One contribution is zero at $t = |\mu|$, and the amplitude $|B|$ is constant at $t = \sqrt{q}$. Points of the dashed line outside the segment are not attainable.

From conditional effects to a population target. The potential-outcome map $Y(x) = A + Bx$ is common to the environments compared at a fixed state (Assumption 1), but its law may differ across states. Environments are therefore the source of identifying variation, whereas states are the dimension over which we aggregate: every result so far is conditional on one w, and different states carry different slope moments and hence different sharp sets. A single population-level statement requires a target distribution over states. Corollary A.5 integrates the state-specific sets over such a target G: it delivers the sharp set for the population functionals $P_G = \int P\, dG$ and $N_G = \int N\, dG$. The weights describe the population being reported and are chosen by the analyst rather than estimated from the environments; they do not alter the state-specific mechanism.

Corollary A.5 (Sharp aggregation in a target population)**.** *Suppose that at every state represented in G the compatibility conditions of Proposition A.2 hold and two comparable observed environments have distinct treatment variances, that no further restrictions link the conditional latent distributions across states, and that $\int q(w)\, dG(w) < \infty$. Write*

$$\mu_G = \int \mu(w)\, dG(w), \qquad L_G = \int |\mu(w)|\, dG(w), \qquad H_G = \int \sqrt{q(w)}\, dG(w),$$

and let π_0 be the target-population probability of states with $\mu(w) = 0$ and $q(w) > 0$. The attainable values of $t_G = \int t(w)\, dG(w)$ are exactly

$$\mathcal{T}_G = \begin{cases} [L_G, H_G], & \pi_0 = 0, \\ (L_G, H_G], & \pi_0 > 0, \end{cases} \tag{9}$$

and the joint set for (P_G, N_G) is the segment $P - N = \mu_G$ traced by $t \in \mathcal{T}_G$.

Proof of Corollary A.5. Write

$$\mathcal{W}_0 = \{w : \mu(w) = 0,\ q(w) > 0\}, \qquad G(\mathcal{W}_0) = \pi_0.$$

Necessity. For any compatible model, Theorem 2 gives

$$|\mu(w)| \le t(w) \le \sqrt{q(w)}.$$

Integrating over the target distribution yields

$$L_G \le t_G \le H_G, \qquad H_G^2 \le \int q(w)\, dG(w) < \infty.$$

On $\mathcal{W}_0$, the same theorem requires $t(w) > 0 = |\mu(w)|$. Consequently, if $\pi_0 > 0$, then

$$t_G - L_G = \int \{t(w) - |\mu(w)|\}\, dG(w) > 0,$$

so the lower endpoint is excluded.

Attainability. First suppose that $H_G = L_G$. Since $\sqrt{q(w)} - |\mu(w)| \geq 0$, equality of the integrals implies $q(w) = \mu(w)^2$ for G-almost every w. Thus $\pi_0 = 0$, and choosing the attainable upper endpoint $t(w) = \sqrt{q(w)}$ at each state gives $t_G = L_G = H_G$.

Now suppose that $H_G > L_G$. Fix

$$t^* \in \begin{cases} [L_G, H_G], & \pi_0 = 0, \\ (L_G, H_G], & \pi_0 > 0, \end{cases}$$

and define

$$\lambda = \frac{t^* - L_G}{H_G - L_G}, \qquad t_\lambda(w) = |\mu(w)| + \lambda\big\{\sqrt{q(w)} - |\mu(w)|\big\}.$$

Here $\lambda \in [0, 1]$, and $\lambda = 0$ is used only when $\pi_0 = 0$. If $\lambda = 0$, redefine $t_\lambda(w) = \sqrt{q(w)}$ on $\mathcal{W}_0$. Because this set has zero target probability in that case, the modification does not change the integral.

The resulting choice satisfies

$$t_\lambda(w) \in \mathcal{T}\big(\mu(w), q(w)\big)$$

at each relevant state, including $t_\lambda(w) = 0$ whenever $q(w) = 0$, and

$$\int t_\lambda(w)\, dG(w) = L_G + \lambda(H_G - L_G) = t^*.$$

The explicit construction in Theorem 2 provides a conditional slope distribution attaining each $t_\lambda(w)$. Its support points and probabilities, together with the embedding construction of Theorem A.4, depend measurably on the conditional moments. Since no additional restrictions link latent distributions across states, these constructions define a single compatible model preserving the state-specific environmental moments.

Hence every value in the stated interval is attainable. When $\pi_0 > 0$, taking $\lambda \downarrow 0$ approaches L_G without attaining it; when $\pi_0 = 0$, the choice $\lambda = 0$ attains L_G. Finally, integrating the identities $P(w) = (t(w) + \mu(w))/2$ and $N(w) = (t(w) - \mu(w))/2$ gives

$$(P_G, N_G) = \left(\frac{t_G + \mu_G}{2}, \frac{t_G - \mu_G}{2}\right),$$

establishing the claimed joint sharp set. □

A two-state example. Suppose the target puts weight $1/2$ on each of two states, with $(\mu, q) = (1, 4)$ in the first and $(-1, 4)$ in the second. Then $\mu_G = 0$, $L_G = 1$, and $H_G = 2$, so $t_G \in [1, 2]$ and

$$(P_G, N_G) = \left(\frac{t_G}{2}, \frac{t_G}{2}\right),$$

a segment from $(1/2, 1/2)$ to $(1, 1)$. Three features are worth noting. First, a zero population mean does not imply negligible signed contributions: despite $\mu_G = 0$, the conditional moment information implies $P_G, N_G \geq 1/2$. Second, the lower endpoint uses $\int |\mu|\, dG$, not $|\mu_G|$, which is why $L_G = 1$ rather than 0. Third, a vanishing population mean does not by itself make the endpoint open; that happens only when the target places positive weight on states with $\mu = 0 < q$, as the case split in equation 9 records.

A.2 Point identification of the signed effects

Theorem 3, with the explicit formulas. Fix a state w at which the environments used have positive conditional probability and observed support. Suppose that Assumptions 1 and 3 hold, that every environment used has $0 < s_e < \infty$, that at least three environments have pairwise distinct values of s_e^2, and that:

(a) for every environment used, $\eta \perp (A, B, U) \mid (E = e, W = w)$ and the conditional law $\mathcal{L}(\eta \mid E = e, W = w)$ is common across environments, symmetric about zero and of second moment one;

(b) $C \perp |B|$ at w;

(c) $\mathbb{E}_w[|U|^4 + |C|^4 + |B|^4 + |\eta|^4] < \infty$.

For $Z \in \{X, Y\}$ define

$$\kappa_4^{(Z)}(e) = \mathbb{E}_w\big[(Z_e - \mathbb{E}_w Z_e)^4\big] - 3\{\mathrm{Var}_w(Z_e)\}^2, \qquad \varphi_{Z,e}(u) = \mathbb{E}_w[e^{iuZ_e}],$$

and let $p_Z(v) = a_Z v^2 + b_Z v + d_Z$ be the unique polynomial of degree at most two interpolating $\big(v_e, \kappa_4^{(Z)}(e)\big)$ at any three environments with pairwise distinct treatment variances. Under the assumptions this polynomial does not depend on the selected triple and fits every environment. Then μ and q are identified by equation 5, whose hypotheses are implied by (a) and (c), the amplitude dispersion

$$\xi := \frac{a_Y - q^2 a_X}{a_X + 3} = \mathrm{Var}_w(B^2) \geq 0 \tag{10}$$

is observable with $a_X + 3 \geq 1$, and the signed effects are point identified in either of two cases.

(i) *Constant amplitude.* If $\xi = 0$, then $B^2 = q$ almost surely and $t = \sqrt{q}$. No characteristic-function condition is needed; this branch covers $q = 0$ and permits variation in the sign of B when $q > 0$.

(ii) *Varying amplitude.* Suppose that $\xi > 0$ and that $\varphi_{X,e}(u) \neq 0$ and $\varphi_{Y,e}(u) \neq 0$ for every real u at some environment used e, equivalently at every environment used. Fix $Z \in \{X, Y\}$ with $a_Z \neq 0$ and environments ℓ, h with $v_\ell < v_h$. Then the latent variance and every scale square are identified by

$$u_0 = -\frac{b_Z}{2a_Z} = \mathrm{Var}_w(U), \qquad s_e^2 = v_e - u_0 \qquad \text{for every environment used } e,$$

and t is computed from the observed characteristic functions as

$$\varrho = \sqrt{\frac{v_\ell - u_0}{v_h - u_0}} = \frac{s_\ell}{s_h} \in (0, 1), \qquad \Psi_Z(u) = \lim_{n\to\infty} \prod_{j=0}^{n} \frac{\varphi_{Z,h}(\varrho^j u)}{\varphi_{Z,\ell}(\varrho^j u)},$$

$$t = \frac{J_Y}{J_X}, \qquad J_Z = \int_0^\infty \frac{1 - \mathrm{Re}\,\Psi_Z(u)}{u^2}\, du, \qquad Z \in \{X, Y\}.$$

The products converge at every real u, the integrals are finite, and $J_X > 0$, so the ratio is well defined.

With μ from equation 5, either branch gives $P = (t + \mu)/2$ and $N = (t - \mu)/2$.

Proof of Theorem 3. Fix w. The proof combines three sources of identifying information. Second-order moments identify the mean slope $\mu = \mathbb{E}_w[B]$ and its second moment $q = \mathbb{E}_w[B^2]$. Fourth-order moments identify $\xi = \mathrm{Var}_w(B^2)$, distinguishing constant from varying amplitude. When $\xi = 0$, the absolute first moment follows directly as $t = \mathbb{E}_w[|B|] = \sqrt{q}$; when $\xi > 0$, the fourth-order information instead identifies the unknown environmental scales. In this varying-amplitude branch, characteristic functions recover the distributions of two scaled disturbances, whose absolute first-moment ratio identifies the remaining quantity t. Together, μ and t determine the signed effects through $P = (t + \mu)/2$ and $N = (t - \mu)/2$.

We begin by establishing $C \perp B\eta$, which provides the independent-sum structure used in the cumulant and characteristic-function arguments, and conclude by verifying uniqueness within the stated model class.

1. Separability of the baseline and the transmitted disturbance. We first show that $C \perp B\eta$ at the fixed state w. Let $R = |B|$. For any bounded Borel function $g : \mathbb{R} \to \mathbb{R}$, define

$$h_g(b) = \mathbb{E}_w[g(b\eta)].$$

Since $C = A + BU$, independence of η from (A, B, U) implies independence from (B, C). Hence

$$\mathbb{E}_w[g(B\eta) \mid C, B] = h_g(B).$$

Moreover, symmetry of η gives

$$h_g(-b) = \mathbb{E}_w[g(-b\eta)] = \mathbb{E}_w[g(b\eta)] = h_g(b).$$

Thus h_g is an even Borel function, and $h_g(B) = h_g(|B|) = h_g(R)$.

Using the tower property and the separability assumption $C \perp R$, we obtain

$$\begin{aligned}\mathbb{E}_w[g(B\eta) \mid C] &= \mathbb{E}_w[\mathbb{E}_w[g(B\eta) \mid C, B] \mid C] \\ &= \mathbb{E}_w[h_g(B) \mid C] \\ &= \mathbb{E}_w[h_g(R) \mid C] \\ &= \mathbb{E}_w[h_g(R)] \\ &= \mathbb{E}_w[g(B\eta)].\end{aligned}$$

Here the fourth equality uses $C \perp R$, and the last equality follows from $h_g(R) = h_g(B)$ and the tower property. Therefore $C \perp B\eta$ at w.

2. The first two slope moments. Conditions (a) and (c) place Theorem 1 in force. Writing $L = (U, C)^\top$ and $V = (\eta, B\eta)^\top$ as in Section 4, the mean and covariance clauses of Condition M follow from independence, standardization and symmetry, and the vanishing of the symmetric cross-covariance follows from

$$\mathrm{Cov}_w(U, \eta) = \mathrm{Cov}_w(U, B\eta) = \mathrm{Cov}_w(C, \eta) = \mathrm{Cov}_w(C, B\eta) = 0,$$

while the second moments are finite by (c). Hence equation 5 identifies μ and q from any pair of environments with distinct treatment variances, exactly as in the proof of that theorem, and the remaining steps take both as known. The fourth moments enter only from step 3 onward.

3. Fourth cumulants and the observable branch criterion. All moments and cumulants below are taken under the conditional law at the fixed state w. Recall that

$$\kappa_4(T) = \mathbb{E}_w\big[(T - \mathbb{E}_w[T])^4\big] - 3\{\mathrm{Var}_w(T)\}^2.$$

For independent random variables T_1, T_2 with finite fourth moments, fourth cumulants satisfy

$$\kappa_4(T_1 + T_2) = \kappa_4(T_1) + \kappa_4(T_2), \qquad \kappa_4(cT_1) = c^4\kappa_4(T_1).$$

Indeed, after centering T_1 and T_2, the cross term in the fourth moment is $6\mathrm{Var}_w(T_1)\mathrm{Var}_w(T_2)$, which is exactly canceled by the corresponding term in $3\{\mathrm{Var}_w(T_1 + T_2)\}^2$.

Since $U \perp \eta$ and step 1 establishes $C \perp B\eta$, applying these properties to

$$X_e = U + s_e\eta, \qquad Y_e = C + s_eB\eta$$

gives

$$\begin{aligned}\kappa_4(X_e) &= \kappa_4(U) + s_e^4\kappa_4(\eta), \\ \kappa_4(Y_e) &= \kappa_4(C) + s_e^4\kappa_4(B\eta).\end{aligned}$$

To evaluate the disturbance cumulants, write

$$m_{\eta,4} = \mathbb{E}_w[\eta^4], \qquad m_{B,4} = \mathbb{E}_w[B^4].$$

Symmetry and standardization imply $\mathbb{E}_w[\eta] = 0$ and $\mathbb{E}_w[\eta^2] = 1$, so

$$\kappa_4(\eta) = m_{\eta,4} - 3.$$

Moreover, independence of B and η gives

$$\begin{aligned}\mathbb{E}_w[B\eta] &= \mathbb{E}_w[B]\mathbb{E}_w[\eta] = 0, \\ \mathrm{Var}_w(B\eta) &= \mathbb{E}_w[B^2]\mathbb{E}_w[\eta^2] = q, \\ \mathbb{E}_w[(B\eta)^4] &= \mathbb{E}_w[B^4]\mathbb{E}_w[\eta^4] = m_{B,4}m_{\eta,4}.\end{aligned}$$

Consequently,

$$\kappa_4(B\eta) = m_{B,4}m_{\eta,4} - 3q^2.$$

Define

$$a_X = m_{\eta,4} - 3, \qquad a_Y = m_{B,4}m_{\eta,4} - 3q^2.$$

The cumulant relations therefore become

$$\kappa_4(X_e) = \kappa_4(U) + a_X s_e^4, \qquad \kappa_4(Y_e) = \kappa_4(C) + a_Y s_e^4.$$

Let $u_0 = \mathrm{Var}_w(U)$. The treatment-variance relation

$$v_e = \mathrm{Var}_w(X_e) = u_0 + s_e^2$$

implies $s_e^4 = (v_e - u_0)^2$. Hence the observed cumulants satisfy $\kappa_4(Z_e) = p_Z(v_e)$ for $Z \in \{X, Y\}$, where

$$\begin{aligned} p_X(v) &= \kappa_4(U) + a_X(v - u_0)^2, \\ p_Y(v) &= \kappa_4(C) + a_Y(v - u_0)^2. \end{aligned} \tag{11}$$

These are polynomials of degree at most two. In particular, expanding the squares shows that their coefficients of v^2 are a_X and a_Y, regardless of the unknown value of u_0.

The three distinct scale squares give three distinct treatment variances. Thus, for each $Z \in \{X, Y\}$, the three observed pairs $(v_e, \kappa_4(Z_e))$ uniquely determine p_Z and its coefficients. In particular, a_X and a_Y are identified without first knowing u_0 or the scales. Every additional environment must satisfy the same polynomial relations.

Finally, Jensen's inequality gives

$$a_X + 3 = m_{\eta,4} = \mathbb{E}_w[(\eta^2)^2] \geq \{\mathbb{E}_w[\eta^2]\}^2 = 1.$$

The denominator below is therefore strictly positive, and the identified coefficients yield

$$m_{B,4} = \frac{a_Y + 3q^2}{a_X + 3}.$$

Since $q = \mathbb{E}_w[B^2]$ is already identified, it follows that

$$\begin{aligned} \xi &:= \frac{a_Y - q^2 a_X}{a_X + 3} \\ &= \frac{a_Y + 3q^2}{a_X + 3} - q^2 \\ &= \mathbb{E}_w[B^4] - \{\mathbb{E}_w[B^2]\}^2 \\ &= \mathrm{Var}_w(B^2) \geq 0. \end{aligned}$$

This proves equation 10 and makes the distinction between constant and varying amplitude identifiable from the observed environmental moments.

4. Constant amplitude, or recovery of the unknown scales. We now consider the two possible values of the identified quantity $\xi = \mathrm{Var}_w(B^2)$.

If $\xi = 0$, then

$$B^2 = \mathbb{E}_w[B^2] = q \qquad \text{almost surely.}$$

Therefore $|B| = \sqrt{q}$ almost surely, and

$$t = \mathbb{E}_w[|B|] = \sqrt{q}.$$

Using $\mu = P - N$ and $t = P + N$, we obtain

$$P = \frac{\sqrt{q} + \mu}{2}, \qquad N = \frac{\sqrt{q} - \mu}{2}.$$

Thus the signed effects are identified directly, without recovering the scales or imposing a characteristic-function condition. This case permits variation in the sign of B. When $q = 0$, it reduces to $B = 0$ almost surely and $P = N = 0$.

Suppose instead that $\xi > 0$. By equation 10,

$$a_Y - q^2 a_X = (a_X + 3)\xi > 0.$$

Hence a_X and a_Y cannot both be zero. Write the identified polynomials as

$$p_Z(v) = a_Z v^2 + b_Z v + d_Z, \qquad Z \in \{X, Y\}.$$

Expanding equation 11 gives

$$p_X(v) = a_X v^2 - 2a_X u_0 v + \kappa_4(U) + a_X u_0^2,$$
$$p_Y(v) = a_Y v^2 - 2a_Y u_0 v + \kappa_4(C) + a_Y u_0^2.$$

Matching the coefficients of v yields

$$b_Z = -2a_Z u_0, \qquad Z \in \{X, Y\}.$$

Consequently, for any Z with $a_Z \neq 0$,

$$u_0 = -\frac{b_Z}{2a_Z}.$$

Both coefficients on the right are identified, so this ratio identifies $u_0 = \mathrm{Var}_w(U)$. If both a_X and a_Y are nonzero, the model requires the two ratios to coincide:

$$-\frac{b_X}{2a_X} = -\frac{b_Y}{2a_Y} = u_0.$$

It follows that every scale square is identified through

$$s_e^2 = v_e - u_0.$$

Because u_0 is a variance and all scales are strictly positive, the model implies

$$0 \leq u_0 < \min_e v_e, \qquad s_e^2 > 0.$$

The positive scales themselves are therefore identified as $s_e = \sqrt{v_e - u_0}$.

Finally, choose environments ℓ, h with $v_\ell < v_h$. Then $0 < s_\ell < s_h$, and their scale ratio is the identified quantity

$$\varrho = \sqrt{\frac{v_\ell - u_0}{v_h - u_0}} = \sqrt{\frac{s_\ell^2}{s_h^2}} = \frac{s_\ell}{s_h} \in (0, 1).$$

This supplies the scale ratio used in the subsequent characteristic-function reconstruction.

5. Recovering the scaled-disturbance distributions. Continue with the varying-amplitude branch $\xi > 0$. Step 4 identifies the positive scales and the ratio $\varrho = s_\ell / s_h \in (0, 1)$. Our next objective is to recover the distributions of

$$T_X = s_h \eta, \qquad T_Y = s_h B \eta.$$

For $Z \in \{X, Y\}$, denote the corresponding, initially unknown, characteristic function by

$$F_Z(u) = \mathbb{E}_w[e^{iuT_Z}].$$

Also write

$$\varphi_U(u) = \mathbb{E}_w[e^{iuU}], \qquad \varphi_C(u) = \mathbb{E}_w[e^{iuC}].$$

Since $s_\ell = \varrho s_h$, the structural representations in the two selected environments are

$$X_h = U + T_X, \qquad X_\ell = U + \varrho T_X,$$
$$Y_h = C + T_Y, \qquad Y_\ell = C + \varrho T_Y.$$

Independence of U and η gives $U \perp T_X$, while step 1 gives $C \perp T_Y$. Consequently,

$$\varphi_{X,h}(u) = \varphi_U(u) F_X(u), \qquad \varphi_{X,\ell}(u) = \varphi_U(u) F_X(\varrho u),$$
$$\varphi_{Y,h}(u) = \varphi_C(u) F_Y(u), \qquad \varphi_{Y,\ell}(u) = \varphi_C(u) F_Y(\varrho u).$$

The nonvanishing assumption ensures that $\varphi_{X,\ell}(u)$ and $\varphi_{Y,\ell}(u)$ are nonzero for every real u. The factorizations therefore imply that $\varphi_U(u)$, $\varphi_C(u)$, $F_X(\varrho u)$, and $F_Y(\varrho u)$ are all nonzero. Since $\varrho > 0$, the argument ϱu ranges over $\mathbb{R}$, so F_X and F_Y are themselves nowhere zero. Thus all cancellations and divisions below are well defined.

Taking ratios across environments removes the unknown baseline characteristic functions. Specifically, define the observable ratios

$$Q_Z(u) := \frac{\varphi_{Z,h}(u)}{\varphi_{Z,\ell}(u)}, \qquad Z \in \{X, Y\}.$$

The factorizations give

$$Q_Z(u) = \frac{F_Z(u)}{F_Z(\varrho u)},$$

or equivalently,

$$F_Z(u) = Q_Z(u) F_Z(\varrho u).$$

Thus the unknown characteristic function at frequency u is related to its value at the smaller frequency ϱu through an identified multiplier $Q_Z(u)$.

Applying the same relation at ϱu gives

$$F_Z(u) = Q_Z(u) Q_Z(\varrho u) F_Z(\varrho^2 u).$$

Iterating, for every integer $n \geq 0$,

$$F_Z(u) = \left\{ \prod_{j=0}^{n} Q_Z(\varrho^j u) \right\} F_Z(\varrho^{n+1} u).$$

Equivalently, the finite products telescope:

$$\begin{aligned} \prod_{j=0}^{n} \frac{\varphi_{Z,h}(\varrho^j u)}{\varphi_{Z,\ell}(\varrho^j u)} &= \prod_{j=0}^{n} \frac{F_Z(\varrho^j u)}{F_Z(\varrho^{j+1} u)} \\ &= \frac{F_Z(u)}{F_Z(\varrho^{n+1} u)}. \end{aligned}$$

Because $0 < \varrho < 1$, we have $\varrho^{n+1} u \to 0$ for every fixed real u. Characteristic functions are continuous at zero and equal one there, so

$$F_Z(\varrho^{n+1} u) \longrightarrow F_Z(0) = 1.$$

It follows that the observable reconstruction

$$\Psi_Z(u) := \lim_{n \to \infty} \prod_{j=0}^{n} \frac{\varphi_{Z,h}(\varrho^j u)}{\varphi_{Z,\ell}(\varrho^j u)}$$

exists for every real u and satisfies

$$\Psi_X(u) = F_X(u) = \mathbb{E}_w[e^{i u s_h \eta}], \qquad \Psi_Y(u) = F_Y(u) = \mathbb{E}_w[e^{i u s_h B \eta}].$$

The argument uses finite-product identities and continuity at zero; no complex logarithm or branch choice is required.

Since characteristic functions uniquely determine probability distributions, this identifies the marginal distributions of $T_X = s_h \eta$ and $T_Y = s_h B \eta$. The next step uses these identified distributions to determine their absolute first moments and, through their ratio, $\mathbb{E}_w[|B|]$.

6. From the recovered distributions to the signed functionals. Once the marginal distributions of T_X and T_Y are identified, their finite absolute first moments are identified as well. Write

$$m_X = \mathbb{E}_w[|T_X|], \qquad m_Y = \mathbb{E}_w[|T_Y|].$$

These moments are finite: independence of B and η and standardization of η give

$$\mathbb{E}_w[T_X^2] = s_h^2, \qquad \mathbb{E}_w[T_Y^2] = s_h^2 q,$$

so Cauchy–Schwarz implies

$$m_X \leq s_h, \qquad m_Y \leq s_h \sqrt{q} < \infty.$$

In the latent representation, $T_Y = BT_X$. Moreover, $B \perp \eta$ implies $B \perp T_X$, because s_h is a fixed constant. Hence

$$\begin{aligned} m_Y &= \mathbb{E}_w[|BT_X|] \\ &= \mathbb{E}_w[|B||T_X|] \\ &= \mathbb{E}_w[|B|]\,\mathbb{E}_w[|T_X|] \\ &= t\,m_X. \end{aligned}$$

Since $\mathbb{E}_w[\eta^2] = 1$, the disturbance is not almost surely zero. Together with $s_h > 0$, this gives

$$m_X = s_h \mathbb{E}_w[|\eta|] > 0.$$

Therefore,

$$t = \mathbb{E}_w[|B|] = \frac{m_Y}{m_X}.$$

This relation uses the structural independence assumption but requires only the two identified marginal absolute moments, not recovery of the joint distribution of (T_X, T_Y).

To express these moments directly in terms of the recovered characteristic functions, recall that, for every real x,

$$|x| = \frac{2}{\pi} \int_0^\infty \frac{1 - \cos(ux)}{u^2}\, du.$$

For an integrable real random variable T, the integrand is nonnegative, so Tonelli's theorem gives

$$\begin{aligned} \mathbb{E}_w[|T|] &= \frac{2}{\pi} \int_0^\infty \frac{\mathbb{E}_w[1 - \cos(uT)]}{u^2}\, du \\ &= \frac{2}{\pi} \int_0^\infty \frac{1 - \operatorname{Re} \varphi_T(u)}{u^2}\, du. \end{aligned}$$

Related positive-part representations are discussed by Pinelis (2011; 2018). This identity computes the required absolute moments directly from characteristic functions, without first reconstructing a probability density.

Define

$$J_Z := \int_0^\infty \frac{1 - \operatorname{Re} \Psi_Z(u)}{u^2}\, du, \qquad Z \in \{X, Y\}.$$

Applying the identity to the two recovered distributions yields

$$m_X = \frac{2}{\pi} J_X, \qquad m_Y = \frac{2}{\pi} J_Y.$$

Equivalently,

$$J_X = \frac{\pi}{2} s_h \mathbb{E}_w[|\eta|], \qquad J_Y = \frac{\pi}{2} s_h \mathbb{E}_w[|B|]\,\mathbb{E}_w[|\eta|].$$

Both integrals are finite, and $J_X > 0$. Their ratio therefore gives

$$t = \frac{m_Y}{m_X} = \frac{J_Y}{J_X}.$$

Finally, the identities

$$B_+ = \frac{|B| + B}{2}, \qquad (-B)_+ = \frac{|B| - B}{2}$$

and the mean $\mu = \mathbb{E}_w[B]$ identified in step 2 imply

$$P = \frac{t + \mu}{2} = \frac{J_Y/J_X + \mu}{2}, \qquad N = \frac{t - \mu}{2} = \frac{J_Y/J_X - \mu}{2}.$$

Every quantity on the right is determined by the observed environmental distributions. Thus the signed effects are point identified without requiring reconstruction of the full distribution of B.

7. Uniqueness. Fix w and consider any two models satisfying the assumptions of the theorem and inducing the same conditional distribution

$$\mathcal{L}(X, Y \mid E = e, W = w)$$

for every environment e.

The two models have the same observed means, covariances, fourth cumulants, and marginal characteristic functions. The identification formulas therefore give the same μ, q, and ξ, placing both models in the same branch of the theorem.

If $\xi = 0$, both models satisfy $t = \sqrt{q}$. If $\xi > 0$, the recovered baseline variance u_0 and all scale squares are the same in both models. Using the same ordered pair of environments, they therefore have the same scale ratio ϱ. The reconstruction in step 5 then gives the same Ψ_X and Ψ_Y, and step 6 gives the same value

$$t = \frac{J_Y}{J_X}.$$

Thus both models have the same μ and t, and consequently the same signed effects:

$$P^{(1)} = P^{(2)} = \frac{t+\mu}{2}, \qquad N^{(1)} = N^{(2)} = \frac{t-\mu}{2}.$$

Hence (P, N) is uniquely determined by the observed conditional environmental distributions within the stated model class. This establishes point identification. □

Mechanism behind the proof. The proof above states which quantity identifies which; the mechanism behind each step is the following. *Symmetry and independence remove the direction.* For bounded g, the conditional expectation $\mathbb{E}_w[g(B\eta) \mid C, B]$ depends on B only through $|B|$, because $h_g(b) = \mathbb{E}_w[g(b\eta)]$ is even; separability then gives $C \perp B\eta$, which is what lets the environment-wise characteristic functions factor as $\varphi_{X,e} = \varphi_U f_X(s_e \cdot)$ and $\varphi_{Y,e} = \varphi_C f_Y(s_e \cdot)$. *Fourth cumulants recover the unknown scales.* Cumulants add over independent sums and are homogeneous of degree four, so $\kappa_4(X_e) = \kappa_4(U) + a_X(v_e - u_0)^2$ and $\kappa_4(Y_e) = \kappa_4(C) + a_Y(v_e - u_0)^2$: the environment-wise fourth cumulants are an exact quadratic in the treatment variance, so three distinct scales determine its coefficients and hence $u_0 = \mathrm{Var}_w(U)$ and every s_e^2. This is where the identifying variation comes from, and with a Gaussian disturbance it reads $a_X = 0$ and $a_Y = 3\mathrm{Var}_w(B^2)$: heterogeneous transmission amplitude is what recovers the scale from the outcome side, whereas if the amplitude does not vary, $\xi = 0$ and the target is already $t = \sqrt{q}$. *Characteristic-function ratios cancel the baseline.* With the scales known, the products $\Psi_Z(u) = \lim_n \prod_{j \le n} \varphi_{Z,h}(\varrho^j u)/\varphi_{Z,\ell}(\varrho^j u)$ telescope to $f_Z(s_h u)$, because the environment-independent factors φ_U and φ_C cancel; the baseline distributions are eliminated rather than estimated. *An absolute-moment identity delivers the target.* For integrable T, $\mathbb{E}|T| = (2/\pi)\int_0^\infty u^{-2}(1 - \Re\varphi_T(u))\,du$, so the two recovered characteristic functions give J_X and J_Y and their ratio is $t = \mathbb{E}_w|B|$.

Additional requirements for point identification. Relative to the second-moment analysis, this result uses additional distributional information under stronger assumptions: an independent, symmetric standardized disturbance; baseline–amplitude separability $C \perp |B|$; finite fourth moments of U, C, B, η; three environments with pairwise distinct, strictly positive disturbance scales; and, in the varying-amplitude branch, marginal characteristic functions of X_e and Y_e that vanish nowhere. It is what makes the telescoping products and the integral representation well defined, and hence enables the reconstruction that identifies $\mathbb{E}_w|B|$ and, with the identified mean slope, both signed contributions. These conditions are sufficient for the proposed route, and their necessity is not established.

Economic interpretation of separability. The condition $C \perp |B| \mid W$ separates baseline repricing from transmission amplitude within a pretreatment state. In the tin–equity application, $C = Y(U) = A + BU$ is the hypothetical opening gap with the treatment disturbance switched off and the latent baseline component retained, while $|B|$ measures the magnitude of the equity's response to a unit tin innovation. The assumption allows market states to affect both quantities, generating dependence across states, but requires that within a given state, baseline repricing provides no additional information about transmission amplitude. The response direction may still depend on C, and B may depend on U, so treatment can remain confounded. Its economic plausibility depends on whether the chosen pretreatment state captures the common market conditions linking baseline repricing and exposure; this remains an identifying assumption for the point-identification result.

B Statistical Supplement: Estimation and Set Consistency

This appendix gives the sampling results for the moment-based estimators and their induced signed-effect sets.

Sampling and notation. There are fixed numbers K of environments and H of states. Each cell (e, h) contains n_{eh} i.i.d. observations from a fixed law; different cells are independent. Sample sizes are deterministic, with $n = \sum_{e,h} n_{eh}$ and $n_{eh}/n \to \pi_{eh} > 0$. All sampling results assume $\mathbb{E}_{eh}[X^4 + Y^4] < \infty$ in every cell (M4). These sampling assumptions are additional to Condition M.

Define

$$T(x, y) = (x, y, x^2, xy, y^2)^\top, \qquad f(z) = (z_1, z_2, z_3 - z_1^2, z_4 - z_1 z_2, z_5 - z_2^2)^\top.$$

Let $\zeta_{eh} = \mathbb{E}_{eh} T(X, Y)$, $\beta_{eh} = f(\zeta_{eh}) = (m_{X,eh}, m_{Y,eh}, v_{eh}, c_{eh}, r_{eh})^\top$, $\hat{\zeta}_{eh} = n_{eh}^{-1} \sum_i T(X_{ehi}, Y_{ehi})$ and $\hat{\beta}_{eh} = f(\hat{\zeta}_{eh})$. Unsubscripted vectors stack all cells. Write

$$\psi_{eh}(x, y) = Df(\zeta_{eh})\{T(x, y) - \zeta_{eh}\}, \qquad \Lambda_{eh} = \mathrm{Cov}_{eh}(\psi_{eh}(X, Y)),$$

and define $\hat{\psi}_{eh}$ by substituting $\hat{\zeta}_{eh}$ for ζ_{eh}.

Theorem B.1 (Sample moments and feasible covariance)**.** *Under the stated sampling assumptions,*

$$\sqrt{n}(\hat{\beta} - \beta) \Rightarrow N(0, V_\beta), \qquad V_\beta = \mathrm{blockdiag}_{e,h}(\pi_{eh}^{-1} \Lambda_{eh}).$$

Moreover,

$$\hat{\Lambda}_{eh} = \frac{1}{n_{eh}} \sum_i \hat{\psi}_{eh}(X_{ehi}, Y_{ehi}) \hat{\psi}_{eh}(X_{ehi}, Y_{ehi})^\top, \qquad \hat{V}_\beta = \mathrm{blockdiag}_{e,h}\left(\frac{n}{n_{eh}} \hat{\Lambda}_{eh}\right) \to_p V_\beta.$$

Proof. M4 makes $TT^\top$ integrable. The cellwise multivariate CLT and delta method (van der Vaart, 1998), independence, and $\sqrt{n/n_{eh}} \to \pi_{eh}^{-1/2}$ give the joint limit. The law of large numbers for T and $TT^\top$, together with continuity of Df, gives covariance consistency; no eighth moments are needed. □

Moment contrasts. For a fixed state h, set

$$d_h = (v_{2h} - v_{1h}, \ldots, v_{Kh} - v_{1h})^\top, \qquad A_h = d_h \otimes I_2,$$
$$b_h = (c_{2h} - c_{1h}, r_{2h} - r_{1h}, \ldots, c_{Kh} - c_{1h}, r_{Kh} - r_{1h})^\top.$$

The model restrictions give $b_h = A_h \theta_h$ for $\theta_h = (\mu_h, q_h)^\top \in \Theta = \{(\mu, q)^\top : q \geq \mu^2\}$. For a candidate $\theta = (\mu, q)^\top$, write $g_h(\theta) = b_h - A_h \theta$ and $\hat{g}_h(\theta) = \hat{b}_h - \hat{A}_h \theta$. Define

$$R(\theta) = \begin{pmatrix} 0 & 0 & -\mu & 1 & 0 \\ 0 & 0 & -q & 0 & 1 \end{pmatrix}, \qquad A_{eh}(\theta) = R(\theta) \Lambda_{eh} R(\theta)^\top.$$

The covariance $S_h(\theta)$ of the limiting distribution of $\sqrt{n}\{\hat{g}_h(\theta) - g_h(\theta)\}$ has blocks

$$[S_h(\theta)]_{ej} = \mathbf{1}\{e = j\} \frac{A_{eh}(\theta)}{\pi_{eh}} + \frac{A_{1h}(\theta)}{\pi_{1h}}, \qquad e, j = 2, \ldots, K.$$

Its sample version $\hat{S}_h(\theta)$ replaces Λ_{eh} by $\hat{\Lambda}_{eh}$ and π_{eh} by n_{eh}/n. The reference-cell term retains dependence between contrasts.

Theorem B.2 (Ratio estimator)**.** *For a prespecified pair with fixed $d_{ef,h} = v_{eh} - v_{fh} \neq 0$, let*

$$\hat{\theta}_{ef,h} = \frac{(\hat{c}_{eh} - \hat{c}_{fh},\, \hat{r}_{eh} - \hat{r}_{fh})^\top}{\hat{d}_{ef,h}}, \qquad \hat{d}_{ef,h} = \hat{v}_{eh} - \hat{v}_{fh},$$

assigning any fixed value when $\hat{d}_{ef,h} = 0$. Then

$$\sqrt{n}(\hat{\theta}_{ef,h} - \theta_h) \Rightarrow N\left(0, \frac{A_{eh}(\theta_h)/\pi_{eh} + A_{fh}(\theta_h)/\pi_{fh}}{d_{ef,h}^2}\right).$$

Stacked prespecified pairs retain their shared-cell covariance. This is a pointwise result for fixed nonzero contrasts, not a uniform normal approximation as the contrasts vanish.

Proof. On $\{\hat{d}_{ef,h} \neq 0\}$,

$$\hat{\theta}_{ef,h} - \theta_h = \frac{R(\theta_h)[(\hat{\beta}_{eh} - \beta_{eh}) - (\hat{\beta}_{fh} - \beta_{fh})]}{\hat{d}_{ef,h}}.$$

Theorem B.1, $\hat{d}_{ef,h} \to_p d_{ef,h} \neq 0$ and Slutsky's lemma give the result. Applying the same argument jointly proves the stacked-pair statement. □

Theorem B.3 (Constrained minimum distance)**.** *Suppose $d_h \neq 0$. Let $\hat{\Omega}_{n,h} \succ 0$ be data-dependent weights, held fixed during minimization, such that $\hat{\Omega}_{n,h} \to_p \Omega_h$ for a deterministic $\Omega_h \succ 0$. Then*

$$\hat{\theta}_h \in \arg\min_{\theta \in \Theta} \hat{g}_h(\theta)^\top \hat{\Omega}_{n,h} \hat{g}_h(\theta)$$

exists and is unique with probability tending to one, and $\hat{\theta}_h - \theta_h = O_p(n^{-1/2})$, including when $q_h = \mu_h^2$.

Proof. Full column rank gives $\hat{J}_h := \hat{A}_h^\top \hat{\Omega}_{n,h} \hat{A}_h \to_p J_h := A_h^\top \Omega_h A_h \succ 0$. Thus the objective is strictly convex and coercive with probability tending to one, giving a unique minimizer on the closed convex set Θ. Optimality and the triangle inequality yield

$$\sqrt{\lambda_{\min}(\hat{J}_h)}\, \|\hat{\theta}_h - \theta_h\| \leq 2\|\hat{g}_h(\theta_h)\|_{\hat{\Omega}_{n,h}} = O_p(n^{-1/2}), \qquad \|z\|_\Omega = (z^\top \Omega z)^{1/2},$$

where the rate follows from Theorem B.1. Only feasibility of θ_h, not interiority, is used. □

Theorem B.4 (Interior normality and GLS)**.** *Under Theorem B.3 and $q_h > \mu_h^2$,*

$$\sqrt{n}(\hat{\theta}_h - \theta_h) \Rightarrow N(0, \mathcal{V}_h), \qquad \mathcal{V}_h = H_h S_h H_h^\top, \quad H_h = (A_h^\top \Omega_h A_h)^{-1} A_h^\top \Omega_h,$$

where $S_h = S_h(\theta_h)$. Sample substitution gives $\hat{\mathcal{V}}_h = \hat{H}_h \hat{S}_h(\hat{\theta}_h) \hat{H}_h^\top \to_p \mathcal{V}_h$. If $S_h \succ 0$, GLS weights converging to S_h^{-1} attain

$$\mathcal{V}_h^{\mathrm{GLS}} = (A_h^\top S_h^{-1} A_h)^{-1},$$

the smallest covariance in the positive semidefinite order among minimum-distance estimators using these contrasts and positive definite limiting weights.

Proof. Interiority and consistency make the constraint inactive with probability tending to one, so $\hat{\theta}_h - \theta_h = \hat{H}_h \hat{g}_h(\theta_h)$. Theorem B.1, Slutsky's lemma and continuity give the limit and plug-in consistency. For the GLS matrix $H_{*,h}$, both $H_h A_h = H_{*,h} A_h = I_2$, and $S_h H_{*,h}^\top = A_h (A_h^\top S_h^{-1} A_h)^{-1}$. Hence

$$H_h S_h H_h^\top - H_{*,h} S_h H_{*,h}^\top = (H_h - H_{*,h}) S_h (H_h - H_{*,h})^\top \succeq 0.$$

□

The estimated covariance of $\hat{\theta}_h$ is $\hat{\mathcal{V}}_h / n$. Theorem B.4 does not assert a Gaussian limit at $q_h = \mu_h^2$.

Wald inference. Under Theorem B.4, suppose $\mathcal{V}_h$ is positive definite. Standard Wald inference (Newey & McFadden, 1994) then gives the asymptotic $(1 - \alpha)$ marginal intervals

$$\hat{\theta}_{h,j} \pm z_{1-\alpha/2} \sqrt{\hat{\mathcal{V}}_{h,jj}/n}, \qquad j = 1, 2,$$

and the joint confidence ellipse

$$\left\{ \theta \in \mathbb{R}^2 : n(\hat{\theta}_h - \theta)^\top \hat{\mathcal{V}}_h^{-1} (\hat{\theta}_h - \theta) \leq \chi^2_{2,1-\alpha} \right\},$$

where $\hat{\mathcal{V}}_{h,jj}$ is the corresponding diagonal entry of the plug-in covariance and $z_{1-\alpha/2}$ and $\chi^2_{2,1-\alpha}$ denote standard normal and chi-square quantiles. These approximations apply to interior parameters with fixed nonzero variance contrasts; the marginal intervals and the joint ellipse are the columns of Table 3 labelled $\mathrm{cov}(\mu)$, $\mathrm{cov}(q)$ and *joint cov*.

Target-population sets. Fix weights $g_h \geq 0$ with $\sum_h g_h = 1$ and let $\mathcal{H}_+ = \{h : g_h > 0\}$. For feasible $\theta = (\mu_h, q_h)_h$, define

$$m_G = \sum_h g_h \mu_h, \qquad L_G = \sum_h g_h |\mu_h|, \qquad H_G = \sum_h g_h \sqrt{q_h},$$

$$\mathcal{K}(\theta) = \left\{ \left(\frac{t + m_G}{2}, \frac{t - m_G}{2} \right) : L_G \leq t \leq H_G \right\}.$$

Under Corollary A.5, this is the closure of the signed-effect set identified from the environmental means and covariances.

Theorem B.5 (Hausdorff identity and consistency)**.** *For feasible θ, θ', write $\Delta m_G = m'_G - m_G$, and similarly for L_G, H_G. The Euclidean Hausdorff distance satisfies*

$$d_H(\mathcal{K}(\theta'), \mathcal{K}(\theta)) = \frac{1}{\sqrt{2}} \sqrt{(\Delta m_G)^2 + \max\{(\Delta L_G)^2, (\Delta H_G)^2\}}.$$

Thus a feasible estimator consistent in every positive-weight state is Hausdorff consistent.

Proof. The orthogonal map $(P, N) \mapsto ((P + N)/\sqrt{2}, (P - N)/\sqrt{2})$ sends $\mathcal{K}(\theta)$ to $[L_G/\sqrt{2}, H_G/\sqrt{2}] \times \{m_G/\sqrt{2}\}$. The identity follows from the distance between parallel intervals. Consistency follows from continuity of the three summaries. □

Summary and endpoint limits. For the remaining result, stack only positive-weight states in $\theta, \hat{\theta}$; all sums below are over $\mathcal{H}_+$. Write $F = (m_G, L_G, H_G)^\top$. The endpoint vector is

$$F_{\text{end}} = (a_P, a_N, b_P, b_N)^\top = MF, \qquad M = \frac{1}{2} \begin{pmatrix} 1 & 1 & 0 \\ -1 & 1 & 0 \\ 1 & 0 & 1 \\ -1 & 0 & 1 \end{pmatrix},$$

where a and b correspond to $t = L_G$ and $t = H_G$.

Proposition B.6 (Summary and endpoint limits)**.** *Let $\theta, \hat{\theta}$ be feasible, with $q_h > 0$ in every positive-weight state. Suppose $\sqrt{n}(\hat{\theta} - \theta) \Rightarrow Z \sim N(0, V_\theta)$ and $\hat{V}_\theta \to_p V_\theta$. If every $\mu_h \neq 0$, then*

$$\sqrt{n}\{F(\hat{\theta}) - F(\theta)\} \Rightarrow N(0, \Sigma_F), \qquad \Sigma_F = J_F V_\theta J_F^\top, \qquad J_{F,h} = g_h \begin{pmatrix} 1 & 0 \\ \operatorname{sgn}(\mu_h) & 0 \\ 0 & 1/(2\sqrt{q_h}) \end{pmatrix},$$

where J_F concatenates the state blocks. The endpoint limit is $N(0, M\Sigma_F M^\top)$, and both covariance matrices admit consistent plug-in estimates.

Allowing zero means, the summary limit is $(Z_m, Z_L, Z_H)^\top$, where

$$Z_m = \sum_h g_h Z_{\mu_h}, \qquad Z_H = \sum_h \frac{g_h Z_{q_h}}{2\sqrt{q_h}},$$

$$Z_L = \sum_{\mu_h \neq 0} g_h \operatorname{sgn}(\mu_h) Z_{\mu_h} + \sum_{\mu_h = 0} g_h |Z_{\mu_h}|.$$

The endpoint limit is $M(Z_m, Z_L, Z_H)^\top$ and need not be Gaussian. Its covariance is singular because $a_P - a_N - b_P + b_N = 0$ identically.

Proof. The smooth case follows from the delta method. At $\mu_h = 0$, $\sqrt{n}|\hat{\mu}_h| = |\sqrt{n}(\hat{\mu}_h - \mu_h)|$; combining this identity with the smooth expansions of the other coordinates gives the general limit by continuous mapping. The endpoint conclusions follow from multiplication by M. □

Under Theorem B.4 in every positive-weight state, its linear representations and independent cell samples give the required joint limit, with

$$V_\theta = \operatorname{blockdiag}_{h \in \mathcal{H}_+}(\mathcal{V}_h), \qquad \hat{V}_\theta = \operatorname{blockdiag}_{h \in \mathcal{H}_+}(\hat{\mathcal{V}}_h) \to_p V_\theta.$$

No parameter boundary limit is established here.

C Additional Simulation Results

Roadmap. Section C.1 fixes the designs. Section C.2 asks whether the moment-based estimator and its Wald inference are reliable, and what the weak- and zero-contrast designs say about the identifying variation. Section C.3 checks the aggregation over states of Corollary A.5 in a two-state design, and Section C.4 asks whether the point-identification construction of Theorem 3 can be computed at realistic sample sizes and what it recovers beyond the first two moments.

C.1 Simulation designs

Moment-based design. These simulations use a single fixed state, $H = 1$, so that the three environments are the three cells of the design. The structural model is drawn directly, $X_e = U + s_e\eta$, $Y_e = C + s_eB\eta$, $C = A + BU$ and $A = C - BU$, so that the environment enters only through the disturbance scale. The latent variables are

$$U \sim N(0, 0.8), \qquad C = \frac{0.5}{0.8}\,U + \varepsilon, \qquad \varepsilon \sim N(0, 1.6875) \;\perp\; U,$$

where $1.6875 = 2 - 0.5^2/0.8$, so that $\mathrm{Var}(U) = 0.8$, $\mathrm{Var}(C) = 2.0$ and $\mathrm{Cov}(U, C) = 0.5$ as stated; unless noted the slope is Gaussian, $B \sim N(\mu, q - \mu^2)$ with $\mu = 0.25$ and $q = 1.25$, drawn independently of (U, ε); and η is independent standard Gaussian. Treatment is therefore confounded, because U enters both C and X. With Gaussian or finite-support latent variables and independent Gaussian η, every moment used below is finite, so the moment conditions of Appendix B hold. The environment treatment variances are $v_e = u_0 + s_e^2$; the strong contrast uses $s^2 = (0.40, 0.25, 3.22)$, giving $v = (1.20, 1.05, 4.02)$, so that $d = (v_2 - v_1, v_3 - v_1) = (-0.15, 2.82) \neq 0$ and the full-rank condition of Theorem B.3 holds, while $q - \mu^2 = 1.25 - 0.0625 = 1.1875 > 0$, so that the parameter is interior and the regular Gaussian limit of Theorem B.4 applies. The weak contrast uses $s^2 = (0.05, 0.10, 0.15)$, giving $v = (0.85, 0.90, 0.95)$, and the zero contrast uses $s^2 = (0, 0, 0)$. Sample sizes are fixed before each run: for a total n the three cells receive n_e, with proportions n_e/n approaching $\pi = (0.145, 0.589, 0.266)$; observations are i.i.d. within a cell and independent across cells. The baseline sample sizes are $n = 500$, 1,000 and 2,000, close to the 615 sign-state dates of the application and spanning a factor of four, and the large-sample rows are $n = 5{,}000$ and 20,000. The point-estimate experiment uses 1200, 800 and 500 Monte Carlo replications at the three baseline sample sizes and the two large ones use 300 and 150 replications, where the Wald coverage returns to the nominal level. The replication count falls as the sample size grows because the cost of a replication grows with it, which keeps the total simulation budget roughly constant; at these counts the Monte Carlo standard error of a coverage near 0.95 is at most 0.010 for $M = 500$ and 0.006 for $M = 1200$. Serial dependence is introduced only in the warning experiment, where η follows a standardised AR(1) process with $\rho = 0.25$ within each cell at the cell sizes of the $n = 500$ design, over 600 replications.

Point-identification design. The point-identification experiment draws

$$X_e = U + s_e\eta, \qquad Y_e = C + s_eB\eta, \qquad C = A + BU,$$

with scales $s = (0.5, 1.0, 2.0)$. The disturbance η is independent, symmetric about zero and standardized, and is either standard Gaussian or standardized Laplace with $\mathbb{E}\eta^4 = 6$. The amplitude $|B| = R$ takes the values 0.4 and 1.6 with equal probability and is drawn independently of C, while the direction $D = \mathrm{sign}(B)$ depends on C through $\Pr(D = 1 \mid C) = \Phi(0.8C + \kappa_0)$, with $\kappa_0 = 0.6$ in the baseline design; and $C \sim N(0, 0.25)$ and $\varepsilon \sim N(0, 1)$ are independent, with $U = 0.3C + 0.3\varepsilon$, so that $u_0 = \mathrm{Var}(U) = 0.1125$ and $\mathrm{Cov}(U, C) = 0.075$. Treatment is confounded through $C \mapsto D(C) \mapsto B$, so U and B are dependent even though the amplitude R is not, and the population values are $\mu = 0.4225$, $q = 1.36$, $\xi = \mathrm{Var}(B^2) = 1.44$, $t = 1$, $P = 0.7113$ and $N = 0.2887$, all in closed form. This construction also keeps the characteristic-function condition of Theorem 3 satisfied: $\varphi_U(u) = e^{-u_0u^2/2}$ vanishes nowhere, and under Gaussian η $\varphi_{Y,e}(u) = \varphi_C(u)\mathbb{E}_R[e^{-s_e^2u^2R^2/2}]$ is a product of two positive factors, with the Laplace case analogous, so no marginal characteristic function has a zero. An amplitude term $\theta_r(R - \mathbb{E}R)$ inside U would instead contribute a factor $\cos(\theta_r\Delta R\,u)$ and hence zeros at $u = \pi/(2\theta_r\Delta R)$, which the theorem does not allow, and the reported designs avoid it. Each environment is drawn with n observations, so a run labelled n uses $3n$ observations in total. Variants replace the Gaussian disturbance by the Laplace one; set the direction intercept to zero, $\Pr(D = 1 \mid C) = \Phi(0.8C)$, so that

$\mathbb{E}D = 0$ by symmetry of C and hence $\mu = 0$ while D still depends on the baseline, which yields the widest sharp set for N, the half-open interval $(0, \sqrt{q}/2\,] = (0, 0.583]$; and set $|B| \equiv 1$ so that $\xi = 0$ and the constant-amplitude branch applies. In these experiments μ and q are the slopes of separate intercept regressions of the environment covariances and outcome variances on the treatment variances, so each regression carries one over-identifying restriction with three environments; it is the fourth-cumulant quadratics that three distinct variance values determine exactly, by interpolation. The constrained minimum-distance estimator of Section 5.1 is not used here, and the two procedures are linked only through the sample-moment central limit theorem.

C.2 Moment-based estimation and inference

Estimators. The point estimates compared in Figure 3 are the constrained minimum-distance estimator of Theorem B.3, which fits the environment contrasts $(\hat{c}_e - \hat{c}_1, \hat{r}_e - \hat{r}_1)$ to $(\hat{v}_e - \hat{v}_1)(\mu, q)$ over $\Theta = \{(\mu, q) : q \geq \mu^2\}$ under identity or feasible-GLS weights, and the unweighted environment ratio $\hat{\mu} = (\hat{c}_e - \hat{c}_f)/(\hat{v}_e - \hat{v}_f)$, $\hat{q} = (\hat{r}_e - \hat{r}_f)/(\hat{v}_e - \hat{v}_f)$. The environment ratio is the moment-difference estimator of the heteroskedasticity-identification literature (Rigobon, 2003; Lewbel, 2012), with limit given by Theorem B.2; in the equal-mean scale-shift case it coincides with the treatment-noise branch of Kivva et al. (2025) whenever the first cross-moment that differs across environments is the second, and it is reported once, on the first and third environments; the second and third pair identifies the same population parameter, though its estimation precision can differ. All three are approximately unbiased in the strong-variance-shift design and their root-mean-square error decays at the parametric rate. GLS weighting is more accurate than identity weighting at every sample size (0.196 against 0.222 for μ and 0.360 against 0.416 for q at $n = 500$), and the environment ratio is comparable to GLS in this strong regime but unstable in the weak one reported below.

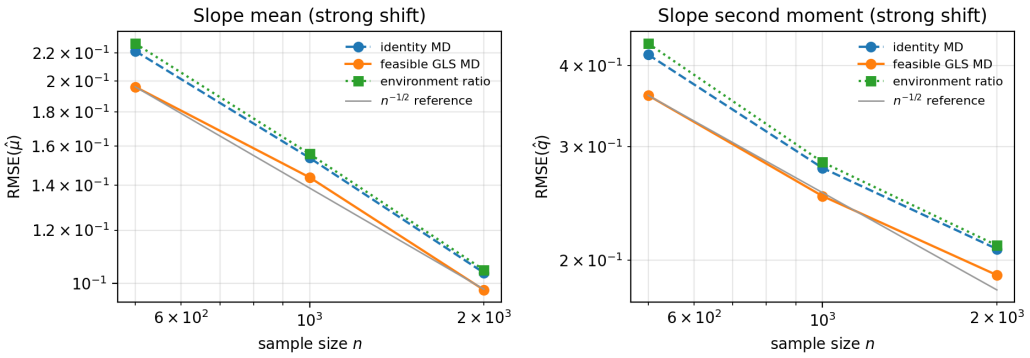


Figure 3: Root-mean-square error of the three estimators of μ (left) and of q (right) under the strong variance shift: identity-weighted minimum distance (dashed), feasible-GLS minimum distance (solid) and the environment ratio (dotted). The grey line is an $n^{-1/2}$ reference through the smallest sample size, so a curve parallel to it decays at the parametric rate.

Feasible covariance and Wald region. The reported intervals and ellipses are feasible in every case: the per-cell influence covariance is estimated within the same cells that produce the point estimates, the constraint rows are evaluated at a first-step identity estimate, and the weights and covariance are $\hat{S}^{-1}$ and $(\hat{A}^\top \hat{S}^{-1} \hat{A})^{-1}$ with $\hat{S} = \hat{C}\,\mathrm{blockdiag}\big((n/n_e)\hat{\Lambda}_e\big)\hat{C}^\top$, where $\hat{\Lambda}_e$ is the sample covariance of the estimated influence function of β in cell e — the feasible covariance of Theorem B.1. With that standardization, $\hat{\mathcal{V}} = (\hat{A}^\top \hat{S}^{-1} \hat{A})^{-1}$ estimates the asymptotic covariance of $\sqrt{n}(\hat{\theta} - \theta)$, so the Wald intervals and ellipse use $\hat{\mathcal{V}}/n$ as the estimated covariance of $\hat{\theta}$. No population covariance enters any reported number; coverage is evaluated at the true parameter, as its definition requires. All reported Wald coverage experiments use interior parameters, $q > \mu^2$. At $q = \mu^2$ the constraint may remain active asymptotically, so the regular Gaussian approximation of Theorem B.4 is not asserted, and no coverage is reported for that case.

Weak and zero variance contrasts. Re-running the design of Section 5.1 with $s^2 = (0.05, 0.10, 0.15)$, whose treatment variances $v = (0.85, 0.90, 0.95)$ differ by at most 12%, leaves the constrained minimum-distance estimator workable but makes the environment ratio unstable. At $n = 500$ the ratio estimator of μ has RMSE 66.83 against 1.12 for the constrained minimum-distance estimator, and even at $n = 5{,}000$ the ratio RMSE is 27.48 against 0.71. The medians tell the same story more quietly — 0.84 against 0.47 at $n = 500$, and 0.39 against 0.36 at $n = 5{,}000$ for μ, with the ratio's median absolute error still decreasing in n — so the ratio's root-mean-square error is dominated by rare extreme draws rather than by its typical behaviour. With $s^2 = (0, 0, 0)$ the treatment variances are identical, the slope moments are not point identified, and the variance contrast is unresolved in every replication and at every sample size examined, including $n = 20{,}000$: the compatible set of (μ, q) is unbounded and carries no information about the slope moments, so no bounded region is available in that regime.

Serial dependence. Appendix B assumes independent observations within a cell. Introducing a within-cell $\mathrm{AR}(1)$ disturbance with $\rho = 0.25$ at the cell sizes of the $n = 500$ design leaves the i.i.d. sandwich interval for μ at 0.942 coverage against a nominal 0.95 over 600 replications, but a naive block-bootstrap percentile interval covers only 0.798. Serial dependence is therefore outside the sampling scheme, and the application, whose treatment is built by a rolling regression, does not meet the conditions of Appendix B.

C.3 Aggregation across states

Design. Corollary A.5 is checked in a two-state design. The target weights g_h are the state shares of the application, $g = (0.470, 0.530)$ for the down and up states; the simulated sample is a separate $n = 1{,}000$ observations. The two states have slope moments $(\mu, q) = (0.5, 1.5)$ and $(-0.2, 1.0)$, environment scales $s^2 = (0.40, 0.25, 3.22)$ and $(0.50, 0.30, 2.60)$, and remaining parameters $(u_0, \mathrm{Cov}(U, C), \mathrm{Var}(C)) = (0.7, 0.4, 1.8)$ and $(0.9, 0.5, 2.2)$. The observations are split across the six cells in proportion to $g_h \pi_e$ with $\pi = (0.145, 0.589, 0.266)$, and each replication estimates (μ_h, q_h) state by state with the feasible-GLS minimum-distance estimator of Section 5.1; the summaries (m_G, L_G, H_G) and the aggregated set $\mathcal{K}(\hat{\theta})$ are then formed with the same weights, over 500 replications. We also refit the same model to the two states pooled, dropping the state index.

Results. Table 6 reports the outcome. The estimated summaries are close to their population values: the biases of $\hat{m}_G$, $\hat{L}_G$ and $\hat{H}_G$ are -0.004, $+0.018$ and -0.014 against standard deviations 0.154, 0.133 and 0.124, the quantities whose joint limit is described by Proposition B.6, and the two components of the target slope dispersion, $\sum_h g_h(q_h - \mu_h^2)$ and $\sum_h g_h(\mu_h - m_G)^2$, have biases -0.043 and $+0.019$. The estimated aggregated set tracks the population one: the mean Euclidean Hausdorff distance is 0.144 (median 0.136), matching the 0.144 that the identity of Theorem B.5 assigns to the sampling errors of the same three summaries, so the set error is accounted for by the summary error rather than by anything specific to the set. Dropping the state index, by contrast, moves the pseudo-true value away from m_G: the pooled mean slope averages 0.165 against 0.129 ($+0.036$, five Monte Carlo standard errors) and the pooled second moment 1.269 against 1.235. Aggregation therefore has to be applied to state-conditional estimates.

C.4 Point identification of the signed effects

Population benchmarks. The design of Section 5.2 is summarized in Table 7: all targets and the sharp set for N are computed in closed form from the data-generating process, so no Monte Carlo error enters the benchmarks.

How the estimator is computed. We compute the six steps of Theorem 3 directly. Sample moments of (X_e, Y_e) give $\hat{\mu}, \hat{q}$ by an intercept regression on $\hat{v}_e$; the fourth cumulants give the interpolating quadratics and hence $\hat{\xi}$; in the varying-amplitude branch the coefficients of those quadratics give $\hat{u}_0$ and $\hat{s}_e^2$, hence $\hat{\varrho} = \hat{s}_\ell / \hat{s}_h$; and the empirical characteristic functions give the telescoping products $\hat{\Psi}_Z$, the integrals $\hat{J}_Z$ and $\hat{t} = \hat{J}_Y / \hat{J}_X$. Each environment pair with $v_\ell < v_h$ is used and the resulting estimates of t are averaged. Two regularizations enter the characteristic-function step, the only place where regularization is needed: the integral is truncated once $|\hat{\varphi}_{Z,e}|$ approaches the sampling noise

Table 6: Aggregation across two states. The target weights $g = (0.470, 0.530)$ are the state shares of the application and the simulated sample is $n = 1{,}000$, with 500 replications. "Population" is the value implied by Corollary A.5 at the true state parameters; mean, bias and standard deviation are Monte Carlo. The Hausdorff distance has population value zero, so its bias equals its mean. The last two rows refit the same model to the two states pooled, dropping the state index; their "Population" entries are the aggregation targets, so the bias column measures the deviation of the pooled fit from those targets rather than from its own pseudo-true value.

Quantity	Population	Mean	Bias	SD
$\hat{m}_G$ (mean slope)	0.129	0.125	-0.004	0.154
$\hat{L}_G$	0.341	0.359	$+0.018$	0.133
$\hat{H}_G$	1.106	1.091	-0.014	0.124
$\sum_h g_h(\hat{q}_h - \hat{\mu}_h^2)$	1.096	1.053	-0.043	0.257
$\sum_h g_h(\hat{\mu}_h - \hat{m}_G)^2$	0.122	0.141	$+0.019$	0.111
Hausdorff distance	0	0.144	$+0.144$	0.068
pooled $\hat{\mu}$, no state index	0.129	0.165	$+0.036$	0.156
pooled $\hat{q}$, no state index	1.235	1.269	$+0.034$	0.275

Table 7: Point-identification experiment: closed-form population values. The last two columns give the sharp set for N implied by (μ, q) alone; the set of the zero-mean design is open at its lower endpoint.

Design	μ	q	ξ	t	P	N	set for N	width
I baseline	0.4225	1.360	1.440	1	0.7113	0.2887	$[0, 0.372]$	0.372
II Laplace η	0.4225	1.360	1.440	1	0.7113	0.2887	$[0, 0.372]$	0.372
III $\mu = 0$	0	1.360	1.440	1	0.5000	0.5000	$(0, 0.583]$	0.583
IV $\xi = 0$	0.4225	1.000	0	1	0.7113	0.2887	$[0, 0.289]$	0.289

floor, and the real part of $\hat{\Psi}_Z$ is projected onto $[-1, 1]$. Because $\hat{\xi}$ is built from fourth cumulants, a replication can produce $\hat{\xi} > 0$ together with an infeasible scale recovery ($\hat{s}_e^2 \leq 0$) even when the population has no amplitude variation. For those replications, and for any replication in which no environment pair yields a usable ratio, we apply the flagged endpoint substitution $t = \sqrt{\hat{q}}$, the largest admissible value of t and hence the largest admissible negative contribution $N = (\sqrt{\hat{q}} - \hat{\mu})/2$. This keeps the algorithm defined in every replication of the designs reported below.

Finite-sample results. Table 8 reports the complete set of estimates for μ, q, ξ, t, P and N, together with the share of replications in which the characteristic-function step is defined, the share that the criterion $\hat{\xi} \leq 0$ assigns to the constant-amplitude branch, and the bias of the within-environment ordinary least squares slope.

Three features of Table 8 are worth extracting. First, the sample-moment quantities converge at a rate close to the square-root rate: the log-log slope of RMSE($\hat{\mu}$), RMSE($\hat{q}$) and RMSE($\hat{\xi}$) against n is -0.52, -0.48 and -0.51, which is consistent with the $-1/2$ reference implied by the sample-moment central limit theorem of Appendix B for $\hat{\mu}$ and $\hat{q}$. The slope for $\hat{\xi}$ is an empirical trend rather than a consequence of that theorem: $\hat{\xi}$ is built from fourth sample moments, so a formal rate statement for it would require the corresponding higher moments — for instance $\mathbb{E}[X_e^8 + Y_e^8] < \infty$ — which these designs satisfy because every latent variable is Gaussian or of finite support, but which the stated assumptions do not impose. The slopes for $\hat{t}$, $\hat{P}$ and $\hat{N}$ are attenuated (-0.23, -0.31, -0.26) because the characteristic-function step, not the moment step, is the binding constraint in these sample sizes. Second, the failures are concentrated where the scale recovery is ill-conditioned: with $\xi = 0$ the curvature a_Y vanishes in population, so only 50 percent of the replications satisfy $\hat{\xi} \leq 0$; 42 percent have $\hat{\xi} > 0$ with an infeasible scale recovery and take the endpoint substitution, and 8 percent reach the characteristic-function route, where the same target is recovered. Third, the within-environment slope is badly biased in every design, by $+0.076$ in the baseline design and

Table 8: Full point-identification results, from the run with the endpoint substitution, so "valid" is one in every cell. "endpoint" is the share in which $\hat{\xi} > 0$ but the scale recovery is infeasible or no environment pair yields a usable ratio, and the largest admissible $t = \sqrt{\hat{q}}$ is reported instead; the complement, less the share with $\hat{\xi} \le 0$, reaches the characteristic-function route. "OLS" is the mean within-environment ordinary least squares slope.

Design	n	$\hat{\mu}$	$\hat{q}$	$\hat{\xi}$	$\hat{t}$	$\hat{P}$	$\hat{N}$	endpoint	$\hat{\xi} \le 0$	OLS
I	1,000	0.4272	1.3660	1.4451	1.0898	0.7585	0.3313	0.15	0.00	0.5022
I	4,000	0.4243	1.3564	1.4657	1.0535	0.7389	0.3146	0.01	0.00	0.4973
I	16,000	0.4211	1.3623	1.4518	1.0562	0.7387	0.3175	0.00	0.00	0.4990
II	16,000	0.4258	1.3610	1.4493	1.0139	0.7198	0.2940	0.00	0.00	0.4953
III	16,000	-0.0005	1.3618	1.4344	1.0576	0.5285	0.5291	0.00	0.00	0.2076
IV	16,000	0.4212	1.0011	0.0009	1.0004	0.7108	0.2896	0.42	0.50	0.5001

$+0.208$ in the symmetric one, while the differencing estimator's bias in μ stays below 0.005 in every design. The residual upward bias of $\hat{t}$, $+0.056$ in Design I at $n = 16{,}000$, does not shrink with n: it is inherited from truncating the Fourier integral at the frequency where the empirical characteristic function reaches its sampling noise floor, not from a failure of identification, and replacing the estimated characteristic functions by their population values on the same grid with the same cutoff leaves it essentially unchanged ($+0.057$).

D Data Construction and Empirical Protocol

Data. The analysis is built from one-minute bars for four traded instruments — three A-share stocks (000960, 000426, 600301) and one exchange-traded fund (512400) — for the CSI 300 index (000300), and for four futures products (Shanghai tin, copper, aluminum and zinc). The estimation itself uses daily observations constructed from these series together with a few intraday summaries of 000960: its 14:30–15:00 return, realized volatility and log volume. The target firm is Yunnan Tin Co., Ltd. (000960), whose opening gap in basis points is the outcome: 10^4 times the log difference between the open of the first one-minute bar of d — the 09:30 print, that is, the opening price set in the opening call auction (09:15–09:25) — and the close of the last bar in the 13:30–15:00 window of the preceding trading date, so the gap spans the overnight break, including weekends and holidays. 000426 and 600301 enter as peer controls and 000300 and 512400 as market benchmarks. Among the futures products tin (SN) carries the treatment and copper (CU), aluminum (AL) and zinc (ZN) enter the same-window adjustment. Because the futures files are organized by contract, each trading date is read from the contract traded over that window, and dates on which the contract changes are flagged as rolls. A daily table aligns all series to the A-share trading date d; night sessions are matched by the exchange calendar, so the session that opens on the preceding evening — including Friday evenings and the eves of holidays — belongs to the next A-share trading date. The fixed-scale treatment is defined from 8 December 2023 to 4 September 2026; the analysis sample used throughout is the 633 dates described below.

Pretreatment controls. W^{treat} collects the 18 pretreatment controls that construct the treatment, all observed from $d-1$ and earlier and listed in Table 9: four 000960 features (previous-day close-to-close return, 14:30–15:00 return, realized volatility, log volume); previous-day returns of the peers and the market, 000426, 600301, 000300 and 512400; previous-day day-session returns of tin, copper, aluminum and zinc; four day-of-week indicators (Monday to Thursday); and the holiday and contract-roll indicators. The tin overnight returns at lags one to three and the 20- and 60-day tin volatilities are deliberately excluded, because the sign state is defined from the lagged tin overnight return and the volatility state from the same series' rolling standard deviation; among tin-related controls, the previous-day day-session return is the only one that is retained.

Treatment definition. The treatment is the fixed-scale residual

$$X_d = 100\{R_d^{\mathrm{SN}} - \widehat{g}_d(M_d, W_d^{\mathrm{treat}})\},$$

where $\widehat{g}_d$ is refitted for each date d by ordinary least squares on the preceding 120 complete dates (at least 60),

$$(\widehat{a}_d, \widehat{b}_d, \widehat{\gamma}_d) = \arg\min_{a,b,\gamma} \sum_{s \in \mathcal{S}_d} \left\{ R_s^{\mathrm{SN}} - a - b' M_s - \gamma' W_s^{\mathrm{treat}} \right\}^2,$$

and $\widehat{g}_d(M_d, W_d^{\mathrm{treat}}) = \widehat{a}_d + \widehat{b}_d' M_d + \widehat{\gamma}_d' W_d^{\mathrm{treat}}$ evaluates that fit at the current window. Here $\mathcal{S}_d$ is the set of the preceding 120 complete dates on which the first stage can be estimated, so the coefficients are known before the treatment window opens and only M_d is realized inside it. M_d collects the same-window returns of the three other metal futures — copper (CU), aluminum (AL) and zinc (ZN) — measured over the same window as R_d^{SN}, and W_d^{treat} is the control set listed above. The same-window metal returns and the pretreatment features are distinct: the former are contemporaneous with the treatment window, whereas the latter are observed before it. No index, peer, or other futures series enters X_d beyond those listed in Table 9, and the treatment is not standardized by a rolling volatility estimate. The factor 100 converts log returns to percentage points: holding the adjustment variables and the fitted rule fixed, a one-percentage-point increase in X_d corresponds to an increase of 0.01 in the raw tin log return.

Why the treatment is adjusted. Three considerations motivate the residual construction. First, tin, copper, aluminum and zinc are jointly driven by a common base-metals factor — global growth expectations, the dollar, risk appetite and Chinese demand — and the equity of a tin producer is exposed to that factor as well. Regressing the opening gap on the raw overnight tin return would therefore mix tin-specific transmission with a common component whose association with the equity is confounded by any shock that moves both markets; the identity of the shock behind a commodity price move is what Kilian & Park (2009) show matters for the equity response. Second, the first stage also removes the part of the overnight return that prior public information already predicts, so what remains is closer to a surprise than to a routine fluctuation. Third, the object of interest is the exposure of a producer's equity to the price of its output, the phenomenon studied for gold miners by Tufano (1998); tin net of the other metals is the version of that exposure that is specific to tin. We therefore read X_d as a tin-specific, unanticipated overnight move, and in the model of Section 3 it plays the role of the treatment: the question is how the opening gap would move if this component were one percentage point larger, holding the other metals and prior information fixed. This is an economic reading of the construction, not a claim of exogeneity; Section 3 states the assumptions under which the causal reading is licensed.

Table 9: Every series that enters the treatment X_d: the same-window metal returns M_d and the 18 pretreatment controls W_d^{treat}. The rolling first stage is estimated on the preceding 120 valid dates (at least 60); the tin columns that are excluded from it are described in the text. Parentheses give the number of columns in each row.

Series	Variable	Enters
SHFE futures: copper, aluminum, zinc	same-window log return (3)	M_d
000960	14:30–15:00 return, realized volatility, log volume, previous-day return (4)	W^{treat}
000426, 600301	previous-day return (2)	W^{treat}
000300, 512400	previous-day return (2)	W^{treat}
SHFE tin, copper, aluminum, zinc	previous-day day-session return (4)	W^{treat}
calendar	day-of-week indicators, Monday–Thursday (4)	W^{treat}
calendar	holiday and contract-roll indicators (2)	W^{treat}

State variables. Two pretreatment states are used, a sign state and a lagged-volatility state. The sign state is the sign of the earlier, completed overnight tin return associated with d^-, so it is known before the treatment window. The lagged-volatility state is the rolling standard deviation of the previous overnight tin returns over 20 trading dates (at least 10 observations, shifted by one date): a date is labelled high-volatility if that value exceeds its expanding median, computed from at least 20 earlier dates, and low-volatility otherwise. The split is close to balanced (280 low- and 353 high-

volatility dates), but the low state contains only five Era III dates, which is too sparse to estimate reliably, so no state-conditional estimates are reported for it.

Environments. The three environments are long calendar segments associated with fee-rule boundaries: Era I before the June 2024 order-submission fee schedule, which covers a group of fifteen futures products including copper, aluminum, zinc and tin rather than tin alone; Era II from that schedule until the December 2025 boundary; Era III from the December 2025 close-today fee on the tin contract SN2601 onward. These boundaries are candidate environments rather than validated interventions: the notices are not tin-specific in 2024, and Era III overlaps a strong tin-market episode, so the variance change cannot be attributed to the fee rule alone.

Sample. Of the 729 dates in the daily table, 708 have the core night-session and opening-gap variables present and 647 also have a non-missing outcome and treatment; dropping the long-holiday windows (eight windows, 14 dates, when margin requirements rise) leaves the fixed-scale treatment available for 633 trading dates, and the sign-state sample, which is also the sample of the motivating regressions, uses 615 of them. The 18 dates that drop out are those whose preceding trading date has no night session, so the lagged overnight return that defines the sign state is undefined; those dates have night sessions of their own, and the volatility state, whose rolling window reaches back before the sample, is available on all 633.

Motivating regressions. The three rows of Table 1 regress the 000960 opening gap in basis points on the fixed-scale treatment X_d in percentage points, adding controls cumulatively, and Table 10 lists the covariates of each row. These steps are not a confounding correction: the first documents the association, and the others show that it is stable to these particular linear adjustments. No environment indicator enters the ladder, which is a descriptive regression that plays no role in the environment-based contrasts. The state-dependent slopes quoted in Section 2 come from adding an interaction of X_d with the sign state and its own intercept to the last specification: 27.34 (Newey–West SE 8.13) after a negative overnight return against 47.40 (Newey–West SE 9.14) after a positive one, with an interaction coefficient of 20.06 (Newey–West SE 10.76, $p = 0.062$). All of these rows use the same common sample of 615 dates and Newey–West standard errors with five lags (Newey & West, 1987).

Table 10: Right-hand-side variables of the motivating regressions of Table 1, by row. Parentheses give the number of columns each line contributes; the outcome is the 000960 opening gap in basis points and the regressor of interest is X_d in percentage points in every row. Controls are added cumulatively, so row 3 contains all of the lines marked 3 and rows 2–3 contain the lines marked 2–3. The two state-interaction specifications add an interaction of X_d with the state and its own intercept to row 3. The three contemporaneous metal returns are the same series that M_d collects, and the remaining nine controls are columns of the pretreatment control set W^{treat}; every covariate of the ladder therefore also enters the construction of X_d, and no environment indicator enters it.

Rows	Variable	Source series
1–3	treatment X_d (the regressor)	as in Table 9
2–3	same-window log return (3)	SHFE copper, aluminum, zinc
3	previous-day day-session return (3)	SHFE copper, aluminum, zinc
3	day-of-week indicators, Monday–Thursday (4)	calendar
3	holiday and contract-roll indicators (2)	calendar

Estimation details. All sample moments use the bias-corrected denominator $n_e - 1$. For each date the futures returns are read from the contract traded in its window, and dates on which that contract changes are flagged as rolls (36 dates); on those dates the night-session close and the preceding 15:00 close may come from different contracts. The three-era estimates come from the constrained minimum-distance estimator with feasible GLS weights, and the merged estimates from the unweighted ratio of pooled Eras I–II to Era III. The target weights are the sample state shares, $289/615$ in the down state and $326/615$ in the up state.

Overall effects and comparison with regression slopes. The state-conditional estimates aggregate to a target population, as in Corollary A.5, with weights g_h equal to the observed state shares, $g = (0.470, 0.530)$. Three different objects follow, and they should be read separately.

(i) The average effect. The aggregated mean slope, $\mu_G = \sum_h g_h \mu_h$, is 28.63 bp/pp in the merged Eras I–II specification and 26.38 bp/pp in the three-era specification. This is the population's average transmission.

(ii) The set beyond the average. Writing $L_G = \sum_h g_h |\mu_h|$ and $H_G = \sum_h g_h \sqrt{q_h}$, the aggregated sharp set is the segment traced by $t \in [L_G, H_G]$ through $\big((t + \mu_G)/2, (t - \mu_G)/2\big)$. Its projections are $P \in [28.63, 83.41]$ and $N \in [0, 54.78]$ in the merged specification and $P \in [26.38, 81.22]$ and $N \in [0, 54.84]$ in the three-era specification.

(iii) A shape benchmark. The Gaussian columns apply the normal-slope formula $P = \sigma\phi(\mu/\sigma) + \mu\Phi(\mu/\sigma)$ and $N = P - \mu$ with $\sigma^2 = q - \mu^2$ inside each state, and then aggregate with the same weights, giving $(P, N) = (69.47, 40.83)$ in the merged specification and $(67.48, 41.11)$ in the three-era specification. These are points, not sets, and they lie inside the corresponding sharp set; they depend on the Gaussian slope assumption and are reported as a benchmark rather than as identification.

Feasibility of the point-identification route. Table 5 reports moment-based sets only. Whether Theorem 3 could also be applied is a question about the estimator rather than about the tin data, so we simulate it: a mechanism is drawn with the generating equations of Section 5.2 and calibrated to the three-era state-conditional estimates of Table 5, using those state-conditional variances and cell sizes, an amplitude dispersion $\xi/q^2 = 0.78$, standard Gaussian η, direction loading $\kappa = 0.8$ and $\mathrm{Var}(U) = 0.1125$, and each of 400 independent simulated samples per state has the state's own sample size (289 observations in the down state, 326 in the up state). These are simulated samples, not resamples of the observed data, which is what makes the target known. Table 11 classifies the replications by the branch the estimator takes. A nonpositive implied scale square ends about half of them (0.4775 in the down state and 0.4250 in the up state) and the constant-amplitude branch takes a further third (0.3725 and 0.2900), so only 0.1500 and 0.2850 of replications reach the characteristic-function step. Those replications do not recover the target: the median $\hat{t}$ is 10,778 against a population value of 134 in the down state and 6,683 against 101 in the up state. Two mechanisms account for the gap. First, the fourth-cumulant inputs are noisy — in the down state $\hat{\xi}$ has median 7.0×10^8 against a population value of 4.7×10^8 and an interquartile range of 8.0×10^8, and the recovered intercept $\hat{u}_0$ has median 0.42 against 0.11; this is what makes the implied-scale check fail in about half of the replications. Second, the cut-off truncates the two integrals unequally. At small u each integrand in J_Z sits at the flat value $\mathbb{E}[Z^2]/2$ and turns down only when u approaches the reciprocal of that disturbance's scale; the two reconstructed disturbances, $s_h\eta$ and $s_h B\eta$, have scales in the ratio $1{:}\sqrt{q}$, so the Y integrand turns down at a $\sqrt{q}$ times smaller frequency than the X integrand. At the cut-off that the empirical characteristic functions support here ($u_{\max} \approx 0.015$), the Y integral has accumulated two thirds to four fifths of its value and the X integral under one per cent, so the truncated denominator shrinks far more than the numerator and the ratio is inflated about eightyfold. Population characteristic functions on the same grid reproduce this (11,054 against a population value of 134), so the offset is truncation rather than sampling noise.

Table 11: Branches of the point-identification estimator at the empirical cell sizes, 400 replications per state. Each row partitions the replications into three exhaustive outcomes: the $\hat{\xi} \le 0$ branch, in which $t = \sqrt{\bar{q}}$; replications with $\hat{\xi} > 0$ and a nonpositive implied scale square, which also fall back to $t = \sqrt{\bar{q}}$; and replications that produce a characteristic-function estimate.

State	$\hat{\xi} \le 0$	$\hat{s}^2 \le 0$	CF estimate
Down	0.3725	0.4775	0.1500
Up	0.2900	0.4250	0.2850