Simultaneous Self-Normalized Inference for Time-Varying Linear IV Models

Jiajing Sun Haiqi Li Yongmiao Hong Jin Zhou*

August 2026

Abstract

We develop simultaneous self-normalized inference for smooth coefficient paths in locally stationary linear instrumental variable (IV) models estimated by local generalized method of moments (GMM). The inferential targets are confidence bands on a prespecified finite grid of dates and tests of whether the entire coefficient path is constant or linear. Pointwise intervals do not control these joint statements, while repeated local covariance estimation introduces date-specific nuisance parameters and must preserve dependence across overlapping windows.

Our construction starts from the feasible influence path associated with a one-step identity-weight local-quadratic GMM estimator. A score formed from residuals of the same local fit contains a leading estimation component; a kernel-weighted quartic projection removes this component. Cumulative squared equivalent-kernel loadings then measure how variance accumulates inside the local window. The adjusted range of the projected path is the primary self-normalizer, and a quadratic root-mean-square functional of the same path provides a matched benchmark. Common calendar-time multipliers calibrate the maximum over dates and coordinates. Identification requires unconditional orthogonality between the IVs and the structural innovation. For the bandwidth order $b_T = (c_b/\sqrt{12})T^{-1/5}$, an additional martingale-difference condition justifies independent and identically distributed (iid) Gaussian multipliers; a separate result for serially dependent scores retains unconditional orthogonality and uses dependent multipliers to reproduce the local long-run covariance. In the martingale-score simulations, simultaneous coverage is close to its nominal level and adjusted-range bands are narrower than matched quadratic bands in every reported design, while relative coverage and power vary across designs. In an application to ten Fama–French industry portfolios, both multiplier specifications reject system-wide constancy of factor exposures, whereas conclusions about linear paths depend on the dependence calibration.

Keywords: self-normalization; adjusted range; quadratic self-normalization; local GMM; finite-grid path inference; time-varying parameters; locally stationary processes.

Journal of Economic Literature (JEL) classification: C12, C13, C14, C22.

*Corresponding author. Email: `jinzhou@xmu.edu.cn`.

1 Introduction

Many economic relationships evolve gradually rather than remaining constant over a long sample. Policy responses may change, market frictions may shift, and factor exposures may move as firms and industries change composition. In such settings, the object of interest is an entire coefficient path: which values are compatible with the data at all reported dates, and can a constant or linear path explain the observed variation?

We study these questions in the smooth time-varying linear IV model $Y_{t,T} = R'_{t,T}\theta_0(t/T)+\varepsilon_{t,T}$, identified by the unconditional orthogonality condition $E(Z_{t,T}\varepsilon_{t,T}) = 0$. Restrictions on the serial dependence of the instrumented score are introduced later for inference and are not part of the identifying moment. The model contains time-varying exogenous regression, exactly identified IV, and overidentified IV. Generalized method of moments (GMM) combines the moment conditions locally; the maintained model is linear in the coefficient path, and GMM supplies its estimator (Hansen, 1982; Giraitis et al., 2021). Local stationarity formalizes gradual change while retaining a stationary approximation around each date (Dahlhaus, 1996; Chen and Hong, 2012; Koo and Linton, 2012).

Pointwise confidence intervals do not answer path-level questions. Even when each interval has nominal marginal coverage, the probability that all reported intervals cover can be substantially smaller. We therefore target a prespecified finite grid of interior dates and calibrate one maximum over dates and coefficient coordinates. The grid is part of the inferential target, matching the empirical practice of reporting coefficients at selected dates without treating a plotted curve as a continuum confidence band. The same maximum also yields tests of whether the full path is constant or linear.

The main challenge is uncertainty calibration. A local IV estimator has a date-specific scale determined by the local design, the instrumented score, and the smoothing window. Adjacent estimates share observations, and serially dependent scores introduce the local long-run covariance $\Lambda(u) = \sum_h E\{U_0(u)U_h(u)'\}$. Estimating and inverting a separate local covariance matrix at every date can be unstable, especially in an overidentified model, and a simultaneous procedure must preserve dependence across windows and coordinates.

Self-normalization provides an alternative by constructing the numerator and denominator from the same influence-score path. Their common scale cancels, reducing reliance on a separate long-run variance estimate for every date and coefficient. For joint inference, the remaining correlation across coordinates and overlapping windows is calibrated by common multipliers. Under the additional martingale-difference condition, the nonzero-lag score covariances vanish, $\Lambda(u)$ equals the lag-zero covariance, and iid multipliers require no long-run-variance bandwidth or block-length choice. Under the serial-dependence conditions studied below, the identifying moment remains unconditional and dependent multipliers approximate $\Lambda(u)$, with a correlation length that must be chosen in finite samples. This advantage builds on recursive and quadratic self-normalization (Lobato, 2001; Shao, 2010b; Shao and Zhang, 2010; Shao, 2015), but a feasible local GMM path cannot be inserted into those constructions directly.

The choice of local polynomial is tied to this inferential target. At the maintained bandwidth $b_T = (c_b/\sqrt{12})T^{-1/5}$, the usual $O(b_T^2)$ local-linear bias is of the same order as the standard error for a generally curved coefficient path. A local-quadratic fit is the lowest-order one-step fit whose approximation bias is negligible at that scale. It has a larger variance constant and requires an interior grid and additional smoothness; the paper uses it because simultaneous bands for unrestricted smooth paths are a primary target, not because derivatives of the coefficient path are being estimated.

The first difficulty comes from using the same local fit twice. The feasible score replaces the structural innovation by a residual from the local polynomial estimator, creating a leading fitting component at the local inference scale. For local-quadratic estimation, the intercept loading and the coefficient estimation error are both quadratic in the local coordinate, so their product lies in the span of a quartic polynomial. We remove this component by a kernel-weighted quartic projection and prove that the projected feasible path is first-order equivalent to the ideal influence path. A heteroskedasticity-consistent type 3 (HC3)-inspired adjustment, based on the diagonal derivative of the local fitted-value map rather than an ordinary least-squares hat matrix, reduces finite-sample leverage shrinkage (MacKinnon and White, 1985).

The second difficulty is that a local-quadratic equivalent kernel may change sign. Calendar time then fails to describe how variance accumulates within a local window. We instead index partial sums by cumulative squared equivalent-kernel loadings. One calendar-time multiplier sequence is shared across all dates and coordinates, preserving the covariance induced by overlapping windows and jointly observed outcomes.

Our primary self-normalizer is the adjusted range of the projected, tied-down influence path. The adjusted range records both its largest upward and downward excursion and adapts the range-based methods of Sun et al. (2022) and Hong et al. (2024) to local GMM path inference. We also study a matched quadratic self-normalizer based on the root mean square of the same path. The estimator, feasible score, projection, variance-time ordering, grid, and multiplier draws are identical across the two procedures, so their comparison isolates the final self-normalizing functional.

The paper makes three contributions. First, it constructs a feasible local influence path for smooth linear IV coefficients and uses its adjusted range to form simultaneous finite-grid bands and tests of constancy or linearity. The construction includes a fitted-score expansion, the quartic projection that removes its leading fitting component, and variance-time ordering. Second, it separates the unconditional IV identifying moment from the dependence conditions used for calibration. It establishes validity for the iid multiplier under martingale instrumented scores and for a dependent multiplier under unconditionally centered serially dependent scores. It also derives the restricted-estimator expansion needed for the path tests and characterizes fixed-grid local power. Third, it provides matched simulation evidence and an application to a fixed system of industry factor-exposure equations.

The numerical evidence distinguishes the main martingale-score results from the serial-dependence sensitivity analysis. Across the martingale-score designs, simultaneous coverage is close to nominal and adjusted-range bands are narrower than the matched quadratic bands in every comparison; relative coverage and rejection rates have no uniform ordering. Tests have strong power against smooth nonlinear paths but, as expected, less power against a localized departure. Experiments with serially correlated scores show that performance depends on both feasible-score construction and dependence calibration. In the Fama–French application, iid and dependent-score calibrations both reject system-wide constancy, while the conclusion for a linear path depends on the multiplier specification.

The paper connects locally stationary moment models, self-normalization, and simultaneous inference for time-varying regressions. The first literature develops time-varying IV estimation, local GMM inference, and tests of smooth structural change (Chen, 2015; Giraitis et al., 2021; Li et al., 2024; Bai, 2026). Our contribution is simultaneous studentization and calibration of a feasible local influence path. Relative to adjusted-range and quadratic self-normalization, the new feature is that the local fitted score must first be projected and then indexed by equivalent-kernel variance contribution (Kim et al., 2015; Sun et al., 2022; Hong et al., 2024). Relative to Gaussian approximation, wild bootstrap, blockwise wild bootstrap, and sieve methods, studentization comes from the same feasible influence path as the local estimator (Zhou and Wu, 2010; Zhang and Wu, 2012; Karmakar et al., 2022; Friedrich and Lin, 2024; Lin et al., 2025).

The analysis focuses on smooth coefficient paths, fixed dimensions, and a fixed interior grid. Regressors may be contemporaneously endogenous, and innovations may be conditionally heteroskedastic. Identification uses the unconditional IV moment throughout. The iid-multiplier benchmark adds a martingale-difference restriction on the instrumented score; the dependent-score extension drops that restriction and uses dependent multipliers. The empirical estimates are retrospective, two-sided local-polynomial paths.

Section 2 defines the linear IV model and inferential targets. Section 3 constructs the local estimator, feasible influence path, self-normalizers, bands, and path tests. Section 4 establishes the theoretical results. Section 5 describes implementation, Section 6 reports the simulations, and Section 7 presents the factor-exposure application. Section 8 concludes. Proofs and additional numerical results are collected in the Appendix.

2 Model and Inferential Targets

This section fixes three objects that play different roles in the analysis: the linear IV model, the unconditional moment that identifies its coefficient path, and the joint statements to be made about that path. Conditions on serial dependence are introduced only when the multiplier calibration is developed.

2.1 Model, identifying moment, and local stationarity

For $t = 1, \dots, T$, let $(Y_{t,T}, R'_{t,T}, Z'_{t,T})'$ form a triangular array, where $R_{t,T} \in \mathbb{R}^p$ is the regressor vector and $Z_{t,T} \in \mathbb{R}^q$ is the instrument vector. We study

$$Y_{t,T} = R'_{t,T}\theta_0(u_t) + \varepsilon_{t,T}, \qquad u_t = t/T, \tag{2.1}$$

where the unknown map $\theta_0 : [0,1] \to \mathbb{R}^p$ changes smoothly with rescaled time. Some components of $R_{t,T}$ may be contemporaneously endogenous, so $E(R_{t,T}\varepsilon_{t,T})$ need not vanish.

For a candidate value $\theta \in \mathbb{R}^p$, define

$$\varepsilon_{t,T}(\theta) = Y_{t,T} - R'_{t,T}\theta, \qquad U_{t,T}(\theta) = Z_{t,T}\varepsilon_{t,T}(\theta).$$

The true instrumented score is $U_{t,T} = U_{t,T}\{\theta_0(u_t)\} = Z_{t,T}\varepsilon_{t,T}$. The identifying condition is the unconditional moment

$$E[U_{t,T}\{\theta_0(u_t)\}] = 0, \qquad t = 1, \dots, T. \tag{2.2}$$

We allow exact identification ($q = p$) and overidentification ($q > p$). Local identification will follow from full column rank of the population regressor–instrument cross moment defined below.

Model (2.1) is globally nonstationary because both its coefficient and the distribution of its data may change with u_t. Local stationarity makes that change estimable: within a shrinking neighborhood of a fixed date u, the data are close to a stationary process indexed by u (Dahlhaus, 1996, 2012). A convenient primitive formulation is the following.

Definition 2.1 (Local stationary approximation)**.** Let $W_{t,T} = (R'_{t,T}, Z'_{t,T}, \varepsilon_{t,T})'$. The triangular array is locally stationary in L^s if, for each $u \in [0,1]$, there is a strictly stationary process $\{W_t(u) : t \in \mathbb{Z}\}$ such that, uniformly in t, T, u,

$$\|W_{t,T} - W_t(u)\|_s \le C\{|u_t - u| + T^{-1}\},$$

and $W_t(u)$ varies Lipschitz continuously with u in L^s.

The definition says that observations in a window centered at u may be analyzed through one stationary approximation, while that approximation is allowed to evolve across windows. Assumption 4.1 gives the moment, smoothness, and dependence conditions used in the proofs. Related local stationarity formulations for time-varying moment models appear in Chen and Hong (2012); Li et al. (2024); Bai (2026).

Write $(R_t(u), Z_t(u), \varepsilon_t(u))$ for the stationary approximation and define

$$D(u) = E\{Z_t(u)R_t(u)'\}, \qquad \Omega(u) = E\{U_t(u)U_t(u)'\}, \qquad U_t(u) = Z_t(u)\varepsilon_t(u).$$

Full column rank of $D(u)$ identifies $\theta_0(u)$ locally. The matrix $\Omega(u)$ is the lag-zero score covariance. It is the relevant covariance under the martingale-score specialization; with serially dependent scores it is replaced by the local long-run covariance in (4.15). This distinction separates identification, which uses (2.2), from calibration.

2.2 Finite-grid bands and path restrictions

Let $\mathcal{U}_M = \{u_1, \ldots, u_M\}$ be a prespecified set of distinct interior dates. The first objective is a collection of intervals $\{\mathcal{C}_{\ell j} : \ell \le M, j \le p\}$ satisfying

$$P\{\theta_{0,j}(u_\ell) \in \mathcal{C}_{\ell j} \text{ for every } (\ell, j)\} \longrightarrow 1 - \alpha. \tag{2.3}$$

Thus the probability statement covers every reported date and coefficient at once. It does not assert coverage at unreported dates and is not a discretized continuum band.

The second objective is to test whether the full coefficient path is constant or linear:

$$H_{0,C}: \quad \theta_0(u) = \gamma_0 \qquad\qquad \text{for every } u \in [0,1], \tag{2.4}$$

$$H_{0,L}: \quad \theta_0(u) = \gamma_{0,0} + u\gamma_{0,1} \qquad\qquad \text{for every } u \in [0,1]. \tag{2.5}$$

The restricted coefficients use the full sample, while the discrepancy from the unrestricted path is measured on $\mathcal{U}_M$. The nulls therefore concern the complete path even though the maximum test statistic is evaluated on a finite grid.

These targets require more than a sequence of pointwise standard errors. The local windows overlap, their scales can change with u, and fitted residuals reuse the same observations as the local estimator. The construction below addresses these three features in turn.

3 Self-Normalized Finite-Grid Inference

The procedure has four components, each with a separate purpose. A local-quadratic GMM fit estimates the coefficient path while keeping smoothing bias negligible at the maintained bandwidth. A quartic projection then removes the leading fitted-score component caused by reusing the local window. Cumulative squared equivalent-kernel loadings describe how variance is distributed through that window and supply the time index for the self-normalizer. Finally, common calendar-time multipliers calibrate one maximum over dates and coefficient coordinates.

3.1 Local GMM estimator and the polynomial-order choice

We first define the local estimator before turning to its influence representation. At $u_\ell \in \mathcal{U}_M$, let

$$z_{t\ell} = \frac{u_t - u_\ell}{b_T}, \qquad w_{t\ell} = b_T^{-1} K(z_{t\ell}).$$

For a local polynomial of degree $r \in \{1, 2\}$, put $p_r(z) = (1, z, \dots, z^r)'$ and define the parameter-major augmented variables

$$X_{t\ell}^{(r)} = R_{t,T} \otimes p_r(z_{t\ell}), \qquad V_{t\ell}^{(r)} = Z_{t,T} \otimes p_r(z_{t\ell}).$$

For $\beta \in \mathbb{R}^{(r+1)p}$, the local weighted sample moment is

$$\widehat{g}_\ell^{(r)}(\beta) = T^{-1} \sum_{t=1}^{T} w_{t\ell} V_{t\ell}^{(r)} \{Y_{t,T} - (X_{t\ell}^{(r)})'\beta\}. \tag{3.1}$$

With $E_{0,r} = I_p \otimes e_{0,r+1}'$, where $e_{0,d} = (1, 0, \dots, 0)' \in \mathbb{R}^d$, the identity-weight local GMM estimator is

$$\widehat{\beta}_\ell^{(r)} = \arg\min_\beta \left\{ \|\widehat{g}_\ell^{(r)}(\beta)\|^2 + \rho_T \|\beta\|^2 \right\}, \qquad \widehat{\theta}_r(u_\ell) = E_{0,r} \widehat{\beta}_\ell^{(r)}. \tag{3.2}$$

The ridge is a numerical guard and may be set to zero when the local system is nonsingular; its rate in (4.5) makes it asymptotically negligible.

The maintained bandwidth is

$$b_T = \frac{c_b}{\sqrt{12}} T^{-1/5}, \qquad n_T = T b_T, \tag{3.3}$$

where $c_b > 0$ is fixed. The simulations vary c_b but not the exponent. This rate makes the polynomial order consequential for inference on a curved path. At an interior point and with a symmetric kernel, the local-linear intercept has the expansion

$$\widehat{\theta}_1(u) - \theta_0(u) = \frac{1}{n_T} \sum_{t=1}^{T} K(z_t) H_D(u) U_{t,T} + \frac{\mu_2}{2} b_T^2 \theta_0''(u) + o_p(n_T^{-1/2} + b_T^2), \tag{3.4}$$

where $\mu_2 = \int z^2 K(z)\, dz$ and $H_D(u) = \{D(u)'D(u)\}^{-1} D(u)'$. Hence $\sqrt{n_T} b_T^2 = (c_b/\sqrt{12})^{5/2}$ does not converge to zero. A centered multiplier law would therefore miss a first-order curvature drift in bands for a general smooth path.

A local-quadratic fit leaves a generic $O(b_T^3)$ approximation remainder, for which $\sqrt{n_T} b_T^3 = O(T^{-1/5}) \to 0$. Under the fixed interior grid and symmetric kernel used here, the odd cubic contribution to the intercept cancels, and Proposition 4.3 sharpens the bias to $O(b_T^4)$. Thus degree two is the lowest one-step local polynomial that makes smoothing bias negligible for the paper's unrestricted bands without changing the $T^{-1/5}$ bandwidth order. It is not chosen to estimate derivatives of θ_0.

This bias reduction has a variance cost. With the Epanechnikov kernel, the intercept equivalent-kernel variance constants are 3/5 for local linear and 5/4 for local quadratic, a variance ratio of 25/12. Local linear also requires less smoothness and is often preferable at boundaries. The present analysis instead targets a fixed interior grid and includes bands for curved paths, for which the nonvanishing local-linear bias is decisive. Constant and linear null paths are reproduced exactly by either fit; using one local-quadratic estimator throughout keeps the bands and path tests on the same feasible score and multiplier calibration. These bias–variance facts follow standard local-polynomial calculations (Fan and Gijbels, 1996; Zhou and Wu, 2010; Kristensen and Lee, 2026); the paper-specific comparison is stated in Lemma A.1, and the local-quadratic expansion is proved in Appendix A. This comparison is analytic. All reported experiments hold the local-quadratic fit fixed, so they do not measure the coverage cost of an uncorrected local-linear band at the maintained bandwidth.

We now specialize (3.2) to $r = 2$. Write

$$p_2(z) = (1, z, z^2)', \qquad q_4(z) = (1, z, z^2, z^3, z^4)',$$

and suppress the superscript (2). If $\beta = (a_1, b_1, c_1, \ldots, a_p, b_p, c_p)'$, then

$$X_{t\ell} = R_{t,T} \otimes p_2(z_{t\ell}), \qquad V_{t\ell} = Z_{t,T} \otimes p_2(z_{t\ell}),$$

so that $X'_{t\ell}\beta = R'_{t,T}\{a + bz_{t\ell} + cz^2_{t\ell}\}$. This parameter-major ordering agrees with $E_0 = I_p \otimes e'_{0,3}$ and with the column-major vectorization used below.

Define

$$\widehat{G}_\ell = T^{-1}\sum_{t=1}^{T} w_{t\ell}V_{t\ell}X'_{t\ell}, \qquad \widehat{s}_\ell = T^{-1}\sum_{t=1}^{T} w_{t\ell}V_{t\ell}Y_{t,T}.$$

Then $\widehat{g}_\ell(\beta) = \widehat{s}_\ell - \widehat{G}_\ell\beta$, and

$$\widehat{H}_\ell = \{\widehat{G}'_\ell\widehat{G}_\ell + \rho_T I_{3p}\}^{-1}\widehat{G}'_\ell, \qquad \widehat{\beta}_\ell = \widehat{H}_\ell\widehat{s}_\ell, \qquad \widehat{\theta}(u_\ell) = E_0\widehat{\beta}_\ell. \tag{3.5}$$

The closed form in (3.5) is simply the solution to the local GMM problem in (3.2); the ridge is not an additional economic restriction.

Let

$$B^0_\ell = \begin{pmatrix} \theta_0(u_\ell)' \\ b_T\{\partial_u\theta_0(u_\ell)\}' \\ \frac{1}{2}b_T^2\{\partial_u^2\theta_0(u_\ell)\}' \end{pmatrix} \in \mathbb{R}^{3\times p}, \qquad \beta^0_\ell = \mathrm{vec}(B^0_\ell),$$

where $\mathrm{vec}(\cdot)$ denotes column-major vectorization. Let

$$Q_2 = \int_{-1}^{1} K(z)p_2(z)p_2(z)'\,dz$$

denote the kernel moment matrix. Define

$$H_D(u) = \{D(u)'D(u)\}^{-1}D(u)', \qquad \ell_2(z) = e'_{0,3}Q_2^{-1}p_2(z), \qquad k_2(z) = K(z)\ell_2(z).$$

The matrix $H_D(u)$ maps the q IV moments into the p coefficient coordinates, while $\ell_2(z)$ extracts the intercept from the local-polynomial fit. The ideal raw influence score and its contribution to the estimator are

$$\begin{aligned} a_{t\ell} &= \ell_2(z_{t\ell})H_D(u_\ell)U_{t,T}, \\ \phi_{t\ell} &= T^{-1}w_{t\ell}a_{t\ell} = n_T^{-1}k_2(z_{t\ell})H_D(u_\ell)U_{t,T}. \end{aligned} \tag{3.6}$$

The vector $a_{t\ell}$ is the score after the IV and intercept loadings have been applied. Multiplication by $T^{-1}w_{t\ell}$ gives the observation's contribution $\phi_{t\ell}$, and $\sum_t \phi_{t\ell}$ is the leading stochastic term in $\widehat{\theta}(u_\ell) - \theta_0(u_\ell)$. The sign is positive because the moment is $V_{t\ell}(Y_{t,T} - X_{t\ell}'\beta)$ and its derivative with respect to β is $-V_{t\ell}X_{t\ell}'$.

3.2 Feasible fitted score

The ideal contribution depends on the unobserved structural innovation. We replace it with a residual from the local GMM fit. Because the same window is used to estimate $\widehat{\beta}_\ell$ and to form the score, this substitution creates a leading fitted-score component. The leverage adjustment below reduces finite-sample residual shrinkage; the quartic projection in the next subsection removes that component at first order.

Let

$$\widehat{e}_{t\ell} = Y_{t,T} - X_{t\ell}'\widehat{\beta}_\ell$$

be the fitted residual. Define the raw leverage

$$\widetilde{h}_{t\ell} = T^{-1} w_{t\ell} X_{t\ell}' \widehat{H}_\ell V_{t\ell},$$

and apply the two distinct truncations

$$h_{t\ell}^c = \min\{\max(\widetilde{h}_{t\ell}, 0), 1 - \epsilon_h\}, \qquad d_{t\ell}^{\mathrm{HC3}} = \max\{1 - h_{t\ell}^c, \epsilon_d\}.$$

This adjustment is inspired by the HC3 heteroskedasticity correction of MacKinnon and White (1985), but it is not an ordinary least-squares hat-matrix correction. Specifically, it uses the diagonal element of the Jacobian of the local-GMM fitted-value map, which measures the sensitivity of each fitted value to its corresponding observation. The resulting leverage-adjusted residual, raw score, and contribution are

$$\widehat{e}_{t\ell}^{\mathrm{HC3}} = \widehat{e}_{t\ell} / d_{t\ell}^{\mathrm{HC3}}, \qquad \widehat{a}_{t\ell} = E_0 \widehat{H}_\ell V_{t\ell} \widehat{e}_{t\ell}^{\mathrm{HC3}}, \qquad \widehat{\phi}_{t\ell} = T^{-1} w_{t\ell} \widehat{a}_{t\ell}. \tag{3.7}$$

It is useful to distinguish the unweighted score contribution $\widehat{a}_{t\ell}/T$ from the kernel-weighted contribution $w_{t\ell}\widehat{a}_{t\ell}/T = \widehat{\phi}_{t\ell}$.

3.3 Removing same-window fitted-score contamination

The polynomial order used in the correction is implied by the structure of the local-quadratic estimator and therefore does not introduce an additional tuning parameter. Specifically, the equivalent-kernel loading for the local intercept is a quadratic function of the local coordinate, while the leading parameter-estimation error also enters through a quadratic polynomial. Their product is therefore a polynomial of degree at most four and lies in the span of

$$q_4(z) = (1, z, z^2, z^3, z^4)'.$$

Accordingly, we project the residual-based fitted score onto this polynomial space and retain the projection residual. This orthogonalization removes the leading component induced by using the same local window to estimate the coefficients and construct the fitted score.

Define

$$\widehat{Q}_{4\ell} = T^{-1} \sum_{t=1}^{T} w_{t\ell} q_4(z_{t\ell}) q_4(z_{t\ell})'.$$

For each coefficient coordinate $j = 1, \ldots, p$, define the weighted least squares coefficient from projecting the feasible raw score $\widehat{a}_{t\ell,j}$ on $q_4(z_{t\ell})$:

$$\widehat{\delta}_{\ell j} = \widehat{Q}_{4\ell}^{-1} T^{-1} \sum_{t=1}^{T} w_{t\ell} q_4(z_{t\ell}) \widehat{a}_{t\ell,j}.$$

The projected score contribution is

$$\widehat{\phi}_{t\ell,j}^{Q4} = T^{-1} w_{t\ell} \{\widehat{a}_{t\ell,j} - q_4(z_{t\ell})' \widehat{\delta}_{\ell j}\}. \tag{3.8}$$

The term in braces is the weighted projection residual. The projection acts on the raw score before multiplication by $T^{-1} w_{t\ell}$; this ordering agrees with the fitted-score expansion and is also used inside every multiplier draw.

3.4 Indexing the path by variance contribution

The equivalent-kernel representation shows how observations enter the leading stochastic component of the estimator. Local-quadratic loadings vary in magnitude and may change sign, so equal increments of calendar time do not represent equal increments of variance. After freezing the smoothly varying score covariance within a local window, the squared equivalent-kernel loading determines each observation's leading relative variance contribution. We keep the observations in chronological order but index partial sums by the cumulative share of these squared loadings.

To construct the cumulative squared-loading index, define the regularized local-quadratic moment matrix

$$\widehat{Q}_{2\ell}^{\mathrm{eq}} = T^{-1} \sum_{t=1}^{T} w_{t\ell} p_2(z_{t\ell}) p_2(z_{t\ell})' + \rho_T^Q I_3$$

and the corresponding estimated equivalent-kernel loading

$$\widehat{\kappa}_{t\ell} = w_{t\ell} e_{0,3}' (\widehat{Q}_{2\ell}^{\mathrm{eq}})^{-1} p_2(z_{t\ell}).$$

The sign of $\widehat{\kappa}_{t\ell}$ determines the direction of the observation's contribution, whereas $\widehat{\kappa}_{t\ell}^2$ measures its relative squared-loading contribution. Let $t_{\ell,1} < \cdots < t_{\ell,N_\ell}$ denote, in chronological order, the observations satisfying $w_{t\ell} > 0$. Define

$$\widehat{\mathcal{V}}_\ell(m) = \frac{\sum_{i=1}^{m} \widehat{\kappa}_{t_{\ell,i},\ell}^2}{\sum_{i=1}^{N_\ell} \widehat{\kappa}_{t_{\ell,i},\ell}^2}, \qquad m = 1, \ldots, N_\ell. \tag{3.9}$$

Thus, $\widehat{\mathcal{V}}_\ell(m)$ records the proportion of the total squared equivalent-kernel loading accumulated through the mth observation in the local window, while preserving the original chronological ordering. Using this index, define the tied-down partial-sum path for the jth parameter coordinate by

$$\widehat{\mathbb{B}}_{\ell j}(m) = \sum_{i=1}^{m} \widehat{\phi}_{t_{\ell,i},\ell,j}^{Q4} - \widehat{\mathcal{V}}_\ell(m) \sum_{i=1}^{N_\ell} \widehat{\phi}_{t_{\ell,i},\ell,j}^{Q4}. \tag{3.10}$$

The first term is the cumulative corrected contribution through the mth observation, whereas the second term subtracts the corresponding squared-loading share of the full-window sum. This transformation ensures that $\widehat{\mathbb{B}}_{\ell j}(N_\ell) = 0$. Moreover, because the quartic projection includes an intercept, its weighted normal equations imply that $\sum_{i=1}^{N_\ell} \widehat{\phi}_{t_{\ell,i},\ell,j}^{Q4} = 0$ in exact arithmetic. The tied-down term is retained because the observed and multiplier paths must use the same transformation.

The adjusted range measures the largest upward and downward excursion of the projected, tied-down path over the trimmed portion of variance time. It is the path-based self-normalizer used by the primary method. Formally,

$$\widehat{R}^{Q4}_{\ell j} = \max_{m:\, \tau\le \widehat{\mathcal{V}}_\ell(m)\le 1-\tau} \widehat{\mathbb{B}}_{\ell j}(m) - \min_{m:\, \tau\le \widehat{\mathcal{V}}_\ell(m)\le 1-\tau} \widehat{\mathbb{B}}_{\ell j}(m). \tag{3.11}$$

Let $\Delta\widehat{\mathcal{V}}_\ell(m) = \widehat{\mathcal{V}}_\ell(m) - \widehat{\mathcal{V}}_\ell(m-1)$ and $\widehat{\mathcal{V}}_\ell(0) = 0$. The matched quadratic denominator is

$$\widehat{Q}^{Q4}_{\ell j} = \left[\frac{\sum_m \widehat{\mathbb{B}}_{\ell j}(m)^2 \Delta\widehat{\mathcal{V}}_\ell(m)\mathbf{1}\{\tau \le \widehat{\mathcal{V}}_\ell(m)\le 1-\tau\}}{\sum_m \Delta\widehat{\mathcal{V}}_\ell(m)\mathbf{1}\{\tau \le \widehat{\mathcal{V}}_\ell(m)\le 1-\tau\}}\right]^{1/2}. \tag{3.12}$$

Because the range and quadratic measures are distinct functionals of the tied-down path, their multiplier distributions generally differ. We therefore compute separate multiplier critical values for the two resulting statistics.

3.5 Multiplier construction

The multiplier procedure approximates the conditional distributions of the range- and quadratic-based statistics by applying random multipliers to the corrected influence contributions and replicating the cumulative-indexing, tied-down, and path-functional transformations used to construct the observed statistics. A single time-indexed multiplier sequence is used for all evaluation points and parameter coordinates. Assigning the same multiplier to a given observation across local windows and coordinates preserves the covariance induced by overlapping windows and, in multivariate systems, the cross-equation dependence among jointly observed outcomes. For each studentization scheme, the multiplier procedure approximates the conditional distribution of the corresponding maximum statistic. Because the adjusted-range and quadratic statistics use distinct path functionals and studentizers, their critical values must be computed separately.

Let

$$\overline{a}_{\ell j} = \frac{\sum_t w_{t\ell}\widehat{a}_{t\ell,j}}{\sum_t w_{t\ell}}$$

be the kernel-weighted raw-score mean. In one draw, generate one iid $N(0,1)$ variable ξ_t per calendar observation. The same ξ_t is used at every u_ℓ and every coordinate. For the supplementary serial-dependence extension, Section 4.4 replaces the iid variables by a stationary Gaussian sequence. The multiplier numerator is

$$\widehat{N}^*_{\ell j} = \sum_{t=1}^{T} T^{-1} w_{t\ell}\xi_t(\widehat{a}_{t\ell,j} - \overline{a}_{\ell j}). \tag{3.13}$$

For the multiplier denominator, the quartic projection is reapplied inside the draw. Multiplication by a nonconstant ξ_t and weighted projection do not commute, so projecting the observed score once and then multiplying it would not reproduce the feasible transformation:

$$\widehat{\delta}^*_{\ell j} = \widehat{Q}^{-1}_{4\ell} T^{-1}\sum_{t=1}^{T} w_{t\ell} q_4(z_{t\ell})\xi_t(\widehat{a}_{t\ell,j} - \overline{a}_{\ell j}),$$

$$\widehat{\phi}^{*,Q4}_{t\ell,j} = T^{-1} w_{t\ell}\left[\xi_t(\widehat{a}_{t\ell,j} - \overline{a}_{\ell j}) - q_4(z_{t\ell})'\widehat{\delta}^*_{\ell j}\right]. \tag{3.14}$$

The projection is therefore recomputed from the multiplied, centered raw scores. Apply (3.10)–(3.12) to

(3.14), holding the estimated variance-time index fixed, to obtain $\widehat{R}^{*,Q4}_{\ell j}$ and $\widehat{Q}^{*,Q4}_{\ell j}$. Define

$$\widehat{S}^*_R = \max_{\ell \le M,\, j \le p} \frac{|\widehat{N}^*_{\ell j}|}{\widehat{R}^{*,Q4}_{\ell j}},$$
$$\widehat{S}^*_Q = \max_{\ell \le M,\, j \le p} \frac{|\widehat{N}^*_{\ell j}|}{\widehat{Q}^{*,Q4}_{\ell j}}.$$

Their conditional $(1-\alpha)$ quantiles are denoted $\widehat{c}_R(1-\alpha)$ and $\widehat{c}_Q(1-\alpha)$. The theoretical procedure uses the unscaled method-specific quantile; the Appendix also shows how scalar changes affect finite-sample coverage and rejection.

3.6 Simultaneous bands and path tests

The calibrated maximum has two uses. Centering it on the unrestricted local estimate gives simultaneous bands; replacing the unknown path by a globally estimated constant or linear restriction gives the corresponding path test. Both uses retain the same method-specific self-normalizer and critical value.

For $F \in \{R, Q\}$, let $\widehat{F}_{\ell j}$ denote $\widehat{R}^{Q4}_{\ell j}$ or $\widehat{Q}^{Q4}_{\ell j}$ and let $\widehat{c}_F(1-\alpha)$ be the corresponding multiplier quantile. The simultaneous intervals are

$$\widehat{\theta}_j(u_\ell) \ \pm\ \widehat{c}_F(1-\alpha)\widehat{F}_{\ell j}, \qquad \ell \le M, \quad j \le p. \tag{3.15}$$

The adjusted-range procedure, obtained by setting $F = R$, is the primary specification. The quadratic procedure uses $F = Q$ and its separately simulated critical value; it is not calibrated with the adjusted-range quantile.

The path tests compare the unrestricted local estimates with a restricted full-sample GMM fit. The restricted estimator imposes the null path structure on the complete sample, whereas $\widehat{\theta}(u_\ell)$ allows the coefficient to vary across evaluation dates. Let K_0 denote the dimension of the path basis: $K_0 = 1$ for a constant path and $K_0 = 2$ for a linear path. Define

$$b_1(u) = 1, \qquad b_2(u) = (1, u)', \qquad L_{K_0}(u) = I_p \otimes b_{K_0}(u)'.$$

Let $\widehat{\gamma}_{K_0}$ be the restricted GMM estimator defined in Section 4. The test statistic is

$$\widehat{T}_{F,K_0} = \max_{\ell \le M,\, j \le p} \frac{|\widehat{\theta}_j(u_\ell) - [L_{K_0}(u_\ell)\widehat{\gamma}_{K_0}]_j|}{\widehat{F}_{\ell j}}. \tag{3.16}$$

The null is rejected when $\widehat{T}_{F,K_0} > \widehat{c}_F(1-\alpha)$. Thus bands and path tests use the same method-specific critical value. The restricted estimator is not recomputed in each multiplier draw because its root-T error is negligible relative to the local rate; Lemma 4.11 establishes this result under primitive conditions for the linear IV model.

4 Finite-Grid Asymptotic Validity

The unconditional moment in (2.2) identifies the linear IV model. The dependence assumptions enter at the inference stage. We first give an iid-multiplier result under the additional martingale-difference condition and then a dependent-score result under unconditional centering and a dependent multiplier. The theory justifies three links in the construction. The first is an influence expansion for the local-quadratic estimator. The second is first-order equivalence between the ideal influence path and the feasible fitted-score path after leverage adjustment and quartic projection. The third is validity of the multiplier approximation for the studentized maximum. Together these results yield simultaneous bands and tests of constant or linear paths. The iid result covers martingale instrumented scores; Section 4.4 gives the serial-dependence result with dependent Gaussian multipliers.

4.1 Regularity conditions

For the true instrumented score and its conditional second moment, define

$$C_{t,T} = U_{t,T}U'_{t,T}, \qquad \Gamma_{t,T} = E(C_{t,T} \mid \mathcal{F}_{t-1,T}).$$

For a causal stationary process $X_t(u) = G_X(u, \ldots, \eta_{t-1}, \eta_t)$, let $X_t^{*(0)}(u)$ be obtained by replacing η_0 by an independent copy and put

$$\delta_s^X(k) = \sup_{u\in[0,1]} \left\| X_k(u) - X_k^{*(0)}(u) \right\|_s, \qquad k \geq 0.$$

All vector norms are Euclidean, all unqualified matrix norms are Frobenius norms, and $\|\cdot\|_{\text{op}}$ denotes the operator norm.

Assumption 4.1 (Fixed-grid linear IV and local GMM conditions)**.** The following conditions hold.

(i) **Bandwidth and evaluation grid.**

The bandwidth is the maintained sequence in (3.3), with a fixed constant $c_b > 0$. Thus $b_T \to 0$ and $n_T = Tb_T \asymp T^{4/5} \to \infty$. The number of evaluation dates M is fixed, and the distinct points in $\mathcal{U}_M = \{u_1, \ldots, u_M\}$ lie in $[\eta, 1-\eta]$ for some $\eta > 0$. Eventually $b_T < \eta$, so every local window is two-sided; distinct fixed-grid windows are eventually disjoint.

(ii) **Primitive causal processes and moments.** There is an independent and identically distributed (iid) innovation sequence $\{\eta_t : t \in \mathbb{Z}\}$ with $\mathcal{F}_t = \sigma(\eta_s : s \leq t)$, and the triangular array can be realized on this innovation space with $\sigma(R_{t,T}, Z_{t,T}, \varepsilon_{t,T}) \subseteq \mathcal{F}_t$ and $\mathcal{F}_{t,T} = \mathcal{F}_t$. There are causal stationary approximations

$$(R_t(u), Z_t(u), \varepsilon_t(u)), \qquad u \in [0,1],$$

and the derived processes are

$$U_t(u) = Z_t(u)\varepsilon_t(u), \quad A_t(u) = Z_t(u)R_t(u)', \quad C_t(u) = U_t(u)U_t(u)', \quad \Gamma_t(u) = E\{C_t(u) \mid \mathcal{F}_{t-1}\}.$$

For some $r > 8$ and a finite constant C, uniformly in T, t, u, v,

$$\begin{aligned} &\|R_{t,T}\|_{2r} + \|Z_{t,T}\|_{2r} + \|\varepsilon_{t,T}\|_{2r} \leq C, \\ &\|X_{t,T} - X_t(u_t)\|_{2r} \leq CT^{-1}, \qquad \|X_t(u) - X_t(v)\|_{2r} \leq C|u-v|, \end{aligned} \tag{4.1}$$

coordinatewise for $X \in \{R, Z, \varepsilon\}$. Lemma 4.2 below derives the corresponding moment, triangular-approximation, time-smoothness, and dependence bounds for U, A, C, and Γ.

(iii) **Primitive physical-dependence summability.** The three primitive stationary approximations satisfy, coordinatewise,

$$\sum_{k=0}^{\infty}(k+1)^2\left\{\delta_{2r}^{R}(k)+\delta_{2r}^{Z}(k)+\delta_{2r}^{\varepsilon}(k)\right\}<\infty. \tag{4.2}$$

Consequently the required summability holds for U, A, C, and Γ. The polynomial factor $(k+1)^2$ is stronger than is needed for a pointwise law of large numbers; it supplies one condition for the maximal prefix laws used below.

(iv) **Additional condition for iid-multiplier inference.**

The unconditional identifying moment in (2.2) is strengthened for this specialization. The data variables $R_{t,T}$, $Z_{t,T}$, and $\varepsilon_{t,T}$ are $\mathcal{F}_{t,T}$-measurable and

$$E(U_{t,T}\mid \mathcal{F}_{t-1,T})=0.$$

For every u, $E\{U_t(u)\mid \mathcal{F}_{t-1}\}=0$. No predictability condition is imposed separately on $Z_{t,T}$ or $R_{t,T}$. Thus a regressor may be contemporaneously endogenous, and an instrument may be contemporaneous, provided the instrumented innovation itself satisfies the displayed conditional orthogonality restriction.

(v) **Smooth path and local population matrices.** The map $\theta_0:[0,1]\to\mathbb{R}^p$ has four bounded continuous derivatives on an open neighborhood of the fixed grid. The matrices $D(u)$ and $\Omega(u)$ introduced in Section 2 satisfy $\Omega(u)=E\{C_t(u)\}=E\{\Gamma_t(u)\}$. The map D has two bounded continuous derivatives and Ω has one bounded continuous derivative on that neighborhood. For constants $c_D, c_\Omega>0$,

$$\inf_u \lambda_{\min}\{D(u)'D(u)\}\ge c_D, \qquad \inf_u \lambda_{\min}\{\Omega(u)\}\ge c_\Omega, \tag{4.3}$$

where the infima are over a fixed neighborhood of $\mathcal{U}_M$.

(vi) **Kernel and polynomial designs.** The kernel K is symmetric, Lipschitz, supported on $[-1,1]$, strictly positive on $(-1,1)$, and normalized by $\int_{-1}^{1}K(z)\,dz=1$. The matrices

$$Q_2=\int_{-1}^{1}K(z)p_2(z)p_2(z)'\,dz, \qquad Q_4=\int_{-1}^{1}K(z)q_4(z)q_4(z)'\,dz \tag{4.4}$$

are positive definite. Define $\ell_2(z)=e_{0,3}'Q_2^{-1}p_2(z)$, $k_2(z)=K(z)\ell_2(z)$, and $c_2=\int_{-1}^{1}k_2(z)^2\,dz$; then $c_2>0$.

(vii) **Numerical regularization and leverage adjustment.** The local criterion ridge $\rho_T\ge 0$ and equivalent-kernel ridge $\rho_T^Q\ge 0$ satisfy

$$\sqrt{n_T}\,\rho_T\longrightarrow 0, \qquad \rho_T^Q\longrightarrow 0. \tag{4.5}$$

Exact nonsingular solves correspond to $\rho_T=\rho_T^Q=0$. For fixed $\epsilon_h,\epsilon_d\in(0,1/2)$, raw leverage is truncated to $[0,1-\epsilon_h]$, after which the residual denominator is bounded below by ϵ_d. The upper truncation and denominator floor are asymptotically inactive because the maximum raw leverage is $o_p(1)$. Section 5.1 records the constants used in the computations.

Any numerical floor $\epsilon_T^{\mathrm{num}}$ used only to guard division by an adjusted-range or quadratic self-normalizer satisfies

$$\epsilon_T^{\mathrm{num}} = o(n_T^{-1/2}). \tag{4.6}$$

The mathematical statistic uses the unfloored denominator; the implementation uses a vanishing guard satisfying (4.6).

Remark 4.1 (Consequences of the maintained bandwidth)**.** Under (3.3),

$$\sqrt{n_T} b_T^2 = O(1), \qquad \sqrt{n_T} b_T^3 = O(T^{-1/5}), \qquad \sqrt{n_T} b_T^4 = O(T^{-2/5}).$$

The first relation explains why an uncorrected local-linear fit is not centered for general curved-path bands. The second is sufficient for the generic local-quadratic remainder, and the third applies to the sharper symmetric interior intercept expansion. Compact support of K restricts observations to $|u_t - u_\ell| \leq b_T$; symmetry supplies the odd-moment cancellation, and positive definiteness of Q_2 identifies the local level, slope, and curvature coefficients.

(viii) **Independent multipliers for the martingale-score case.** Conditionally on the sample, the theorem multipliers are iid $N(0,1)$, independent of the data, and one multiplier is indexed by each calendar observation and reused at every grid point and coordinate. The number of Monte Carlo draws satisfies $B_T \to \infty$.

Lemma 4.2 (Regularity inherited by the score and conditional covariance)**.** *Under Assumption 4.1(ii)–(iii), the product processes U, A, C, and Γ satisfy, coordinatewise,*

$$\sup_{T,t}\{\|U_{t,T}\|_r + \|A_{t,T}\|_r + \|C_{t,T}\|_{r/2} + \|\Gamma_{t,T}\|_{r/2}\} < \infty,$$

$$\|X_{t,T} - X_t(u_t)\|_{s_X} = O(T^{-1}), \qquad \|X_t(u) - X_t(v)\|_{s_X} = O(|u-v|),$$

where $s_X = r$ for $X \in \{U, A\}$ and $s_X = r/2$ for $X \in \{C, \Gamma\}$. Moreover,

$$\sum_{k=0}^{\infty}(k+1)^2\{\delta_r^U(k) + \delta_r^A(k) + \delta_{r/2}^C(k) + \delta_{r/2}^\Gamma(k)\} < \infty.$$

Thus the conditional-covariance conditions used below are consequences of the primitive causal-process conditions rather than additional assumptions on Γ.

Assumption 4.2 (Conditions for global restricted estimators)**.** For $K_0 \in \{1,2\}$ and b_{K_0} and L_{K_0} defined in Section 3, set

$$\mathcal{R}_{K_0,t,T} = R_{t,T} \otimes b_{K_0}(u_t), \qquad \mathcal{Z}_{K_0,t,T} = Z_{t,T} \otimes b_{K_0}(u_t).$$

Assumption 4.1(ii)–(v) holds on all of $[0,1]$, rather than only on a neighborhood of the evaluation grid. Define

$$\mathcal{D}_{K_0} = \int_0^1 D(u) \otimes \{b_{K_0}(u) b_{K_0}(u)'\}\, du,$$
$$\mathcal{S}_{K_0} = \int_0^1 \Omega(u) \otimes \{b_{K_0}(u) b_{K_0}(u)'\}\, du. \tag{4.7}$$

The matrix $\mathcal{D}_{K_0}$ has full column rank and $\lambda_{\min}(\mathcal{S}_{K_0}) > 0$. The number of covariance-weight updates after the identity-weight fit is fixed; the implementation uses two. Let ρ_T^g be the criterion ridge, ρ_T^S the ridge added to the fitted moment covariance, and $\widetilde{\rho}_T^S$ the guard used when that covariance is inverted. They obey

$$\sqrt{T}\,\rho_T^g \longrightarrow 0, \qquad \rho_T^S + \widetilde{\rho}_T^S \longrightarrow 0.$$

The first condition makes criterion-ridge bias negligible at the root-T scale. The second is sufficient because the covariance weight need only converge to $\mathcal{S}_{K_0}^{-1}$ for the fixed-update expansion. Under the corresponding full-path null there is a unique $\gamma_0 \in \mathbb{R}^{pK_0}$ such that

$$\theta_0(u) = L_{K_0}(u)\gamma_0 \quad \text{for every } u \in [0,1].$$

The uniform L^{2r} bounds in Assumption 4.1(ii) imply, in particular,

$$\sup_{T,t} E\{\|Z_{t,T}\|^2|\varepsilon_{t,T}|\|R_{t,T}\| + \|Z_{t,T}\|^2\|R_{t,T}\|^2\} < \infty,$$

which are the explicit envelope moments used when fitted residuals replace innovations in the global covariance updates.

4.2 Preliminary asymptotic approximations

The first proposition gives the local-quadratic influence expansion. The second verifies that the feasible fitted score and its projected partial-sum path are first-order equivalent to their ideal counterparts.

Proposition 4.3 (Local-quadratic influence expansion)**.** *Under Assumption 4.1, uniformly over the fixed grid,*

$$\|\widehat{\beta}_\ell - \beta_\ell^0\| = O_p(n_T^{-1/2} + b_T^3 + \rho_T), \tag{4.8}$$

$$\begin{aligned} \widehat{\theta}(u_\ell) - \theta_0(u_\ell) &= \sum_{t=1}^{T} \phi_{t\ell} + O_p(b_T^4 + b_T^3 n_T^{-1/2} + \rho_T) + o_p(n_T^{-1/2}) \\ &= \sum_{t=1}^{T} \phi_{t\ell} + o_p(n_T^{-1/2}). \end{aligned} \tag{4.9}$$

The frozen-design cubic term cancels exactly because K is symmetric and ℓ_2 is even. Local variation in D multiplies that cubic term by an additional b_T, and the remaining intercept bias is $O(b_T^4)$. Under (3.3), $\sqrt{n_T}b_T^4 = O(T^{-2/5}) \to 0$.

Remark 4.4 (Fixed separable-weight extension)**.** Proposition 4.3 remains true if the identity augmented weight in (3.5) is replaced by a fixed matrix $W_Z \otimes W_P$, where $W_Z \in \mathbb{R}^{q\times q}$ and $W_P \in \mathbb{R}^{3\times 3}$ are symmetric positive definite. Replace $H_D(u)$ by

$$H_{D,W_Z}(u) = \{D(u)'W_Z D(u)\}^{-1} D(u)'W_Z.$$

The local-polynomial factor remains Q_2^{-1} because

$$\{Q_2 W_P Q_2\}^{-1} Q_2 W_P = Q_2^{-1}.$$

An unrestricted estimated $3q \times 3q$ local augmented-moment weight need not factor and therefore lies outside this separable-weight formulation.

Proposition 4.5 (Feasible-score and projected-path equivalence)**.** *Under Assumption 4.1, uniformly over the fixed grid,*

$$n_T \sum_{t=1}^{T} \|\widehat{\phi}_{t\ell} - \phi_{t\ell}\|^2 = o_p(1), \tag{4.10}$$

$$\sqrt{n_T} \max_{t \le T} \|\widehat{\phi}_{t\ell} - \phi_{t\ell}\| = o_p(1). \tag{4.11}$$

Specifically, with

$$r_T^s = b_T + n_T^{-1/2} + T^{1/r}/n_T + \rho_T,$$

the left side of (4.10) is $O_p\{(r_T^s)^2\}$. Let $\phi_{t\ell}^{Q4}$ be the sample-Q_4 projection of the ideal raw score, formed by the same formula as (3.8). Then

$$\sqrt{n_T} \max_{m \le N_\ell} \left\| \sum_{i=1}^{m} \{\widehat{\phi}_{t_{\ell,i},\ell}^{Q4} - \phi_{t_{\ell,i},\ell}^{Q4}\} \right\| = o_p(1). \tag{4.12}$$

The L^2 and maximum-summand bounds remain valid after multiplying the raw score difference by $z_{t\ell}^d$ for each $d = 0, \ldots, 4$. These five weighted full-window sums are the coordinates entering the sample quartic projection coefficient. Writing a projected prefix as its unprojected prefix minus the partial quartic-moment matrix times that full-window coefficient, the preceding bounds imply (4.12) uniformly over all prefixes. Quartic projection is applied once to the raw fitted-score sequence; multiplying the leading degree-four polynomial by an additional power of z need not leave it in the quartic span.

Moreover,

$$\sup_{m \le N_\ell} |\widehat{\mathcal{V}}_\ell(m) - \mathcal{V}(z_{t_{\ell,m},\ell})| = o(1), \qquad \mathcal{V}(z) = c_2^{-1} \int_{-1}^{z} k_2(v)^2 \, dv.$$

The scalar map $\mathcal{V}(z)$ is the variance-time index and is unrelated to the moment-design matrix $A_{t,T} = Z_{t,T} R_{t,T}'$. Quartic projection is sufficient because the local-quadratic intercept loading $\ell_2(z)$ has degree two and the local-polynomial coefficient error enters through $p_2(z)'(\widehat{\beta}_\ell - \beta_\ell^0)$, also of degree two. Their frozen-design product therefore has degree at most four and is annihilated exactly by the sample weighted projection. Terms of degree five or higher can arise both from the cubic-and-higher path remainder and from first-order local variation in the design matrix. They are not projected out. Instead, their normalized prefix processes are negligible: the path remainder is $O_p(\sqrt{n_T} b_T^3) = O_p(T^{-1/5})$, while the local-design and loading remainders are bounded by $O_p\{b_T + \sqrt{n_T} b_T^4 + \sqrt{n_T} b_T \rho_T\} = o_p(1)$.

4.3 Gaussian limits

The endpoint of the limiting path supplies the numerator, while its projected and tied-down trajectory supplies the self-normalizer. The notation below keeps these two roles distinct: G_ℓ denotes the unprojected limiting influence process and Z_ℓ its projected counterpart.

For each ℓ, let $\mathbb{W}_\ell$ be a p-dimensional Brownian random measure on $[-1, 1]$ satisfying

$$\mathrm{Cov}\{d\mathbb{W}_\ell(v)\} = \Sigma_\ell \, dv, \qquad \Sigma_\ell = H_D(u_\ell)\Omega(u_\ell)H_D(u_\ell)'.$$

The measures for distinct grid points are independent. This independence is a conclusion from the calendar-ordered martingale covariance calculation and eventual disjointness of distinct fixed-grid windows; it is not imposed by reordering observations in variance time. Define the p-vector and $5 \times p$ matrix

$$G_\ell(z) = \int_{-1}^{z} k_2(v)\, d\mathbb{W}_\ell(v), \qquad L_{4\ell} = \int_{-1}^{1} q_4(v) k_2(v)\, d\mathbb{W}_\ell(v)',$$

$$\Psi_4(z) = \int_{-1}^{z} K(v) q_4(v)\, dv, \qquad Z_\ell(z) = G_\ell(z) - L_{4\ell}' Q_4^{-1} \Psi_4(z).$$

The index $\mathcal{V}$ is continuous and strictly increasing. Let $z(s) = \mathcal{V}^{-1}(s)$, put $\mathcal{G}_\ell(s) = G_\ell\{z(s)\}$ and $\mathcal{Z}_\ell(s) = Z_\ell\{z(s)\}$, and note that $\mathcal{Z}_\ell(1) = 0$ because $\Psi_4(1) = Q_4 e_{0,5}$ and $L_{4\ell}' e_{0,5} = G_\ell(1)$. For coordinate j, set

$$\mathcal{R}_{\ell j} = \sup_{s \in [\tau, 1-\tau]} \mathcal{Z}_{\ell j}(s) - \inf_{s \in [\tau, 1-\tau]} \mathcal{Z}_{\ell j}(s),$$
$$\mathcal{Q}_{\ell j} = \left\{ \frac{1}{1-2\tau} \int_{\tau}^{1-\tau} \mathcal{Z}_{\ell j}(s)^2\, ds \right\}^{1/2}.$$

Both quantities are strictly positive almost surely. Define

$$S_R^0 = \max_{\ell, j} \frac{|G_{\ell j}(1)|}{\mathcal{R}_{\ell j}}, \qquad S_Q^0 = \max_{\ell, j} \frac{|G_{\ell j}(1)|}{\mathcal{Q}_{\ell j}}. \tag{4.13}$$

Lemma 4.6 (Continuity of central and drifted ratio maxima)**.** *Under Assumption 4.1(v)–(vi), let $d = \{d_{\ell j} : \ell \le M, j \le p\}$ be any deterministic finite vector and, for $F \in \{R, Q\}$, define*

$$S_F(d) = \max_{\ell \le M, j \le p} \frac{|G_{\ell j}(1) + d_{\ell j}|}{\mathcal{F}_{\ell j}}, \qquad \mathcal{F}_{\ell j} = \begin{cases} \mathcal{R}_{\ell j}, & F = R, \\ \mathcal{Q}_{\ell j}, & F = Q. \end{cases}$$

Then $S_F(d)$ has a continuous distribution, and its distribution function is strictly increasing wherever it lies between zero and one. In particular, for every $\alpha \in (0,1)$ the central quantile

$$c_F(1-\alpha) = \inf\{x : P(S_F(0) \le x) \ge 1 - \alpha\}$$

is unique and, for every $\epsilon > 0$, satisfies

$$P\{S_F(0) \le c_F(1-\alpha) - \epsilon\} < 1 - \alpha < P\{S_F(0) \le c_F(1-\alpha) + \epsilon\}. \tag{4.14}$$

The same continuity statement holds for every deterministic drift d.

Thus $S_R^0 = S_R(0)$ and $S_Q^0 = S_Q(0)$, with the corresponding quantiles $c_R(1-\alpha)$ and $c_Q(1-\alpha)$.

Theorem 4.7 (Joint sample and conditional multiplier limits)**.** *Under Assumption 4.1, the normalized local-estimator endpoints and the polygonal interpolants of all ideal and feasible quartic-projected prefix paths converge jointly, over the fixed grid and coordinates, to $\{G_\ell(1), Z_\ell(\cdot)\}_{\ell \le M}$. Conditionally on the data, in probability, the joint multiplier numerator and the draw-specific quartic-projected multiplier paths converge weakly to an independent copy of the same Gaussian collection. Common calendar-time multipliers retain finite-sample window overlap, while fixed distinct interior windows become disjoint asymptotically after their joint covariance has been formed.*

More precisely, all sample and multiplier process convergence takes place in the product space

$$\mathbb{H}_M = \{\mathbb{R}^p \times C[-1,1]^p\}^M,$$

equipped with the maximum of the Euclidean and uniform metrics. If $\widehat{\mathbb{X}}_T^*$ *denotes the stacked multiplier endpoints and paths and* $\mathbb{X}^0$ *is an independent copy of the Gaussian limit, then*

$$\sup_{f\in \mathrm{BL}_1(\mathbb{H}_M)} \left| E^* f(\widehat{\mathbb{X}}_T^*) - Ef(\mathbb{X}^0)\right| \longrightarrow_p 0,$$

where $\mathrm{BL}_1(\mathbb{H}_M)$ *is the unit bounded-Lipschitz class.*

For $F \in \{R, Q\}$ *and every continuity point* x *of the law of* S_F^0*,*

$$P^*(\widehat{S}_F^* \le x) \longrightarrow_p P(S_F^0 \le x).$$

If the empirical multiplier quantile is based on $B_T \to \infty$ *draws, then*

$$\widehat{c}_R(1-\alpha) \longrightarrow_p c_R(1-\alpha), \qquad \widehat{c}_Q(1-\alpha) \longrightarrow_p c_Q(1-\alpha).$$

The next theorem states the simultaneous coverage result that motivates the construction. It combines the feasible-score equivalence, Gaussian limit, and multiplier approximation established above.

Theorem 4.8 (Feasible simultaneous adjusted-range bands)**.** *Under Assumption 4.1, the adjusted-range bands in* (3.15) *have simultaneous fixed-grid coverage*

$$P\Big[\theta_{0,j}(u_\ell) \in \{\widehat{\theta}_j(u_\ell) \pm \widehat{c}_R(1-\alpha)\widehat{R}_{\ell j}^{Q4}\} \text{ for all } (\ell, j)\Big] \longrightarrow 1-\alpha.$$

The critical value is the unscaled iid-multiplier quantile. The numerical floor in (4.6) *is asymptotically inactive. A fixed scalar applied in the numerical work is a finite-sample calibration and is not part of this theorem.*

Corollary 4.9 (Matched quadratic bands)**.** *Under the same conditions, replacing* $\widehat{R}_{\ell j}^{Q4}$ *by* $\widehat{Q}_{\ell j}^{Q4}$ *and using the separately simulated critical value* $\widehat{c}_Q(1-\alpha)$ *gives simultaneous coverage* $1-\alpha+o(1)$*. The matched quadratic functional requires its separately simulated critical value.*

4.4 Supplementary extension to serially dependent scores

This subsection retains the unconditional identifying moment (2.2) and drops the additional martingale-difference restriction used by the iid-multiplier result. Serial dependence changes the covariance matrix of the local score process, but not the deterministic time transformation in (3.9). To see the distinction, define

$$\Gamma_U(u,h) = E\{U_0(u)U_h(u)'\}, \qquad h \in \mathbb{Z},$$

and the local long-run covariance

$$\Lambda(u) = \sum_{h\in\mathbb{Z}} \Gamma_U(u,h). \tag{4.15}$$

For every fixed lag h, the product of two shifted local kernel loadings has the same limit as the squared loading because $h/(Tb_T) \to 0$. Summing over lags therefore replaces the lag-zero covariance $\Omega(u)$ by $\Lambda(u)$, while the limiting variance-time index remains

$$\mathcal{V}(z) = c_2^{-1} \int_{-1}^{z} k_2(v)^2 \, dv.$$

Assumption 4.3 (Serially dependent score and dependent multiplier)**.** Assumption 4.1(i)–(iii) and (v)–(vii) holds. In place of Assumption 4.1(iv) and (viii), impose the following conditions.

(a) The instrumented score is unconditionally centered: $E(U_{t,T}) = 0$ for every t, T, and $E\{U_t(u)\} = 0$ for every u.

(b) The matrix in (4.15) is uniformly positive definite on a fixed neighborhood of the evaluation grid. Thus, for some $c_\Lambda > 0$,

$$\inf_u \lambda_{\min}\{\Lambda(u)\} \geq c_\Lambda.$$

The moment, local-stationarity, and weighted physical-dependence conditions in Assumption 4.1(ii)–(iii) imply that the series in (4.15) is absolutely and uniformly convergent and that Λ is continuous.

(c) Independently of the data, generate a stationary Gaussian sequence from

$$\xi_{1,T} \sim N(0,1), \qquad \xi_{t,T} = \varrho_T \xi_{t-1,T} + \sqrt{1-\varrho_T^2}\, e_t, \qquad \varrho_T = \exp(-1/L_T), \tag{4.16}$$

where the e_t are iid $N(0,1)$. One realization of this sequence is used at every grid point and for every coordinate. Its dependence length obeys

$$L_T \to \infty, \qquad \frac{L_T(\log T)^2}{\sqrt{n_T}} \to 0, \qquad L_T b_T^2 (\log T)^2 \to 0, \qquad \frac{L_T^2(\log T)^2}{T} \to 0. \tag{4.17}$$

The number of multiplier draws satisfies $B_T \to \infty$. Centering and standardizing each realized multiplier sequence, as in the numerical implementation, is permitted.

With $b_T = (c_b/\sqrt{12})T^{-1/5}$, one admissible class is

$$L_T = \max\{2, \lfloor c_\ell T^a \rfloor\}, \qquad 0 < a < 2/5, \tag{4.18}$$

for any fixed $c_\ell > 0$. The numerical analysis sets $a = 1/4$ and reports several values of c_ℓ rather than treating the rate condition as a finite-sample tuning rule. Combining (4.17) with the explicit transfer rate in Proposition 4.5 gives

$$L_T n_T \sum_t \|\widehat{\phi}_{t\ell} - \phi_{t\ell}\|^2 = o_p(1),$$

the fitted-to-ideal score bound needed by the dependent multiplier.

For the dependent-score limit, replace Σ_ℓ in the Gaussian construction above by

$$\Sigma_\ell^D = H_D(u_\ell)\Lambda(u_\ell)H_D(u_\ell)'. \tag{4.19}$$

Let G_ℓ^D, Z_ℓ^D, $\mathcal{R}_{\ell j}^D$, and $\mathcal{Q}_{\ell j}^D$ denote the resulting processes and functionals, and define S_R^D and S_Q^D as in (4.13) with these substitutions.

Theorem 4.10 (Dependent-score multiplier validity)**.** *Under Assumption 4.3, the conclusions of Propositions 4.3 and 4.5 continue to hold. Jointly over the fixed grid, the normalized local-estimator endpoints* *and the feasible quartic-projected prefix paths converge to $\{G_\ell^D(1), Z_\ell^D(\cdot)\}_{\ell \le M}$. Distinct grid-point limits are independent. The sample and conditional multiplier convergence take place in the product space $\mathbb{H}_M$ defined in Theorem 4.7.*

Conditionally on the data, the multiplier numerator and the draw-specific quartic-projected paths formed with (4.16) converge in probability to an independent copy of the same Gaussian collection. Consequently, for $F \in \{R, Q\}$ and every continuity point x of the corresponding limit law,

$$P^*(\widehat{S}_F^* \le x) \longrightarrow_p P(S_F^D \le x),$$

and the empirical multiplier quantile consistently estimates the $(1-\alpha)$ quantile of S_F^D. The adjusted-range and matched quadratic bands therefore have simultaneous fixed-grid coverage $1-\alpha+o(1)$ when each uses its own dependent-multiplier critical value.

The iid Gaussian multiplier in Theorem 4.7 is the calibration used for the martingale-score specialization. For general serially dependent scores, a fixed value of L_T estimates the exponentially tapered covariance $\sum_h \exp(-|h|/L_T)\Gamma_U(u,h)$ rather than $\Lambda(u)$. Fixed lengths and nonoverlapping block multipliers are useful comparisons in a given sample, whereas (4.18) supplies an asymptotically valid class of dependent-score calibrations. The rate restrictions do not determine c_ℓ, and the finite-sample experiments show that coverage can be sensitive to this constant and to feasible-score construction.

The path-test results below use Assumption 4.1 with the iid multiplier for martingale scores and Assumption 4.3 with the dependent multiplier for serially dependent scores.

4.5 Restricted estimators and path tests

For $K_0 \in \{1,2\}$ let b_{K_0}, $\mathcal{R}_{K_0,t,T}$, $\mathcal{Z}_{K_0,t,T}$, and $L_{K_0}(u)$ be as in Assumption 4.2. Define

$$\widehat{\mathcal{D}}_{K_0} = T^{-1}\sum_{t=1}^{T} \mathcal{Z}_{K_0,t,T}\mathcal{R}'_{K_0,t,T}, \qquad \widehat{s}_{K_0} = T^{-1}\sum_{t=1}^{T} \mathcal{Z}_{K_0,t,T}y_{t,T}.$$

For a positive-definite W, put

$$\mathcal{H}_T(W) = \{\widehat{\mathcal{D}}'_{K_0} W \widehat{\mathcal{D}}_{K_0} + \rho_T^g I\}^{-1}\widehat{\mathcal{D}}'_{K_0} W.$$

Begin with $\widehat{\gamma}^{(0)} = \mathcal{H}_T(I)\widehat{s}_{K_0}$. For a fixed number $J \ge 1$ of covariance-weight updates, recursively form

$$\widehat{e}_t^{(j-1)} = y_{t,T} - \mathcal{R}'_{K_0,t,T}\widehat{\gamma}^{(j-1)},$$

$$\widehat{\mathcal{S}}_{K_0}^{(j)} = T^{-1}\sum_{t=1}^{T} \mathcal{Z}_{K_0,t,T}\mathcal{Z}'_{K_0,t,T}(\widehat{e}_t^{(j-1)})^2 + \rho_T^S I,$$

$$\widehat{W}^{(j)} = \{\widehat{\mathcal{S}}_{K_0}^{(j)} + \widetilde{\rho}_T^S I\}^{-1}, \qquad \widehat{\gamma}^{(j)} = \mathcal{H}_T(\widehat{W}^{(j)})\widehat{s}_{K_0},$$

where the second inversion guard $\widetilde{\rho}_T^S$ is zero under an exact solve and otherwise satisfies Assumption 4.2. Let J denote the number of covariance-weight updates after the identity-weight fit and set $\widehat{\gamma}_{K_0} = \widehat{\gamma}^{(J)}$; the implementation uses $J = 2$.

Lemma 4.11 (Global constant and linear-path estimators)**.** *Suppose Assumptions 4.1 and 4.2 hold and the full-path null $\theta_0(u) = L_{K_0}(u)\gamma_0$ holds for every $u \in [0,1]$. For $K_0 = 1$ (constancy) and $K_0 = 2$ (a linear* *path),*

$$\sqrt{T}(\widehat{\gamma}_{K_0}-\gamma_0)=\mathcal{H}_{K_0}T^{-1/2}\sum_{t=1}^{T}\{U_{t,T}\otimes b_{K_0}(u_t)\}+o_p(1), \tag{4.20}$$

where

$$\mathcal{H}_{K_0}=\{\mathcal{D}_{K_0}'\mathcal{S}_{K_0}^{-1}\mathcal{D}_{K_0}\}^{-1}\mathcal{D}_{K_0}'\mathcal{S}_{K_0}^{-1}.$$

Consequently, $\widehat{\gamma}_{K_0}\to_p\gamma_0$, $\widehat{\gamma}_{K_0}-\gamma_0=O_p(T^{-1/2})$, and

$$\sqrt{T}(\widehat{\gamma}_{K_0}-\gamma_0)\Rightarrow N\big(0,\{\mathcal{D}_{K_0}'\mathcal{S}_{K_0}^{-1}\mathcal{D}_{K_0}\}^{-1}\big)\,.$$

In particular,

$$\sqrt{n_T}\max_{\ell\le M}\|L_{K_0}(u_\ell)(\widehat{\gamma}_{K_0}-\gamma_0)\|=O_p(\sqrt{b_T})=o_p(1). \tag{4.21}$$

A restriction imposed only at the grid points is insufficient: the global estimator uses all observations, so an off-grid violation generally produces a nonzero population moment and an $O(1)$ pseudo-true displacement.

Corollary 4.12 (Restricted estimators with serially dependent scores)**.** *Suppose the conditions of Assumptions 4.3 and 4.2 hold on* $[0,1]$*, except that the conditional score restriction in Assumption 4.1(iv) is replaced by unconditional centering as in Assumption 4.3(a). Under the full-path constant or linear null, define*

$$\mathcal{L}_{K_0}=\int_0^1\Lambda(u)\otimes\{b_{K_0}(u)b_{K_0}(u)'\}\,du$$

and

$$\mathcal{H}_{K_0}^0=\{\mathcal{D}_{K_0}'\mathcal{S}_{K_0}^{-1}\mathcal{D}_{K_0}\}^{-1}\mathcal{D}_{K_0}'\mathcal{S}_{K_0}^{-1}.$$

Then

$$\sqrt{T}(\widehat{\gamma}_{K_0}-\gamma_0)=\mathcal{H}_{K_0}^0T^{-1/2}\sum_{t=1}^{T}\{U_{t,T}\otimes b_{K_0}(u_t)\}+o_p(1)\Rightarrow N\{0,\mathcal{H}_{K_0}^0\mathcal{L}_{K_0}(\mathcal{H}_{K_0}^0)'\}.$$

In particular, (4.21) continues to hold. The lag-zero fitted moment covariance used in the finite number of GMM weight updates determines $\mathcal{H}_{K_0}^0$; serial correlation enters the sampling covariance through $\mathcal{L}_{K_0}$.

Remark 4.13 (A differentiable nonlinear-in-parameter path)**.** A prespecified path $g(u,\gamma)$ can replace $L_{K_0}(u)\gamma$ only after adding a compact parameter space, uniform identification of the full-sample population moment, twice continuous differentiability in γ, an integrable uniform envelope for the first two derivatives, and full column rank of the integrated Jacobian. The usual Taylor expansion then replaces $L_{K_0}(u)$ by $\partial g(u,\gamma_0)/\partial\gamma'$. The constant and linear paths studied here avoid these additional conditions.

The finite-grid statistic used below is the path-test statistic in (3.16).

Theorem 4.14 (Constancy and prespecified linear-path tests)**.** *Under Assumptions 4.1 and 4.2, the test that rejects the full-path null for $K_0=1$ or $K_0=2$ when $\widehat{T}_{R,K_0}>\widehat{c}_R(1-\alpha)$ has asymptotic size α. The same adjusted-range multiplier quantile used for a band is valid because the root-T restricted-estimator error is negligible at the local rate and the limiting central maximum is unchanged.*

Under the serial-dependence conditions of Corollary 4.12, the same conclusion holds with the dependent-multiplier critical value.

Under a fixed smooth alternative, define the population limits of the fixed number of global GMM updates recursively by replacing sample moments and covariances by their integrals. Suppose every update *is uniquely defined and its population matrices are nonsingular, and denote the final limit by* $\gamma^*_{K_0}$. *If* $\theta_{0,j}(u_\ell) \neq [L_{K_0}(u_\ell)\gamma^*_{K_0}]_j$ *for at least one grid coordinate, then* $\widehat{T}_{R,K_0} \to_p \infty$ *and the test is consistent.*

Corollary 4.15 (Matched quadratic path tests)**.** *The conclusions of Theorem 4.14 hold with* $\widehat{T}_{Q,K_0}$ *and the separately simulated quadratic critical value* $\widehat{c}_Q(1-\alpha)$.

Corollary 4.16 (Fixed finite system)**.** *Let* $s = 1, \ldots, S$ *index a fixed number of equations,*

$$y_{s,t,T} = R'_{s,t,T}\theta_{s,0}(u_t) + \varepsilon_{s,t,T}, \qquad U_{s,t,T} = Z_{s,t,T}\varepsilon_{s,t,T},$$

and stack $U^{\mathrm{sys}}_{t,T} = (U'_{1,t,T}, \ldots, U'_{S,t,T})'$. *For each equation, let* $H_{D,s}(u) = \{D_s(u)'D_s(u)\}^{-1}D_s(u)'$ *and define*

$$H^{\mathrm{sys}}_D(u) = \mathrm{diag}\{H_{D,1}(u), \ldots, H_{D,S}(u)\}.$$

Write

$$\Omega^{\mathrm{sys}}(u) = E\{U^{\mathrm{sys}}_0(u)U^{\mathrm{sys}}_0(u)'\}, \qquad \Lambda^{\mathrm{sys}}(u) = \sum_{h\in\mathbb{Z}} E\{U^{\mathrm{sys}}_0(u)U^{\mathrm{sys}}_h(u)'\}.$$

The mapped covariance at u_ℓ *is*

$$\Sigma^{\mathrm{sys}}_\ell = H^{\mathrm{sys}}_D(u_\ell)\Omega^{\mathrm{sys}}(u_\ell)H^{\mathrm{sys}}_D(u_\ell)',$$

in the martingale-score case and is

$$\Sigma^{\mathrm{sys},D}_\ell = H^{\mathrm{sys}}_D(u_\ell)\Lambda^{\mathrm{sys}}(u_\ell)H^{\mathrm{sys}}_D(u_\ell)'$$

in the serial-dependence case. Suppose all equations are adapted to one filtration and the causal local stationarity, moment, and physical-dependence conditions in Assumption 4.1(ii)–(iii) hold for the jointly stacked array. In the martingale-score case, also suppose

$$E(U^{\mathrm{sys}}_{t,T} \mid \mathcal{F}_{t-1,T}) = 0.$$

Alternatively, impose the unconditional centering and long-run covariance conditions of Assumption 4.3 on the jointly stacked score and use the dependent multiplier in (4.16)*. Suppose the equation-specific smoothness and rank conditions in Assumption 4.1(v)–(vii) hold uniformly over the fixed system and that the relevant mapped covariance,* $\Sigma^{\mathrm{sys}}_\ell$ *or* $\Sigma^{\mathrm{sys},D}_\ell$*, is positive definite on the displayed stacked coordinates at every grid point. Then Theorem 4.7 in the martingale-score case, or Theorem 4.10 in the serial-dependence case, and Theorem 4.8, together with Corollary 4.9, remain valid when the maximum is taken jointly over equations, grid points, and coefficient coordinates. If the constant or linear restriction holds on every equation's full path and the corresponding global conditions hold uniformly, Theorem 4.14 and Corollary 4.15 extend to the same system maximum. One common calendar-time multiplier is reused across every equation, grid point, and coordinate.*

4.6 Fixed-grid local alternatives

Let $a_T = n_T^{-1/2}$ and consider the triangular structural path

$$\theta_T(u) = L_{K_0}(u)\gamma_0 + a_T\Delta(u), \qquad u \in [0,1], \tag{4.22}$$

where Δ has four bounded continuous derivatives and

$$y_{t,T} = R'_{t,T}\theta_T(u_t) + \varepsilon_{t,T}.$$

The law of $(R_{t,T}, Z_{t,T}, \varepsilon_{t,T})$ and all constants in the preceding assumptions are uniform along this triangular family. Define

$$h_{K_0,\Delta} = \int_0^1 \{D(u)\Delta(u)\} \otimes b_{K_0}(u)\, du, \qquad \pi_{K_0,\Delta} = \mathcal{H}_{K_0} h_{K_0,\Delta},$$
$$d_{\ell,\Delta} = \Delta(u_\ell) - L_{K_0}(u_\ell)\pi_{K_0,\Delta}. \tag{4.23}$$

The second term is the component of the local drift absorbed by estimation of the restricted global path.

Proposition 4.17 (Fixed-grid local power)**.** *Under* (4.22) *and the assumptions of Lemma 4.11,*

$$a_T^{-1}(\widehat{\gamma}_{K_0} - \gamma_0) \longrightarrow_p \pi_{K_0,\Delta}, \tag{4.24}$$

and jointly over the fixed grid,

$$\sqrt{n_T}\{\widehat{\theta}(u_\ell) - L_{K_0}(u_\ell)\widehat{\gamma}_{K_0}\} \Rightarrow G_\ell(1) + d_{\ell,\Delta}. \tag{4.25}$$

The observed and multiplier self-normalizers retain their central limits. Therefore, for $F \in \{R, Q\}$,

$$\widehat{T}_{F,K_0} \Rightarrow S_F(d_\Delta) = \max_{\ell,j} \frac{|G_{\ell j}(1) + d_{\ell,\Delta,j}|}{\mathcal{F}_{\ell j}}, \tag{4.26}$$

where $\mathcal{F}_{\ell j} = \mathcal{R}_{\ell j}$ for $F = R$ and $\mathcal{F}_{\ell j} = \mathcal{Q}_{\ell j}$ for $F = Q$. Let d_Δ denote the Mp-vector obtained by stacking $d_{\ell,\Delta,j}$ over (ℓ, j) in lexicographic order. By Lemma 4.6, the exact asymptotic local-power function is

$$\pi_F(\Delta) = P\{S_F(d_\Delta) > c_F(1-\alpha)\}. \tag{4.27}$$

It equals α when $d_\Delta = 0$.

Under the serial-dependence specification, the same conclusions hold with G_ℓ, $\mathcal{F}_{\ell j}$, S_F, and c_F replaced by their dependent-score counterparts defined after (4.19)*.*

For every $d_\Delta \neq 0$, $P\{S_F(\lambda d_\Delta) > c_F(1-\alpha)\} \to 1$ as $|\lambda| \to \infty$; hence power exceeds size for all sufficiently large local-drift magnitudes. The two normalizers induce different noncentral laws through $\mathcal{R}_{\ell j}$ and $\mathcal{Q}_{\ell j}$, so their local-power comparison is design-specific.

Under the fixed alternatives in Theorem 4.14, a nonzero displayed pseudo-true discrepancy yields consistency. The two normalizers can therefore be compared through their design-specific noncentral laws and the matched numerical experiments.

The analysis focuses on smooth coefficient paths in linear IV models, fixed dimensions, and a fixed interior grid. It uses a one-step identity-weight local-quadratic estimator, a same-window leverage-adjusted fitted score, and a quartic projection under martingale-difference or serially dependent instrumented scores. The corresponding calibrations use iid or growing-length dependent multipliers. Boundary inference, growing grids, discontinuous paths, nonlinear GMM, weak identification, long-memory scores, and unrestricted estimated local augmented weights require additional arguments.

5 Implementation

Implementation follows the construction in Section 3: estimate the local coefficient path, form the same-window feasible score, remove its quartic fitting component, order the projected path by variance contribution, and simulate the joint maximum with common multipliers. Algorithm 1 records these steps for the bands, and Algorithm 2 shows how the same critical value is used for the path restrictions.

Algorithm 1 Finite-grid self-normalized bands

Require: Data $\{y_{t,T}, R_{t,T}, Z_{t,T}\}_{t=1}^{T}$, interior grid $\mathcal{U}_M$, kernel K, bandwidth constant c_b, trimming fraction τ, multiplier draws B, and level α

Ensure: Local-quadratic GMM path and simultaneous finite-grid bands

1: Set $b_T = (c_b/\sqrt{12})T^{-1/5}$. For each u_ℓ, form $z_{t\ell} = (u_t - u_\ell)/b_T$, $p_2(z) = (1, z, z^2)'$, $X_{t\ell} = R_{t,T} \otimes p_2(z_{t\ell})$, and $V_{t\ell} = Z_{t,T} \otimes p_2(z_{t\ell})$. Estimate $\widehat{\beta}_\ell$ by the one-step identity-weight criterion in (3.5) and retain $\widehat{\theta}(u_\ell) = E_0\widehat{\beta}_\ell$.

2: Compute the fitted residual $\widehat{e}_{t\ell}$, the diagonal fitted-value derivative $\widetilde{h}_{t\ell}$, its truncated value $h^c_{t\ell}$, the leverage-adjusted residual denominator $d^{\mathrm{HC3}}_{t\ell}$, and the raw score $\widehat{a}_{t\ell}$ in (3.7). Observations outside the support of the local kernel have zero contribution.

3: Form the local-quadratic equivalent-kernel weight $\widehat{\kappa}_{t\ell}$. Order the observations in the local window by calendar time and compute the variance-time index $\widehat{\mathcal{V}}_\ell(m)$ in (3.9).

4: For each coordinate j, regress $\widehat{a}_{t\ell,j}$ on $q_4(z_{t\ell}) = (1, z_{t\ell}, z^2_{t\ell}, z^3_{t\ell}, z^4_{t\ell})'$ using the kernel weights. Accumulate the projected contributions $\widehat{\phi}^{Q4}_{t\ell,j}$ in calendar order and compute the adjusted range $\widehat{R}^{Q4}_{\ell j}$ over $\tau \le \widehat{\mathcal{V}}_\ell(m) \le 1-\tau$.

5: Center the raw score by its kernel-weighted mean. For each draw, generate one iid Gaussian multiplier per calendar observation and use the same sequence at every grid point. Apply the quartic projection to $\xi_t(\widehat{a}_{t\ell,j} - \overline{a}_{\ell j})$, and calculate the multiplier numerator and adjusted-range denominator defined in Section 3.

6: Let $\widehat{c}_R(1-\alpha)$ be the empirical $(1-\alpha)$-quantile of

$$\max_{\ell \le M} \max_{j \le p} \frac{|\widehat{N}^*_{\ell j}|}{\widehat{R}^{*,Q4}_{\ell j}}.$$

Report the simultaneous bands

$$\widehat{\theta}_j(u_\ell) \pm \widehat{c}_R(1-\alpha)\widehat{R}^{Q4}_{\ell j}, \qquad \ell = 1, \ldots, M, \quad j = 1, \ldots, p.$$

Algorithm 1 states the iid-multiplier specialization under the additional martingale-score condition. For the serial-dependence extension, the iid sequence is replaced by the stationary Gaussian sequence in (4.16); all other steps are unchanged. Theorem 4.8 concerns the unscaled multiplier quantile, which is also the main specification in the simulations and application. Factors 1.05, 1.10, and 1.15 multiply the same quantile in the band and both path tests only in the reported finite-sample sensitivity calculations.

The matched quadratic procedure replaces each adjusted range $\widehat{R}^{Q4}_{\ell j}$ by

$$\widehat{Q}^{Q4}_{\ell j} = \left[\frac{\sum_m \widehat{\mathbb{B}}_{\ell j}(m)^2 \Delta\widehat{\mathcal{V}}_\ell(m) \mathbf{1}\{\tau \le \widehat{\mathcal{V}}_\ell(m) \le 1-\tau\}}{\sum_m \Delta\widehat{\mathcal{V}}_\ell(m) \mathbf{1}\{\tau \le \widehat{\mathcal{V}}_\ell(m) \le 1-\tau\}} \right]^{1/2}, \qquad \Delta\widehat{\mathcal{V}}_\ell(m) = \widehat{\mathcal{V}}_\ell(m) - \widehat{\mathcal{V}}_\ell(m-1),$$

where $\widehat{\mathcal{V}}_\ell(0) = 0$, and applies the same replacement within every multiplier draw. The quadratic procedure uses its separately simulated critical value $\widehat{c}_Q(1-\alpha)$; the adjusted-range critical value is not reused. Thus the two procedures differ only in the final functional applied to the quartic-projected partial-sum path.

Algorithm 2 Constrained-null path tests

Require: Output of Algorithm 1, restriction index $K_0 \in \{1, 2\}$, and significance level α

Ensure: Adjusted-range statistic $\widehat{T}_{R,K_0}$ and rejection decision

1: Set $b_1(u) = 1$, $b_2(u) = (1, u)'$, and $L_{K_0}(u) = I_p \otimes b_{K_0}(u)'$. Impose $\theta_0(u) = L_{K_0}(u)\gamma_0$ for every $u \in [0, 1]$. Thus $K_0 = 1$ tests constancy and $K_0 = 2$ tests a linear path.

2: Form $\mathcal{R}_{K_0,t,T} = R_{t,T} \otimes b_{K_0}(u_t)$ and $\mathcal{Z}_{K_0,t,T} = Z_{t,T} \otimes b_{K_0}(u_t)$. Estimate $\widehat{\gamma}_{K_0}$ by the global linear GMM updates in Section 4 and evaluate $L_{K_0}(u_\ell)\widehat{\gamma}_{K_0}$ on $\mathcal{U}_M$.

3: Form

$$\widehat{T}_{R,K_0} = \max_{\ell \le M} \max_{j \le p} \frac{|\widehat{\theta}_j(u_\ell) - [L_{K_0}(u_\ell)\widehat{\gamma}_{K_0}]_j|}{\widehat{R}^{Q4}_{\ell j}}.$$

4: Compare $\widehat{T}_{R,K_0}$ with the same $\widehat{c}_R(1 - \alpha)$ used for the adjusted-range bands. The global restricted estimator is not re-estimated within each multiplier draw because its root-T error is asymptotically negligible at the local rate.

5: Reject the full-path null when the finite-grid statistic satisfies $\widehat{T}_{R,K_0} > \widehat{c}_R(1 - \alpha)$.

5.1 Computational constants

Table 1 records the numerical choices shared by the simulations and application unless a table states otherwise. The ridges and denominator guard vanish with sample size and serve only to protect matrix inversion and division in finite precision.

Table 1: Computational constants in the main implementation

Quantity	Value or rule
Kernel	Epanechnikov
Bandwidth	$(1.5/\sqrt{12})T^{-1/5}$
Local polynomial / score projection	Quadratic / quartic
Variance-time trimming	$\tau = 0.10$
Local GMM fit	One identity-weight step
Restricted constant and linear fits	Two covariance-weight updates
Convergence tolerance for restricted fits	10^{-6}
Criterion and equivalent-kernel ridges	$\min\{10^{-8}, T^{-2}\}$
Leverage truncation	$[0, 1 - 10^{-8}]$
Adjusted-residual denominator floor	10^{-6}
Self-normalizer division guard	$\min\{10^{-10}, T^{-2}\}$
Multiplier quantile	Type-7 empirical quantile
Multiplier draws	499 per Monte Carlo sample; 4,999 in the application

5.2 Grid, trimming, and multiplier simulation

The evaluation grid is part of the inferential target. A band over $\mathcal{U}_M$ controls the probability that every grid coordinate is covered; it is not a discretized claim about every point in $[0, 1]$. The theory keeps M fixed and places the grid inside the sample. In the numerical work, the constants and grid endpoints are chosen so that every kernel window is two-sided. This keeps changes in the bandwidth constant from introducing a separate boundary rule.

All reported adjusted-range calculations set $\tau = 0.10$. Trimming is defined in variance time, not by deleting a fixed number of calendar observations. Windows with different equivalent-kernel shapes therefore use the same fraction of accumulated squared weight. The same trimming rule is applied to the observed denominator and every multiplier denominator. For the quadratic procedure, the denominator in (3.11) is replaced by the variance-time root mean square in the display above, including normalization by the retained variance-time mass.

The multiplier simulation generates one calendar-time sequence per draw and reuses it across the entire grid and all coordinates. Independent sequences at different grid points would erase the dependence created by overlapping windows and would calibrate a different maximum. The nine main Monte Carlo designs have martingale instrumented scores and use iid Gaussian multipliers with $B = 499$ draws within each simulated sample. The serial-dependence experiments compare iid, dependent Gaussian, block, and sieve calibrations. The empirical analysis uses $B = 4{,}999$ and reports iid and growing-length dependent Gaussian results. Each dependent multiplier realization is centered and standardized; Lemma A.15 shows that this finite-sample normalization is asymptotically neutral. In both cases the finite-B quantile uses the type-7 linear interpolation rule for the empirical distribution of the joint maximum. The asymptotic results allow any empirical quantile rule that is consistent as $B \to \infty$. Centering the fitted raw score by its kernel-weighted mean enforces a conditionally centered multiplier numerator; Lemma A.9 shows that this centering does not alter the first-order limit.

The main decision rule uses the unscaled method-specific multiplier quantile. A separate high-replication null experiment evaluates the factors 1.00, 1.05, 1.10, and 1.15 on common samples and multiplier draws. These factors are sensitivity calculations, not theoretical critical values and not data-dependent choices. In every column the same factor is applied to the band, constancy test, and linear-path test.

6 Simulation study

The simulations answer four questions raised by the theory. First, does the adjusted-range statistic deliver accurate simultaneous coverage and correctly sized tests of constant and linear coefficient paths? Second, how much power do the path tests have against smooth and localized departures? Third, what changes when the adjusted range is replaced by a matched quadratic self-normalizer? Fourth, how do the feasible score and multiplier calibration behave when instrumented scores are serially dependent?

Every reported matched comparison uses the same simulated sample, local estimator, fitted-score construction, quartic projection, variance-time ordering, and multiplier draws. Adjusted-range and quadratic self-normalization differ only in the final functional applied to the projected partial-sum path. Additional experiments vary the projection and variance-time ordering separately to examine the two corrections required by the feasible local influence process.

6.1 Data-generating processes

The design families add one difficulty at a time. The exactly identified models establish baseline size and power, the overidentified models introduce endogenous regressors, and the autoregressive and conditionally heteroskedastic models vary persistence and scale. A separate family studies serial correlation in the instrumented score.

The exactly identified designs use

$$y_t = \theta_{1,0}(u_t) + \theta_{2,0}(u_t)x_t + \varepsilon_t, \qquad R_t = Z_t = (1, x_t)', \qquad u_t = t/T.$$

One design uses an iid standard-normal regressor. A second sets $x_t = 0.5x_{t-1} + v_t$, with $v_t \sim N(0,1)$; this is a first-order autoregression, denoted AR(1). We discard 200 initial observations. Except in the conditional-heteroskedastic design, $\varepsilon_t \sim N(0,1)$ independently of the regressor innovations. The four coefficient paths are

$$\begin{aligned}
&\text{constant:} && \theta_0(u) = (1, 0.6)', \\
&\text{linear:} && \theta_0(u) = (1 + 0.30u,\ 0.6 - 0.20u)', \\
&\text{sinusoidal:} && \theta_0(u) = (1 + 0.40\sin(2\pi u),\ 0.6 + 0.25\cos(2\pi u))', \\
&\text{localized:} && \theta_0(u) = (1 + 0.25d(u),\ 0.6 + 0.05d(u))',
\end{aligned} \qquad d(u) = \exp\{-80(u - 0.70)^2\}.$$

The constant path is a null for both tests. The linear path is a null for the linear-path test and an alternative to constancy. The sinusoidal and localized paths are alternatives to both restrictions. The localized design complements the sinusoidal design by concentrating the departure in a relatively short part of the evaluation interval.

Two overidentified designs allow the regressor to be endogenous. Let

$$Z_t = (1, z_{1t}, z_{2t}, z_{3t})', \qquad x_t = 1.2z_{1t} + 0.9z_{2t} + 0.7z_{3t} + v_t,$$

where the excluded instruments are independent standard normal variables and $(\varepsilon_t, v_t)'$ is iid bivariate normal with unit marginal variances and correlation 0.6. We use the constant and sinusoidal paths in this model. Finally, a generalized autoregressive conditional heteroskedasticity model of order $(1,1)$, denoted GARCH(1,1), combines the AR(1) regressor and constant path with

$$\varepsilon_t = h_t^{1/2}\eta_t, \qquad h_t = 0.2 + 0.15\varepsilon_{t-1}^2 + 0.80h_{t-1}, \qquad \eta_t \sim N(0,1).$$

This last design is deliberately a heavy-tail stress test. For Gaussian η_t,

$$E(0.15\eta_t^2 + 0.80)^3 = 1.012625 > 1,$$

so the usual sixth-moment condition for this GARCH recursion fails; it is not covered by the higher-moment assumptions used in the proofs. Together these choices produce nine designs: three iid exactly identified models, two iid overidentified IV models, three homoskedastic AR(1)-regressor models, and one conditional-heteroskedastic model. They vary the path shape, identification structure, regressor dynamics, and conditional variance. The GARCH row is retained as an outside-assumptions stress test rather than as a calibration design for the maintained theory.

An additional experiment examines the finite-sample behavior of the serial-dependence extension. It retains the AR(1) regressor and replaces the independent regression error by

$$\varepsilon_t = \rho_\varepsilon\varepsilon_{t-1} + \sqrt{1 - \rho_\varepsilon^2}\,\eta_t, \qquad \rho_\varepsilon \in \{0.3, 0.6\},$$

where η_t is standard normal and independent of the regressor process. The error has unit variance. Hence the intercept score ε_t and the slope score $x_t\varepsilon_t$ are both serially correlated. We consider the constant and sinusoidal coefficient paths. For the null designs we also vary exact versus overidentification, use both oracle and fitted scores, and compare several dependent Gaussian tapers, block calibrations, and a score sieve. These comparisons assess sensitivity to score construction and dependence calibration; they do not choose a single finite-sample rule for the dependence class covered by the theorem.

6.2 Implementation and evaluation criteria

The main experiment sets $T = 500$, uses the Epanechnikov kernel, and evaluates the path at the 25 points $\{0.20, 0.225, \ldots, 0.80\}$. The bandwidth is $b_T = (1.5/\sqrt{12})T^{-1/5}$. The estimator is one-step identity-weight local quadratic. Scores are computed from the leverage-adjusted fitted residual, the variance-time index uses the squared local-quadratic equivalent-kernel weights, and the score projection is quartic. There is no pilot bandwidth.

The experiments therefore evaluate joint calibration and compare the two self-normalizers while holding the local-quadratic estimator and feasible score fixed. As noted in Section 3.1, they are not a finite-sample comparison of polynomial orders.

Let N_{MC} denote the number of Monte Carlo replications. Each of the nine martingale-score designs uses $N_{\mathrm{MC}} = 300$ independently generated samples and $B = 499$ iid Gaussian multiplier draws per sample. One multiplier is drawn for each observation and reused across all grid points and coordinates. Within each method and replication, the very same multiplier maximum supplies the critical value for the simultaneous band, constancy test, and linear-path test. Thus a reported rejection rate and the corresponding coverage rate implement one decision rule rather than separately simulated critical values.

For each design we report simultaneous coverage over all 50 displayed coordinates, average simultaneous-band width, and rejection frequencies for the two path restrictions. Rejection under a true restriction measures size; rejection under a false restriction measures power. Parenthesized entries are binomial Monte Carlo standard errors. Replication-level paired differences are used when two self-normalizers are compared. We also record the minimum self-normalizer and record direct measures of the fitted-score approximation; those diagnostics are reported in Appendix B.

6.3 Finite-sample performance

Table 2: Adjusted-range finite-grid inference at $T = 500$ using the unscaled multiplier critical value

Design	Coverage	Band width	Constancy reject	Linear-path reject
iid regressor, constant path	0.940 (0.014)	1.207	0.023 (0.009)	0.017 (0.007)
iid regressor, linear path	0.933 (0.014)	1.215	0.070 (0.015)	0.057 (0.013)
iid regressor, sinusoidal path	0.937 (0.014)	1.205	0.753 (0.025)	0.333 (0.027)
iid overidentified IV, constant path	0.963 (0.011)	1.606	0.013 (0.007)	0.013 (0.007)
iid overidentified IV, sinusoidal path	0.973 (0.009)	1.582	0.430 (0.029)	0.190 (0.023)
AR(1) regressor, constant path	0.933 (0.014)	1.140	0.063 (0.014)	0.053 (0.013)
AR(1) regressor, sinusoidal path	0.947 (0.013)	1.137	0.803 (0.023)	0.410 (0.028)
GARCH errors, constant path	0.930 (0.015)	2.267	0.060 (0.014)	0.063 (0.014)
AR(1) regressor, localized path	0.940 (0.014)	1.136	0.083 (0.016)	0.053 (0.013)

Notes: Entries for coverage and rejection are Monte Carlo estimates with binomial standard errors in parentheses. Every design uses $N_{\mathrm{MC}} = 300$ samples, $B = 499$ common iid Gaussian multiplier draws, and 25 grid points from 0.20 to 0.80. The bandwidth is $b_T = (1.5/\sqrt{12})T^{-1/5}$. The estimator is one-step identity-weight local quadratic; the fitted score uses the stated leverage adjustment, the variance-time index uses squared equivalent-kernel weights, and the score projection is quartic. The unscaled method-specific multiplier critical value is used for coverage and both path tests. Scalar multiples of this value are reported only as finite-sample sensitivity calculations in Table B.3. The GARCH row is a heavy-tail stress test outside the maintained moment conditions.

Table 2 reports the main adjusted-range results with the unscaled multiplier critical value. Simultaneous coverage ranges from 0.930 to 0.973 across the nine designs. For a constant path, constancy rejection is

0.023 with an iid regressor, 0.013 in the overidentified IV model, and 0.063 with an AR(1) regressor. The linear-path test rejects 0.057 of the time when the path is linear. In the GARCH stress test, coverage is 0.930, and the constancy and linear-path rejection frequencies are 0.060 and 0.063. Parenthesized Monte Carlo standard errors quantify sampling uncertainty.

The sinusoidal alternatives separate power against constancy from power against a fitted linear path. The respective rejection frequencies are 0.753 and 0.333 with an iid regressor, 0.803 and 0.410 with an AR(1) regressor, and 0.430 and 0.190 in the overidentified IV model. The wider IV bands accompany the lower power in that design. The localized departure is difficult to detect path-wide: rejection is 0.083 against constancy and 0.053 against linearity. Appendix B reports the factors 1.05, 1.10, and 1.15 on the same samples and multiplier draws. The larger factors increase coverage and reduce rejection under both nulls and alternatives.

6.4 Matched comparison of self-normalizers

Table 3: Paired adjusted-range minus quadratic Monte Carlo differences

Design	Coverage	Band width	Constancy reject	Linear-path reject
iid regressor, constant path	-0.017 (0.011)	-0.051 (0.003)	-0.010 (0.009)	0.007 (0.007)
iid regressor, linear path	0.007 (0.012)	-0.045 (0.003)	-0.003 (0.010)	0.000 (0.011)
iid regressor, sinusoidal path	-0.023 (0.013)	-0.051 (0.003)	-0.003 (0.021)	0.007 (0.023)
iid overidentified IV, constant path	0.010 (0.011)	-0.058 (0.004)	-0.010 (0.007)	-0.010 (0.007)
iid overidentified IV, sinusoidal path	-0.003 (0.009)	-0.054 (0.004)	-0.023 (0.022)	-0.013 (0.021)
AR(1) regressor, constant path	0.023 (0.012)	-0.037 (0.002)	-0.010 (0.011)	-0.020 (0.012)
AR(1) regressor, sinusoidal path	0.023 (0.013)	-0.036 (0.002)	0.010 (0.021)	-0.023 (0.021)
GARCH errors, constant path	-0.010 (0.013)	-0.032 (0.005)	-0.003 (0.010)	-0.007 (0.012)
AR(1) regressor, localized path	0.007 (0.014)	-0.036 (0.003)	-0.020 (0.013)	0.013 (0.013)

Notes: Each entry is the paired adjusted-range minus quadratic difference with its paired Monte Carlo standard error in parentheses. Both methods are evaluated on every simulated sample using identical multiplier draws.

Table 3 reports adjusted-range minus quadratic differences on all nine common-sample designs. The two procedures use the same local estimates, fitted scores, quartic projection, variance-time index, and Gaussian draws and their respective unscaled multiplier quantiles. The adjusted-range band is narrower in every row: paired mean width differences range from -0.058 to -0.032, with paired standard errors from 0.002 to 0.005. The corresponding coverage and rejection differences have mixed signs and are generally within two paired Monte Carlo standard errors of zero. The evidence therefore identifies a band-width difference in these designs, not a common coverage or power ordering.

Table 4 repeats five matched designs at $T = 1000$. Adjusted-range bands remain narrower in every pair. Null rejection for the constant paths and for the linear-path test under a linear path lies between 0.025 and 0.045 for the adjusted-range procedure, while power against the sinusoidal path rises to 0.985 for constancy and 0.655 for linearity. A projected-score scale benchmark is reported separately in Appendix B. The same appendix reports matched projection and variance-time comparisons that isolate the two construction choices specific to the feasible local score process.

6.5 Fixed-system calibration

The fixed-system experiment matches the dimension of the empirical joint maximum. Ten equations share three autoregressive factors and have four constant coefficients each. Their disturbances contain a common

Table 4: Larger-sample comparison at $T = 1000$

Design	Method	Coverage	Band width	Constancy reject	Linear-path reject
iid regressor, constant path	Adjusted range	0.940	0.894	0.040	0.040
iid regressor, constant path	Quadratic	0.935	0.941	0.045	0.045
iid regressor, linear path	Adjusted range	0.945	0.888	0.080	0.045
iid regressor, linear path	Quadratic	0.915	0.939	0.085	0.060
iid regressor, sinusoidal path	Adjusted range	0.945	0.895	0.985	0.655
iid regressor, sinusoidal path	Quadratic	0.945	0.947	0.980	0.650
iid overidentified IV, constant path	Adjusted range	0.975	1.049	0.030	0.025
iid overidentified IV, constant path	Quadratic	0.950	1.107	0.035	0.035
GARCH errors, constant path	Adjusted range	0.960	1.680	0.040	0.045
GARCH errors, constant path	Quadratic	0.965	1.715	0.045	0.040

Notes: Results use $N_{\mathrm{MC}} = 200$ samples and $B = 499$ common iid Gaussian multiplier draws. The five designs, 25-point grid, bandwidth constant, estimator, score, projection, and unscaled decision rule match the $T = 500$ experiment. The random-number streams are independent of the main experiment.

component, giving pairwise cross-equation correlation 0.4, and remain independent of the factor innovations. The path is evaluated at 71 points from 0.15 to 0.85, so the multiplier maximum contains $10 \times 71 \times 4 = 2{,}840$ coordinates. The experiment uses $T = 500$, $N_{\mathrm{MC}} = 300$, $B = 499$, and common iid Gaussian multipliers across the full system.

Table 5: Finite-system null experiment

Critical-value scalar	Coverage	Band width	Constancy reject	Linear-path reject
1.00	0.907 (0.017)	1.771	0.077 (0.015)	0.060 (0.014)
1.05	0.937 (0.014)	1.860	0.043 (0.012)	0.043 (0.012)
1.10	0.970 (0.010)	1.948	0.027 (0.009)	0.027 (0.009)
1.15	0.977 (0.009)	2.037	0.023 (0.009)	0.023 (0.009)

Notes: The experiment contains 10 equations, 71 grid points, and 40 coefficient paths, so each maximum covers 2,840 coordinates. Errors have cross-equation correlation 0.400. Results use $T = 500$, $N_{\mathrm{MC}} = 300$, $B = 499$ iid common Gaussian multiplier draws, and $b_T = (1.5/\sqrt{12})T^{-1/5}$. Each scalar is applied to the same within-sample critical value for the bands and both path tests. Parentheses contain binomial Monte Carlo standard errors.

Table 5 reports the unscaled rule and the same three scalar adjustments used in the single-equation calibration. It provides a direct finite-sample check of Corollary 4.16 under cross-equation dependence and a maximum comparable in dimension to the application. The unscaled rule has coverage 0.907, constancy rejection 0.077, and linear-path rejection 0.060. With the 1.05 scalar, coverage is 0.937 with Monte Carlo standard error 0.014, while both null rejection frequencies are 0.043. The 1.10 scalar raises coverage to 0.970 and lowers both rejection frequencies to 0.027. The experiment therefore documents a nonnegligible finite-system coverage error for the unscaled rule and shows how scalar changes trade band width against rejection; it does not select a replacement nominal critical value.

6.6 Serial dependence and feasible-score sensitivity

The serial-dependence experiments in Appendix B examine how the first-order dependent-multiplier extension behaves with an estimated score. They identify fitted-score construction as the main source of the remaining finite-sample error. At $T = 500$, the same-window leverage-adjusted score with an exponential dependent multiplier gives coverage 0.880 and 0.760 in the two exactly identified designs with error correlations 0.3 and 0.6. At $T = 1000$, a Bartlett multiplier gives 0.867 and 0.823 in the corresponding cells. Deleting a neighborhood when estimating the score improves these cells, but does not produce uniform performance across exact and overidentified designs. In the $\rho = 0.3$ comparison, the public local blockwise wild bootstrap attains coverage 0.953 and 0.920 for the constant and sinusoidal paths, versus 0.903 and 0.860 for the dependent adjusted-range implementation, and its bands are narrower. Thus the growing-length dependent-multiplier theorem provides first-order validity, while the reported finite samples favor the martingale-score implementation as the baseline and the dependent version as a sensitivity analysis.

7 Empirical application: industry factor exposures

The application asks whether industry factor exposures remain constant, follow a linear trend, or display richer smooth variation. This is a joint path question: a sequence of individually noisy local estimates can look variable even when the complete path remains compatible with stability. We therefore report simultaneous bands and tests over dates, coefficients, and industries, complementing existing work on time-varying factor models (Su and Wang, 2017).

We use monthly Fama–French three-factor series and value-weighted returns for ten United States industry portfolios from the Kenneth French Data Library (Fama and French, 1993; French, 2026). The merged sample runs from July 1963 to May 2026, for $T = 755$. For each industry we estimate smooth paths for the intercept and market, size, and value loadings. The application is an unconditional factor projection; its purpose is to assess path stability, not to assign a causal interpretation to every movement in an exposure.

The portfolio labels follow the Data Library: NoDur is consumer nondurables; Durbl, consumer durables; Manuf, manufacturing; Enrgy, energy; HiTec, business equipment; Telcm, telecommunications; Shops, wholesale, retail, and selected services; Hlth, health care, medical equipment, and drugs; Utils, utilities; and Other, the remaining industries.

Let R_t^i be the return for industry i, R_t^f the one-month risk-free rate, and $R_t^M - R_t^f$ the market excess return. The small-minus-big (SMB) factor measures the size return spread, and the high-minus-low (HML) factor measures the value return spread. Write the common regressor and instrument vector as

$$R_t = Z_t = \{1, R_t^M - R_t^f, SMB_t, HML_t\}'.$$

For each $i = 1, \ldots, 10$, the time-varying three-factor regression is

$$R_t^i - R_t^f = \alpha_i(u_t) + \beta_{M,i}(u_t)(R_t^M - R_t^f) + \beta_{S,i}(u_t)SMB_t + \beta_{H,i}(u_t)HML_t + \varepsilon_{i,t}, \qquad u_t = t/T,$$

Equivalently,

$$\mathbb{E}\Big[Z_t\{R_t^i - R_t^f - R_t'\theta_{0,i}(u_t)\}\Big] = 0, \qquad \theta_{0,i}(u) = (\alpha_i(u), \beta_{M,i}(u), \beta_{S,i}(u), \beta_{H,i}(u))'.$$

This exactly identified moment system is an unconditional projection on observed factors. It supplies the identifying moment used in the application; a causal interpretation requires additional economic assumptions. The iid multiplier is justified only under the additional martingale-score restriction $E(Z_t\varepsilon_{i,t} \mid \mathcal{F}_{t-1}) = 0$. That condition governs the dependence calibration and is stronger than the displayed identifying moment. The industry portfolios are reconstituted annually at the end of June, so changes in the estimated path can reflect changing portfolio composition as well as changing exposures of constituent firms.

The empirical analysis uses the one-step identity-weight local-quadratic estimator and a prespecified 71-point reporting grid from 0.15 to 0.85, corresponding to November 1972 through December 2016. With

$T = 755$, $b_T = (1.5/\sqrt{12})T^{-1/5} = 0.115$ gives a half-window of approximately 87 months, or roughly 175 monthly observations in a complete two-sided window. Every window lies inside the sample, so no reflection or endpoint extrapolation is used. The fitted residual receives the leverage correction, the variance-time index accumulates squared equivalent-kernel weights, and the fitted score is projected on a quartic polynomial. The adjusted range is evaluated with trimming fraction $\tau = 0.10$.

We compute two joint multiplier specifications with $B = 4{,}999$ draws. The baseline specification uses common iid Gaussian multipliers and therefore invokes the additional martingale-score condition. The dependent-score specification maintains the unconditional identifying moment, permits nonzero score autocovariances, and uses the stationary Gaussian multiplier in (4.16) with $L_T = \lfloor 4T^{1/4} \rfloor = 20$, following the correlation-length rate allowed by the dependent-score theorem. The Appendix compares additional correlation lengths to document finite-sample sensitivity. The rate restrictions behind this result are asymptotic rather than finite-sample diagnostics. At $T = 755$, the quantity $L_T(\log T)^2/\sqrt{n_T}$ is approximately 94.2, so the theorem does not provide a quantitative approximation guarantee for this specification at the application sample size. We therefore report it as a sensitivity analysis. For each specification, Table 6 reports the unscaled critical value and a 1.05 sensitivity value; the unscaled value is primary. The same value is used for the bands, constancy statistic, and linear-path statistic, and one multiplier sequence is shared by every industry, date, and coefficient.

Table 6: Joint multiplier calibration across ten industry portfolios

Multiplier	Scalar	Critical value	Band width	Constancy stat.	Reject	Linear stat.	Reject
iid Gaussian	1.00	7.405	2.455	11.059	Yes	9.189	Yes
iid Gaussian	1.05	7.776	2.577	11.059	Yes	9.189	Yes
Dependent Gaussian	1.00	9.819	3.255	11.059	Yes	9.189	No
Dependent Gaussian	1.05	10.310	3.417	11.059	Yes	9.189	No

Notes: Each statistic takes one maximum over ten industries, 71 dates, and four coefficients. The dependent Gaussian multiplier has exponential correlation and length $L_T = 20 = \lfloor 4T^{1/4} \rfloor$. Both methods use $B = 4{,}999$ draws. The same scalar is applied to the bands and both path tests. The 1.00 rows use the unscaled method-specific multiplier quantiles; the 1.05 rows are finite-sample sensitivity calculations.

The maximum constancy statistic is 11.059, attained by the market beta of Other in August 2005. It exceeds both unscaled system-wide critical values, 7.405 under iid multipliers and 9.819 under dependent multipliers. The maximum linear-path statistic is 9.189, attained by the HML beta of Durbl in June 1985. It exceeds the iid critical value but not the dependent-multiplier critical value. The common system-wide conclusion is therefore that at least one factor exposure changes over time; whether the largest departure also rules out a linear path depends on the score-dependence specification. The 1.05 sensitivity rows lead to the same two decisions. At the system-wide level, the constancy statistics for HiTec and Other exceed both unscaled critical values. Corollary 4.16 applies the corresponding score assumption to the jointly stacked system. Common calendar-time multipliers reproduce cross-industry covariance as well as dependence across grid dates and coefficients. Table 7 also shows the effect of replacing ten separate industry maxima by one system-wide maximum.

The difference between the iid and dependent critical values motivates a separate look at the score process. Table 9 summarizes the local fitted scores entering the joint analysis. Across the 710 industry-window sequences for each coefficient, the median absolute first-order autocorrelation lies between 0.066 and 0.074, and 23.8–27.7 percent of the Ljung–Box tests reject at five percent. Median absolute cross-industry correlation lies between 0.180 and 0.227, with much larger values in the upper tail. Because fitted

Table 7: Industry path tests under iid and dependent-multiplier calibration

	Constancy				Linear path			
	iid		Dependent		iid		Dependent	
Industry	Stat.	Crit.	Stat.	Crit.	Stat.	Crit.	Stat.	Crit.
Panel A: Industry-specific maxima (not adjusted across industries)								
NoDur	7.687*	5.482	7.687*	5.847	6.832*	5.482	6.832*	5.847
Durbl	7.127*	6.011	7.127*	6.824	9.189*	6.011	9.189*	6.824
Manuf	3.707	5.543	3.707	5.964	4.562	5.543	4.562	5.964
Enrgy	6.679*	5.559	6.679	7.492	9.031*	5.559	9.031*	7.492
HiTec	10.353*	5.680	10.353*	6.729	8.641*	5.680	8.641*	6.729
Telcm	7.878*	6.168	7.878*	6.547	6.612*	6.168	6.612*	6.547
Shops	7.094*	5.304	7.094*	5.589	6.991*	5.304	6.991*	5.589
Hlth	6.750*	5.755	6.750*	6.713	6.203*	5.755	6.203	6.713
Utils	6.363*	5.582	6.363*	5.826	6.130*	5.582	6.130*	5.826
Other	11.059*	5.953	11.059*	8.265	8.391*	5.953	8.391*	8.265
Panel B: System-wide maximum across ten industries								
NoDur	7.687*	7.405	7.687	9.819	6.832	7.405	6.832	9.819
Durbl	7.127	7.405	7.127	9.819	9.189*	7.405	9.189	9.819
Manuf	3.707	7.405	3.707	9.819	4.562	7.405	4.562	9.819
Enrgy	6.679	7.405	6.679	9.819	9.031*	7.405	9.031	9.819
HiTec	10.353*	7.405	10.353*	9.819	8.641*	7.405	8.641	9.819
Telcm	7.878*	7.405	7.878	9.819	6.612	7.405	6.612	9.819
Shops	7.094	7.405	7.094	9.819	6.991	7.405	6.991	9.819
Hlth	6.750	7.405	6.750	9.819	6.203	7.405	6.203	9.819
Utils	6.363	7.405	6.363	9.819	6.130	7.405	6.130	9.819
Other	11.059*	7.405	11.059*	9.819	8.391*	7.405	8.391	9.819

Notes: Results use the unscaled method-specific multiplier quantiles. Panel A calibrates a maximum over one industry's 71 dates and four coefficients and does not control multiplicity across industries. Panel B uses one common maximum over all ten industries. The iid and dependent columns use the specifications in Table 6. An asterisk on a statistic marks rejection at five percent.

Table 8: Coordinates attaining the joint path-test statistics

Restriction	Industry	Component and date	Statistic
Constancy	Other	Market beta, 2005-08	11.059
Linear path	Durbl	HML beta, 1985-06	9.189

residual moments can reflect both innovation dependence and smooth coefficient variation, these diagnostics do not select a score model. The iid analysis remains the theorem-aligned baseline, and the dependent calibration measures sensitivity to serial dependence. Appendix C reports additional multiplier lengths and a nonoverlapping block comparison.

Table 9: Joint local fitted-score dependence diagnostics

Component	Local sequences	Median $\lvert\widehat{\rho}(1)\rvert$	90th pct. $\lvert\widehat{\rho}(1)\rvert$	Share $p < 0.05$ Ljung–Box(12)	Median cross-industry $\lvert\widehat{\rho}\rvert$	90th pct. cross-industry $\lvert\widehat{\rho}\rvert$
Alpha	710	0.071	0.191	0.277	0.180	0.544
Market beta	710	0.066	0.185	0.238	0.227	0.694
SMB beta	710	0.074	0.182	0.259	0.227	0.675
HML beta	710	0.072	0.188	0.259	0.224	0.641

Notes: The serial diagnostics use the leverage-adjusted fitted influence score within each of 71 local windows for ten industries, giving 710 sequences per coefficient. The Ljung–Box statistic uses 12 monthly lags. Cross-industry correlations are computed for the same coefficient and local window and summarized over the 45 industry pairs. Because fitted residual moments can reflect both innovation dependence and smooth coefficient variation, the table is descriptive and does not identify the innovation-score dependence model. The iid multiplier is the main specification; the growing-length dependent Gaussian multiplier measures sensitivity to serial dependence.

HiTec provides a detailed illustration. Its local market beta averages 1.126 and ranges from 0.877 in September 1986 to 1.639 in September 2003. The SMB loading ranges from -0.201 to 0.498, while the HML point estimate is negative at all 71 grid points and reaches -1.102 in January 1998. The local alpha varies from -0.671 in August 1988 to 1.351 in January 1998.

Table 10: Time-varying factor exposure of the HiTec portfolio

Component	Mean	Min	Max	Global constant fit
alpha	0.165	-0.671	1.351	0.225
Market beta	1.126	0.877	1.639	1.125
SMB beta	0.153	-0.201	0.498	0.178
HML beta	-0.523	-1.102	-0.031	-0.533

Notes: Monthly data are from the Kenneth French Data Library. The dependent variable is the HiTec industry excess return and regressors are a constant, market excess return, SMB, and HML. Returns are in percent per month. The sample is July 1963 through May 2026 with 755 monthly observations. The one-step identity-weight local-quadratic GMM path is evaluated on 71 interior grid points from 0.15 to 0.85. The bandwidth is $b_T = (1.5/\sqrt{12})T^{-1/5}$.

The simultaneous bands distinguish path shape from pointwise sign. The market-beta band lies above zero at 69 of the 71 grid dates under the unscaled iid system-wide critical value and at 59 dates under the dependent-multiplier value. The HML band lies below zero in July 1995 and March 1996 under the iid calibration but not under the dependent calibration; the alpha and SMB bands do not exclude zero at any grid date under either calibration. The positive market exposure is therefore common to both sets of bands over much of the sample, while a simultaneous negative-HML conclusion is sensitive to the score-dependence specification. HiTec's constancy and linear-path statistics are 10.353 and 8.641. The constancy statistic exceeds both system-wide critical values; its margin over the dependent-multiplier value 9.819 is 0.534. The linear-path statistic exceeds the iid value 7.405 but not the dependent value. The largest normalized departure from both fitted restrictions is the market beta in August 2005.

Figure 2 displays the market and HML paths behind the system results. HiTec's market beta rises sharply around the early 2000s, while the market exposures of Utilities and NoDur remain lower over much of the grid. The HML panel separates the negative HiTec loading around the late 1990s and early 2000s from

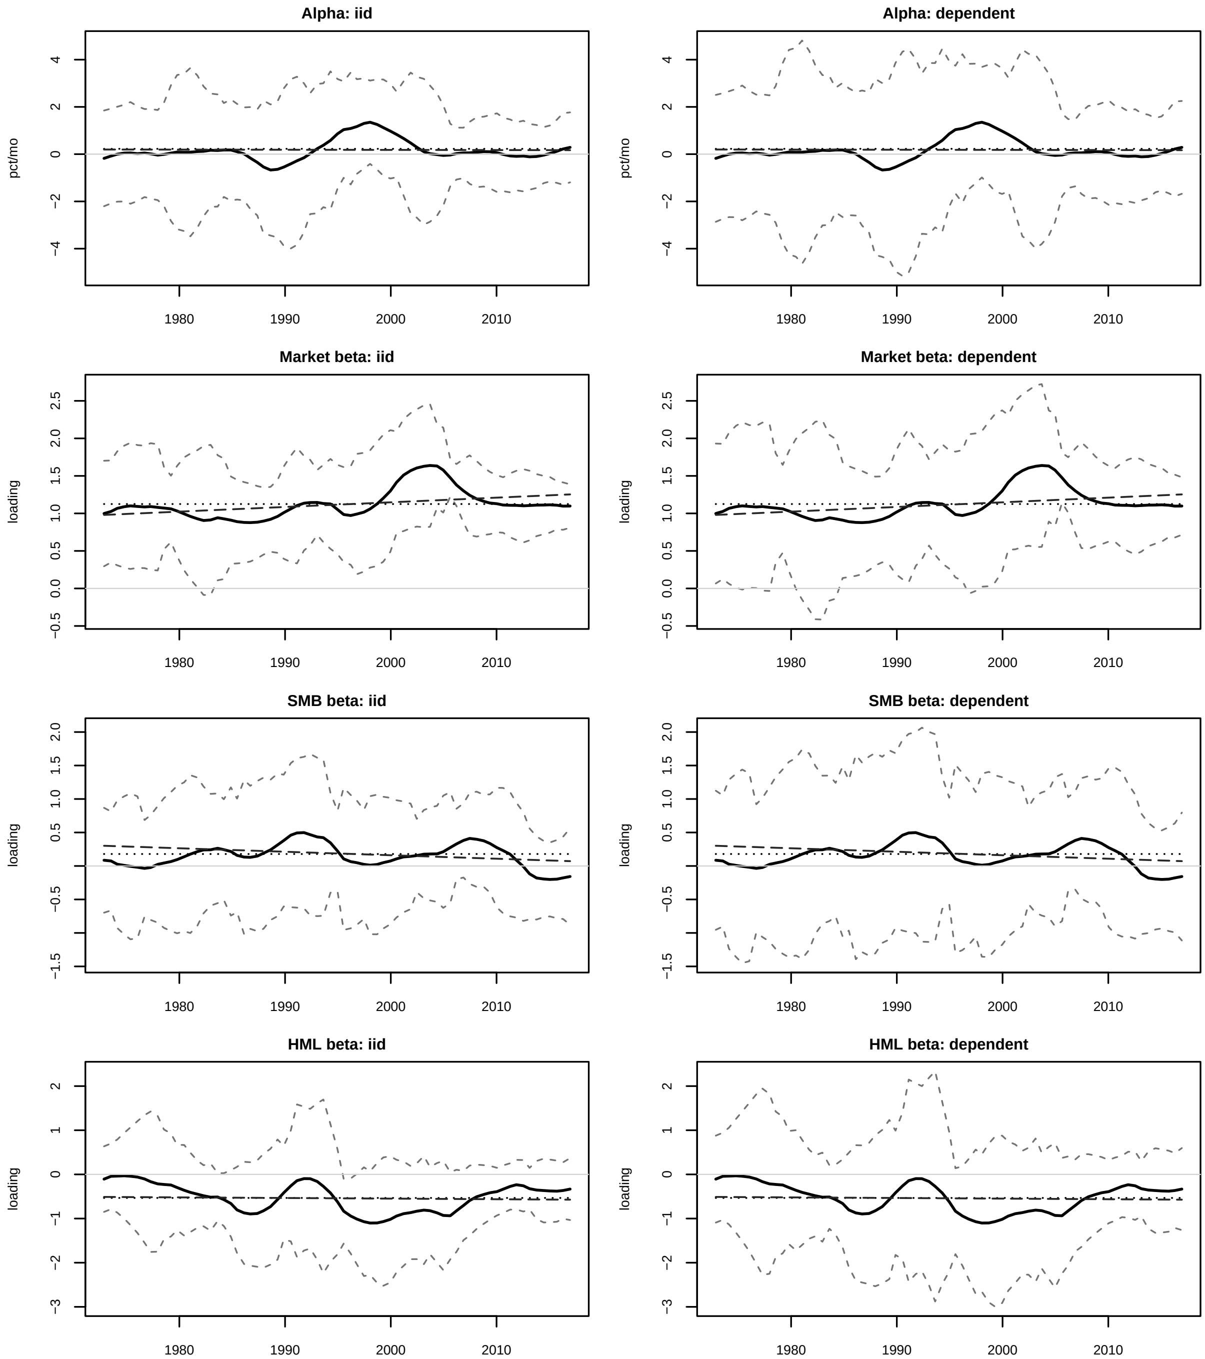

Figure 1: HiTec factor exposures under iid and dependent multiplier specifications

Notes: Solid curves connect the fixed-grid local-GMM estimates, and gray dashed curves connect simultaneous finite-grid band endpoints. The left column uses the unscaled iid system-wide critical value 7.405; the right column uses the unscaled dependent-multiplier value 9.819. Each row has a common vertical scale. Dotted and long-dashed curves show the fitted constant and linear paths. The bands are not continuum confidence bands. The full sample is July 1963 to May 2026; the evaluation grid spans November 1972 to December 2016. Returns are measured in percent per month.

Durbl's positive loading during much of the 1990s and the later positive loading of Other. Enrgy alternates between value and growth exposure rather than following a single monotone trend. Appendix D reports all four coefficient paths for every industry, and Appendix Table C.1 reports the maximizing coefficient and date for each industry-level statistic.

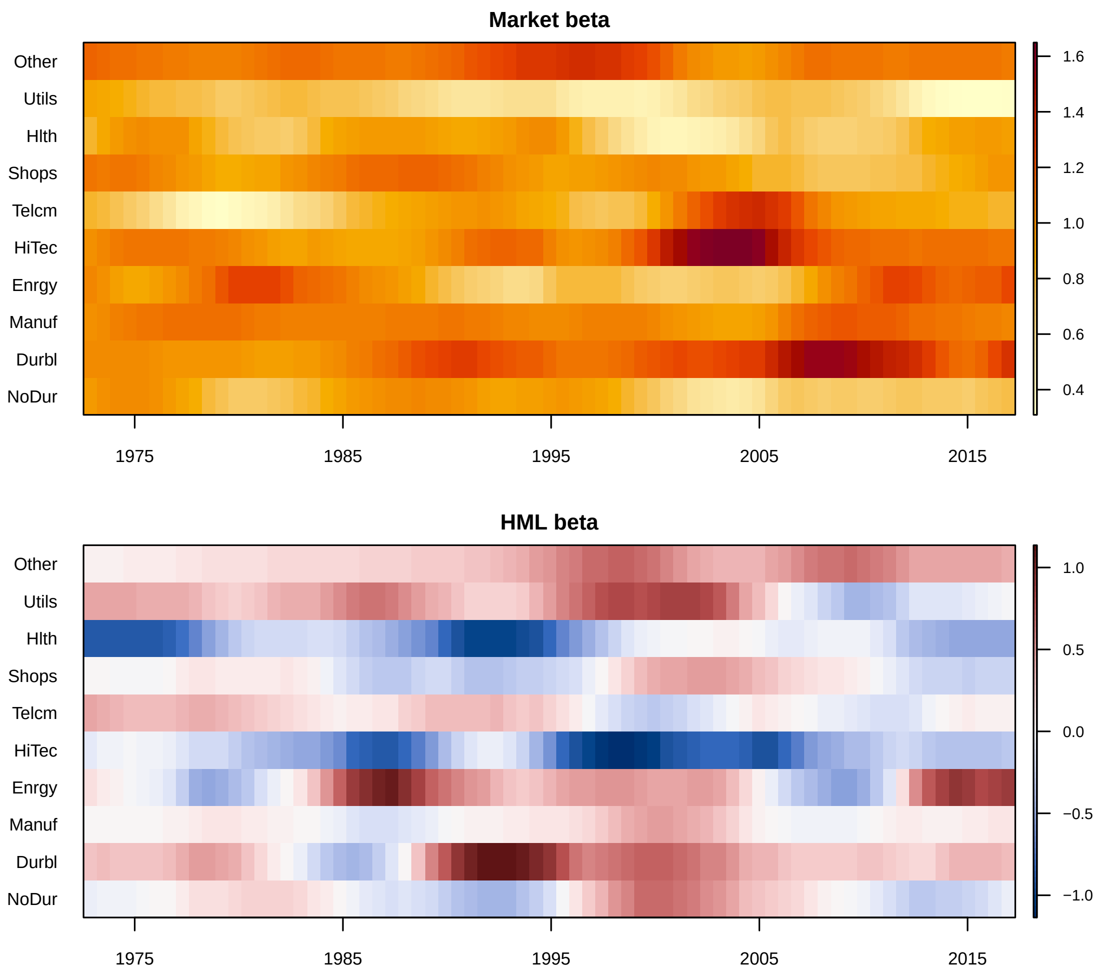

Figure 2: Market and value exposures across ten industry portfolios

Notes: Each row shows one industry's local-quadratic coefficient path on the 71-point evaluation grid. Darker colors in the upper panel indicate a larger market beta. In the lower panel, blue denotes a negative HML loading and red denotes a positive loading; the scale is centered at zero. The paths use the sample, grid, and bandwidth described above. The figure is descriptive; the inferential classifications use the system-wide statistics in Tables 6 and 7.

The baseline results separate a robust finding from a specification-sensitive one: constancy is rejected under both multiplier calibrations, whereas the linear-path conclusion changes with dependence calibration. We next examine whether this distinction persists when the smoothing span, factor model, and sample period change. For HiTec, we vary the bandwidth constant while retaining the $T^{-1/5}$ order, expanding the Fama–French three-factor (FF3) model to the Fama–French five-factor (FF5) model by adding the profitability and investment factors of Fama and French (2015), and splitting the sample at December 1999. Each subsample uses its own rescaled interior grid and the same value $c_b = 1.5$. These robustness calculations use the unscaled iid multiplier rule, so changes across rows reflect the bandwidth, factor model, or sample period rather than a change in multiplier specification. Appendix C reports the additional dependent and block comparisons.

Table 11: HiTec robustness: bandwidth constant, factor model, and sample period

Sample	Model and bandwidth	Critical value	Band width	Constancy	Linear path
Full sample	FF3, $c_b = 1.00$	6.245	2.323	6.392*	6.316*
Full sample	FF3, $c_b = 1.25$	5.955	2.031	8.118*	6.525*
Full sample	FF3, $c_b = 1.50$	5.680	1.807	10.353*	8.641*
Full sample	FF3, $c_b = 1.75$	5.375	1.649	5.909*	5.155
Full sample	FF3, $c_b = 1.90$	5.156	1.547	5.793*	4.923
1963-07 to 1999-12	FF3, $c_b = 1.50$	6.064	2.485	4.010	5.183
2000-01 to 2026-05	FF3, $c_b = 1.50$	6.793	2.048	8.800*	6.765
Full sample	FF5, $c_b = 1.50$	6.301	2.459	7.042*	7.755*

Notes: Every bandwidth equals $(c_b/\sqrt{12})T^{-1/5}$; only c_b varies. The five full-sample FF3 bandwidth rows use the same multiplier draws. Each row uses an interior grid from 0.15 to 0.85, leverage-adjusted fitted scores, the squared local-quadratic equivalent-kernel variance-time index, quartic projection, and $B = 4{,}999$ iid Gaussian multiplier draws. The method-specific multiplier critical value is multiplied by 1.00 in every row. The main specification uses the unscaled value 1.00; other values are sensitivity calculations. An asterisk marks a statistic exceeding its corresponding five-percent critical value.

The full-sample constancy statistic exceeds its critical value for every reported constant from $c_b = 1.00$ through $c_b = 1.90$, with the exponent fixed at $-1/5$; the margin is smallest at $c_b = 1.00$. The linear-path comparison is more sensitive to smoothing: its statistic exceeds the corresponding value for $c_b = 1.00, 1.25$, and 1.50, but not for $c_b = 1.75$ or 1.90. Wider windows reduce average band width and smooth more of the local curvature. The bandwidth exercise therefore locates the time variation over a broad range of window constants, while evidence against a linear path weakens for the two widest windows.

The sample split sharpens that interpretation. Neither statistic exceeds its critical value on the 1963–1999 subsample. On the 2000–2026 subsample, the constancy statistic does, while the linear-path statistic 6.765 remains below its critical value 6.793. In the full-sample FF5 model, both statistics exceed their common critical value. The HiTec full-sample finding therefore survives adding profitability and investment factors, but the split results show that its strength and shape vary across sample periods.

The fitted-score diagnostics may reflect serially correlated innovations, smooth coefficient variation, or both. The supplementary multiplier table therefore varies the dependence span without using the diagnostics to select a single value. It also compares adjusted-range and quadratic self-normalizers under identical multiplier draws. These calculations leave the coefficient paths, observed self-normalizers, and path-test statistics unchanged; only the simulated critical value changes.

8 Conclusion

This paper develops self-normalized finite-grid inference for smooth coefficient paths in linear IV models estimated by local GMM. Its central contribution is a feasible local influence path that can support simultaneous inference. A kernel-weighted quartic projection removes the leading component created by reusing the local fit, while squared local-quadratic equivalent-kernel loadings index the corrected path by accumulated variance. The same local window can then supply the coefficient estimate and a path-based measure of its scale.

The local-quadratic fit is dictated by the bandwidth and the band target. With $b_T \asymp T^{-1/5}$, the local-linear curvature bias is not negligible after normalization, whereas the degree-two remainder is. The method therefore pays the higher local-quadratic variance constant to obtain centered inference for general smooth paths on an interior grid. This argument is specific to the unrestricted bands; constant and linear paths are reproduced by both local linear and local quadratic fits.

The adjusted range is the primary self-normalizer, and a quadratic root-mean-square functional of the same projected path provides a matched benchmark. Common calendar-time multipliers yield simultaneous bands and tests of constant or linear paths on a prespecified finite grid. The IV model is identified by unconditional instrument orthogonality. The martingale-difference restriction is an additional condition for the iid multiplier specialization; it makes the score's long-run covariance equal to its lag-zero covariance. Under the maintained serial-dependence conditions, a growing-length dependent Gaussian multiplier reproduces the local long-run covariance without changing the local influence construction, although its correlation length becomes a finite-sample calibration choice.

With the unscaled multiplier quantile, simultaneous coverage in the nine single-equation designs ranges from 0.930 to 0.973; the GARCH row at the lower end is an outside-assumptions heavy-tail stress test. Adjusted-range bands are narrower than matched quadratic bands in every paired design, while their coverage and rejection differences vary by design. The application-scale system experiment has coverage 0.907 under the unscaled rule, making the finite-system approximation an important numerical consideration. For serially correlated scores, the comparison with oracle scores and a blockwise benchmark identifies feasible-score construction and dependence calibration as the sources of the remaining finite-sample error. These experiments motivate using the martingale-score procedure as the baseline and the dependent calibration as a sensitivity analysis in the current sample sizes.

In the industry portfolio application, the unscaled iid specification and the dependent-score sensitivity analysis both reject system-wide constancy. The linear-path restriction is rejected under iid multipliers but not under the dependent multiplier. HiTec's market exposure is simultaneously positive over most of the formal grid under both specifications, whereas the conclusion for its HML exposure is specification-sensitive. The inferential target is a fixed interior grid of smooth paths. Within that target, the method turns a collection of local IV estimates into joint statements about economically meaningful path restrictions. Extending the feasible influence construction to boundaries, growing grids, discontinuities, weak identification, and long-memory scores remains a direction for further work.

Acknowledgements

This research was supported by the National Natural Science Foundation of China (NSFC) under Grants 71988101 and 72173120.

Data and code availability

The replication archive contains the R source code, simulation seeds, machine-readable numerical summaries, table and figure builders, and TeX sources. The empirical scripts retrieve the public Fama–French files, and a manifest records their source URLs, retrieval dates, and checksums. The files used here were retrieved on July 5, 2026, from archives constructed from the May 2026 Center for Research in Security Prices database; rows carrying the Data Library's missing-value codes are excluded. Unmodified source archives accompany the replication files. The archive provides entry points for rebuilding the reported tables and figures and for rerunning the simulations and empirical analysis.

References

Bai, Y. (2026). Local GMM estimation for nonparametric time-varying coefficient moment condition models. *Journal of Time Series Analysis*, 47(4):774–783.

Bücher, A. and Kojadinovic, I. (2016). A dependent multiplier bootstrap for the sequential empirical copula process under strong mixing. *Bernoulli*, 22(2):927–968.

Chen, B. (2015). Modelling and testing smooth structural changes with endogenous regressors. *Journal of Econometrics*, 185(1):196–215.

Chen, B. and Hong, Y. (2012). Testing for smooth structural changes in time series models via nonparametric regression. *Econometrica*, 80(3):1157–1183.

Dahlhaus, R. (1996). Asymptotic statistical inference for nonstationary processes with evolutionary spectra. In Robinson, P. M. and Rosenblatt, M., editors, *Athens Conference on Applied Probability and Time Series Analysis, Volume II: Time Series Analysis*, pages 145–159. Springer, New York.

Dahlhaus, R. (2012). Locally stationary processes. In Subba Rao, T., Subba Rao, S., and Rao, C. R., editors, *Time Series Analysis: Methods and Applications*, volume 30 of *Handbook of Statistics*, pages 351–413. Elsevier, Amsterdam.

Fama, E. F. and French, K. R. (1993). Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics*, 33(1):3–56.

Fama, E. F. and French, K. R. (2015). A five-factor asset pricing model. *Journal of Financial Economics*, 116(1):1–22.

Fan, J. and Gijbels, I. (1996). *Local Polynomial Modelling and Its Applications.* Chapman and Hall, London.

French, K. R. (2026). Kenneth R. French data library. https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html. Accessed July 5, 2026.

Friedrich, M. and Lin, Y. (2024). Sieve bootstrap inference for linear time-varying coefficient models. *Journal of Econometrics*, 239(1):105345.

Giraitis, L., Kapetanios, G., and Marcellino, M. (2021). Time-varying instrumental variable estimation. *Journal of Econometrics*, 224(2):394–415.

Hall, P. and Heyde, C. C. (1980). *Martingale Limit Theory and Its Application.* Academic Press, New York.

Hansen, L. P. (1982). Large sample properties of generalized method of moments estimators. *Econometrica*, 50(4):1029–1054.

Hong, Y., Linton, O., McCabe, B., Sun, J., and Wang, S. (2024). Kolmogorov–Smirnov type testing for structural breaks: A new adjusted-range based self-normalization approach. *Journal of Econometrics*, 238(2):105603.

Karmakar, S., Richter, S., and Wu, W. B. (2022). Simultaneous inference for time-varying models. *Journal of Econometrics*, 227(2):408–428.

Kim, S., Zhao, Z., and Shao, X. (2015). Nonparametric functional central limit theorem for time series regression with application to self-normalized confidence interval. *Journal of Multivariate Analysis*, 133:277–290.

Koo, B. and Linton, O. (2012). Estimation of semiparametric locally stationary diffusion models. *Journal of Econometrics*, 170(1):210–233.

Kristensen, D. and Lee, Y. J. (2026). Local polynomial estimation of time-varying parameters in nonlinear models. *Econometric Theory*, pages 1–47. First View.

Li, H., Zhou, J., and Hong, Y. (2024). Estimating and testing for smooth structural changes in moment condition models. *Journal of Econometrics*, 246:105896.

Lin, Y., Song, M., and van der Sluis, B. (2025). Bootstrap inference for linear time-varying coefficient models in locally stationary time series. *Journal of Computational and Graphical Statistics*, 34(2):654–667.

Lobato, I. N. (2001). Testing that a dependent process is uncorrelated. *Journal of the American Statistical Association*, 96(455):1066–1076.

MacKinnon, J. G. and White, H. (1985). Some heteroskedasticity-consistent covariance matrix estimators with improved finite sample properties. *Journal of Econometrics*, 29(3):305–325.

Shao, X. (2010a). The dependent wild bootstrap. *Journal of the American Statistical Association*, 105(489):218–235.

Shao, X. (2010b). A self-normalized approach to confidence interval construction in time series. *Journal of the Royal Statistical Society: Series B*, 72(3):343–366.

Shao, X. (2015). Self-normalization for time series: A review of recent developments. *Journal of the American Statistical Association*, 110(512):1797–1817.

Shao, X. and Zhang, X. (2010). Testing for change points in time series. *Journal of the American Statistical Association*, 105(491):1228–1240.

Su, L. and Wang, X. (2017). On time-varying factor models: Estimation and testing. *Journal of Econometrics*, 198(1):84–101.

Sun, J., Hong, Y., Linton, O., and Zhao, X. (2022). Adjusted-range self-normalized confidence interval construction for censored dependent data. *Economics Letters*, 220:110873.

Wu, W. B. (2005). Nonlinear system theory: Another look at dependence. *Proceedings of the National Academy of Sciences*, 102(40):14150–14154.

Zhang, T. and Wu, W. B. (2012). Inference of time-varying regression models. *The Annals of Statistics*, 40(3):1376–1402.

Zhou, Z. and Wu, W. B. (2010). Simultaneous inference of linear models with time varying coefficients. *Journal of the Royal Statistical Society: Series B*, 72(4):513–531.

A Proofs for the finite-grid local-quadratic theory

Throughout, C denotes a finite constant that may change from line to line. Because (M, p, q) are fixed, coordinatewise bounds are automatically uniform over the evaluation grid and parameter coordinates. Put $n = n_T = Tb_T$ and suppress the subscript T on b_T when no confusion can arise. For a fixed grid point, write $z_t = z_{t\ell}$, $w_t = w_{t\ell}$, $D_\ell = D(u_\ell)$, $H_{D\ell} = H_D(u_\ell)$, and $\mathcal{I}_\ell = \{t : |z_{t\ell}| < 1\}$. Since the grid is fixed and interior, $\mathcal{I}_\ell$ is eventually a two-sided calendar interval and, for $\ell \neq \ell'$, $\mathcal{I}_\ell \cap \mathcal{I}_{\ell'} = \varnothing$ for all sufficiently large T.

A.1 Why degree two is used at the maintained bandwidth

The following calculation isolates the smoothing-bias issue from the later self-normalization and multiplier arguments. Let $p_1(z) = (1, z)'$, $Q_1 = \int K(z)p_1(z)p_1(z)'\,dz$, and let $\widehat{\theta}_1(u_\ell)$ be the degree-one version of (3.2).

Lemma A.1 (Local-linear bias at the inferential scale)**.** *Under Assumption 4.1, with the local criterion in* (3.2) *evaluated at $r = 1$, uniformly over the fixed interior grid,*

$$\widehat{\theta}_1(u_\ell) - \theta_0(u_\ell) = \frac{1}{n_T}\sum_{t=1}^{T} K(z_{t\ell})H_D(u_\ell)U_{t,T} + \frac{\mu_2}{2}b_T^2\theta_0''(u_\ell) + o_p(n_T^{-1/2} + b_T^2), \tag{A.1}$$

where $\mu_2 = \int z^2K(z)\,dz$. Consequently,

$$\sqrt{n_T}\,\frac{\mu_2}{2}b_T^2\theta_0''(u_\ell) = \frac{\mu_2}{2}\left(\frac{c_b}{\sqrt{12}}\right)^{5/2}\theta_0''(u_\ell),$$

which is generally nonzero. If the coefficient path is exactly constant or linear on $[0,1]$, the local-linear approximation remainder is identically zero, and so is this smoothing-bias term.

Proof. Fix $u = u_\ell$ and suppress ℓ. Symmetry of the kernel gives

$$Q_1 = \begin{pmatrix} 1 & 0 \\ 0 & \mu_2 \end{pmatrix}, \qquad e_{0,2}'Q_1^{-1} = (1, 0).$$

The local-linear Taylor remainder is

$$\theta_0(u + bz) - \theta_0(u) - bz\theta_0'(u) = \tfrac{1}{2}b^2z^2\theta_0''(u) + o(b^2)$$

uniformly for $|z| \le 1$. The frozen population augmented Jacobian is $D(u) \otimes Q_1$, whose left inverse is $H_D(u) \otimes Q_1^{-1}$. The order-b^2 population residual moment is

$$\tfrac{1}{2}b^2\{D(u)\theta_0''(u)\} \otimes \int K(z)z^2p_1(z)\,dz.$$

Because $e_{0,2}'Q_1^{-1}\int K(z)z^2p_1(z)\,dz = \mu_2$, selecting the intercept produces the second term in (A.1). Smooth variation in D multiplies this quadratic remainder by at least one additional power of b and does not alter the leading coefficient. The weighted kernel laws and left-inverse argument of Lemmas A.3 and A.4, applied with p_1 in place of p_2, supply the stochastic term and the uniform remainder; $\sqrt{n_T}\rho_T \to 0$ removes the ridge at the displayed scale. Finally, substituting (3.3) yields the stated normalized bias. Constant and linear paths are reproduced exactly by the degree-one basis, which proves the last claim. ∎

Proof of Lemma 4.2. We give the argument for scalar coordinates; fixed dimensionality then gives the vector and matrix statements. Uniform L^{2r} bounds and Hölder's inequality yield

$$\begin{aligned}\|Z_t(u)\varepsilon_t(u) - Z_t^*(u)\varepsilon_t^*(u)\|_r \le \|Z_t(u) - Z_t^*(u)\|_{2r}\|\varepsilon_t(u)\|_{2r} \\ + \|Z_t^*(u)\|_{2r}\|\varepsilon_t(u) - \varepsilon_t^*(u)\|_{2r},\end{aligned}$$

and the same calculation, with R in place of ε, applies to $A_t(u) = Z_t(u)R_t(u)'$. Consequently,

$$\begin{aligned}\delta_r^U(k) &\le C\{\delta_{2r}^Z(k) + \delta_{2r}^{\varepsilon}(k)\}, \\ \delta_r^A(k) &\le C\{\delta_{2r}^Z(k) + \delta_{2r}^R(k)\}.\end{aligned} \tag{A.2}$$

The same product inequalities, with the coupled process replaced by the triangular array or by the process evaluated at another rescaled time, give the asserted $O(T^{-1})$ approximation and Lipschitz bounds for U and A.

For $C_t(u) = U_t(u)U_t(u)'$, Hölder's inequality gives

$$\|C_t(u) - C_t^*(u)\|_{r/2} \le \{\|U_t(u)\|_r + \|U_t^*(u)\|_r\}\|U_t(u) - U_t^*(u)\|_r.$$

It also gives the corresponding triangular-array and time-smoothness bounds. Conditional Jensen's inequality supplies the $L^{r/2}$ moment bound for $\Gamma_t(u) = E\{C_t(u) \mid \mathcal{F}_{t-1}\}$ and, because the conditioning sigma-field is common across u,

$$\|\Gamma_t(u) - \Gamma_t(v)\|_{r/2} \le \|C_t(u) - C_t(v)\|_{r/2}.$$

The same contraction proves the triangular-array approximation.

It remains to verify the coupling bound for Γ. For $k = 0$ the conditional covariance is measurable with respect to the pre-0 past, so its coupled difference is zero. For $k \ge 1$, a causal version of $\Gamma_k(u)$ is obtained by integrating the current innovation in $C_k(u)$ while holding $(\ldots, \eta_{k-2}, \eta_{k-1})$ fixed. Construct its coupled version by replacing η_0 in that past and use the same integration variable for the current innovation. Conditional Jensen's inequality on this joint coupling gives

$$\delta_{r/2}^{\Gamma}(k) \le \delta_{r/2}^{C}(k).$$

Combining this inequality with (A.2) and the primitive weighted summability in (4.2) proves the four physical-dependence series. The moment, approximation, and smoothness claims follow from the same inequalities, completing the proof. □

A.2 A weighted physical-dependence inequality

For an integrable random variable Y, let $\mathcal{P}_j Y = E(Y \mid \mathcal{F}_j) - E(Y \mid \mathcal{F}_{j-1})$.

Lemma A.2 (Weighted maximal physical-dependence bound)**.** *Let $X_t(u) = G(u, \ldots, \eta_{t-1}, \eta_t)$ be scalar and centered for every u, let $s \ge 2$, and suppose $\Theta_s^X = \sum_{k\ge0} \delta_s^X(k) < \infty$. For deterministic coefficients $c_{t,T}$,*

$$\left\| \max_{m\le T} \left| \sum_{t=1}^{m} c_{t,T} X_t(u_t) \right| \right\|_s \le C_s \Theta_s^X \left(\sum_{t=1}^{T} c_{t,T}^2 \right)^{1/2}. \tag{A.3}$$

If $X_{t,T}$ satisfies $\max_{t\le T}\|X_{t,T}-X_t(u_t)\|_s \le CT^{-1}$, then the left side with $X_{t,T}-EX_{t,T}$ is bounded by the right side of (A.3) plus $CT^{-1}\sum_t |c_{t,T}|$. The result holds coordinatewise for fixed-dimensional vectors and matrices.

Proof. The causal representation and martingale convergence give the projection expansion

$$X_t(u_t)=\sum_{k=0}^{\infty}\mathcal{P}_{t-k}X_t(u_t) \quad \text{in } L^s.$$

The coupling/projection inequality underlying the physical-dependence measure of Wu (2005), specifically the projection bound in his Theorem 1(ii), gives

$$\|\mathcal{P}_{t-k}X_t(u_t)\|_s \le \|X_t(u_t)-X_t^{*(t-k)}(u_t)\|_s \le \delta_s^X(k).$$

For fixed k, $\{c_{t,T}\mathcal{P}_{t-k}X_t(u_t),\mathcal{F}_{t-k}\}$ is a martingale-difference sequence in t. Doob's maximal inequality and the Burkholder inequality therefore imply

$$\begin{aligned}
&\left\|\max_{m\le T}\left|\sum_{t=1}^{m} c_{t,T}\mathcal{P}_{t-k}X_t(u_t)\right|\right\|_s \\
&\quad \le C_s\left\{\sum_t c_{t,T}^2\|\mathcal{P}_{t-k}X_t(u_t)\|_s^2\right\}^{1/2} \le C_s\delta_s^X(k)\|c_T\|_2.
\end{aligned}$$

Minkowski's inequality and summation over k prove (A.3).

For the triangular array, add and subtract $X_t(u_t)-EX_t(u_t)$. The maximum absolute approximation remainder is no larger than

$$\sum_t |c_{t,T}|\big\{|X_{t,T}-X_t(u_t)|+E|X_{t,T}-X_t(u_t)|\big\}.$$

Minkowski's inequality gives the stated $CT^{-1}\|c_T\|_1$ bound. Applying the scalar result to each coordinate proves the fixed-dimensional version. □

A.3 Kernel sums, weighted laws of large numbers, and local design

Let $\mathscr{H}$ denote any fixed finite collection of uniformly bounded functions on $[-1,1]$ that are piecewise Lipschitz with a uniformly bounded number of jumps. The class used below contains products of K, K^2, ℓ_2, monomials through degree eight, and one prefix indicator $\mathbf{1}\{z\le x\}$, uniformly in $x\in[-1,1]$.

Lemma A.3 (Kernel sums and weighted local laws)**.** *For every $h\in\mathscr{H}$, uniformly over the fixed grid and, when present, the prefix endpoint x,*

$$T^{-1}\sum_{t=1}^{T} w_{t\ell}h(z_{t\ell})=\int_{-1}^{1}K(z)h(z)\,dz+O(n^{-1}). \tag{A.4}$$

Moreover,

$$N_\ell=2n+O(1), \qquad \max_t w_{t\ell}=O(b^{-1}), \qquad \sum_t w_{t\ell}^2=O(T/b).$$

If X is any coordinate of A, C, or Γ, with the moment exponent specified in Lemma 4.2, then

$$\max_{m\le T}\left|T^{-1}\sum_{t=1}^{m} w_{t\ell}h(z_{t\ell})\{X_{t,T}-EX_{t,T}\}\right| = O_p(n^{-1/2}+T^{-1}), \tag{A.5}$$

$$\max_{m\le T}\left|n^{-1}\sum_{t=1}^{m} h(z_{t\ell})\mathbf{1}\{|z_{t\ell}|<1\}\{X_{t,T}-EX_{t,T}\}\right| = O_p(n^{-1/2}+T^{-1}). \tag{A.6}$$

The same bounds hold for U without centering because $EU_{t,T}=0$. If $\mu_X(u)=EX_t(u)$ is Lipschitz, then

$$T^{-1}\sum_t w_{t\ell}h(z_{t\ell})EX_{t,T} = \int_{-1}^{1} K(z)h(z)\mu_X(u_\ell+bz)\,dz + O(n^{-1}+T^{-1}). \tag{A.7}$$

Proof. The local z-grid has mesh $(Tb)^{-1}=n^{-1}$. Since $T^{-1}w_{t\ell}=n^{-1}K(z_{t\ell})$, the first assertion is the Riemann-sum bound for a compactly supported function of bounded variation. A single prefix indicator adds one jump and leaves the $O(n^{-1})$ bound uniform in its location. The support contains $2Tb+O(1)$ observations, and boundedness of K gives the remaining deterministic weight orders.

For (A.5), apply Lemma A.2 with

$$c_{t,T}=T^{-1}w_{t\ell}h(z_{t\ell})=n^{-1}K(z_{t\ell})h(z_{t\ell}).$$

Then $\sum_t c_{t,T}^2 = O(n^{-1})$ and $\sum_t |c_{t,T}| = O(1)$. Equation (A.6) is the same calculation with $c_{t,T} = n^{-1}h(z_{t\ell})\mathbf{1}\{|z_{t\ell}|<1\}$. Finally, the $O(T^{-1})$ approximation of the means, Lipschitz continuity of μ_X, and (A.4) prove (A.7). □

Lemma A.4 (Augmented design and left inverse)**.** *Let*

$$G_\ell^0 = D_\ell\otimes Q_2, \qquad H_\ell^0=\{(G_\ell^0)'G_\ell^0\}^{-1}(G_\ell^0)'.$$

Then, uniformly over the fixed grid,

$$\widehat{G}_\ell - G_\ell^0 = O_p(b+n^{-1/2}+T^{-1}), \tag{A.8}$$

$$\inf_{\ell\le M}\sigma_{\min}(\widehat{G}_\ell) > c > 0 \quad \textit{with probability tending to one}, \tag{A.9}$$

$$\widehat{H}_\ell - H_\ell^0 = O_p(b+n^{-1/2}+T^{-1}+\rho_T). \tag{A.10}$$

For the identity augmented weight,

$$H_\ell^0 = H_D(u_\ell)\otimes Q_2^{-1}. \tag{A.11}$$

For a fixed separable weight $W_Z\otimes W_P$, the population weighted left inverse is $H_{D,W_Z}(u_\ell)\otimes Q_2^{-1}$.

Proof. The mixed-product rule gives

$$V_{t\ell}X_{t\ell}' = (Z_{t,T}R_{t,T}')\otimes\{p_2(z_{t\ell})p_2(z_{t\ell})'\} = A_{t,T}\otimes\{p_2(z_{t\ell})p_2(z_{t\ell})'\}.$$

Equations (A.5), (A.7), and the Lipschitz property of D therefore yield

$$\widehat{G}_\ell = D_\ell\otimes Q_2 + O_p(b+n^{-1/2}+T^{-1}),$$

which proves (A.8). Assumption 4.1(v)–(vi) implies that every G_ℓ^0 has full column rank and that the minimum singular value over the finite grid is positive. Weyl's inequality proves (A.9).

On a neighborhood of the finite collection $\{G_\ell^0\}$, the map

$$(G, \rho) \longmapsto (G'G + \rho I)^{-1} G'$$

is continuously differentiable at $\rho = 0$. The inverse perturbation identity and (A.8) therefore prove (A.10). Finally,

$$(G_\ell^0)' G_\ell^0 = (D_\ell' D_\ell) \otimes Q_2^2,$$

so

$$\{(G_\ell^0)' G_\ell^0\}^{-1} (G_\ell^0)' = \{(D_\ell' D_\ell)^{-1} D_\ell'\} \otimes Q_2^{-1},$$

which is (A.11). With $W_Z \otimes W_P$ in the augmented criterion, the two factors are

$$\{D_\ell' W_Z D_\ell\}^{-1} D_\ell' W_Z \quad \text{and} \quad \{Q_2 W_P Q_2\}^{-1} Q_2 W_P = Q_2^{-1}.$$

This proves the separable-weight extension. □

A.4 Proof of the local-quadratic influence expansion

Proof of Proposition 4.3. Fix ℓ. Define

$$r_{t\ell} = \theta_0(u_t) - \theta_0(u_\ell) - b z_{t\ell} \theta_0'(u_\ell) - \tfrac{1}{2} b^2 z_{t\ell}^2 \theta_0''(u_\ell).$$

Then

$$y_{t,T} = X_{t\ell}' \beta_\ell^0 + \varepsilon_{t,T} + R_{t,T}' r_{t\ell}.$$

Set

$$\overline{U}_{V,\ell} = T^{-1} \sum_t w_{t\ell} V_{t\ell} \varepsilon_{t,T}, \qquad \overline{r}_{V,\ell} = T^{-1} \sum_t w_{t\ell} V_{t\ell} R_{t,T}' r_{t\ell}.$$

The exact regularized normal equations imply

$$\widehat{\beta}_\ell - \beta_\ell^0 = \widehat{H}_\ell (\overline{U}_{V,\ell} + \overline{r}_{V,\ell}) - \rho_T \{\widehat{G}_\ell' \widehat{G}_\ell + \rho_T I\}^{-1} \beta_\ell^0. \tag{A.12}$$

Indeed, $\widehat{H}_\ell \widehat{G}_\ell = I - \rho_T (\widehat{G}_\ell' \widehat{G}_\ell + \rho_T I)^{-1}$.

Because $V_{t\ell} \varepsilon_{t,T} = U_{t,T} \otimes p_2(z_{t\ell})$, Lemma A.3 gives

$$\|\overline{U}_{V,\ell}\| = O_p(n^{-1/2} + T^{-1}).$$

Taylor's theorem and bounded kernel support give $\sup_{t \in \mathcal{I}_\ell} \|r_{t\ell}\| = O(b^3)$. Since

$$V_{t\ell} R_{t,T}' r_{t\ell} = \{A_{t,T} r_{t\ell}\} \otimes p_2(z_{t\ell}),$$

the weighted laws and moment bounds imply $\|\overline{r}_{V,\ell}\| = O_p(b^3)$. Lemma A.4, boundedness of β_ℓ^0, and (A.12) now yield

$$\|\widehat{\beta}_\ell - \beta_\ell^0\| = O_p(n^{-1/2} + b^3 + \rho_T),$$

which proves (4.8).

We next select the intercept and sharpen the bias. The exact Kronecker factorization gives

$$E_0 H_\ell^0 V_{t\ell}\varepsilon_{t,T} = \ell_2(z_{t\ell}) H_{D\ell} U_{t,T},$$

and therefore

$$E_0 H_\ell^0 \overline{U}_{V,\ell} = \sum_t \phi_{t\ell}.$$

By (A.10),

$$\begin{aligned}\|E_0(\widehat{H}_\ell - H_\ell^0)\overline{U}_{V,\ell}\| &= O_p\{(b + n^{-1/2} + T^{-1} + \rho_T)(n^{-1/2} + T^{-1})\}\\ &= o_p(n^{-1/2}).\end{aligned}$$

The selected regularization term in (A.12) is $O_p(\rho_T) = o_p(n^{-1/2})$.

Four-times continuous differentiability gives, uniformly for $|z| \le 1$,

$$r_{t\ell} = \tfrac{1}{6} b^3 z_{t\ell}^3 \theta_0^{(3)}(u_\ell) + \tfrac{1}{24} b^4 z_{t\ell}^4 \theta_0^{(4)}(u_\ell) + o(b^4). \tag{A.13}$$

Using $H_{D\ell} D_\ell = I_p$,

$$E_0 H_\ell^0 \overline{r}_{V,\ell} = T^{-1} \sum_t w_{t\ell} \ell_2(z_{t\ell}) H_{D\ell} A_{t,T} r_{t\ell}.$$

Decompose

$$A_{t,T} = \{A_{t,T} - EA_{t,T}\} + \{EA_{t,T} - D(u_t)\} + \{D(u_t) - D_\ell\} + D_\ell.$$

The first component is $O_p(b^3 n^{-1/2} + b^3 T^{-1})$ by (A.5); the second is $O(b^3 T^{-1})$ by the primitive triangular approximation. For the frozen component D_ℓ, symmetry of K makes Q_2 block diagonal between odd and even monomials, so ℓ_2 is even and

$$\int_{-1}^{1} K(z)\ell_2(z) z^3 \, dz = 0.$$

The corresponding discrete cubic sum is $O(b^3 n^{-1})$ by (A.4). The frozen fourth-order term is $O(b^4)$. Finally, $D(u_\ell + bz) - D_\ell = bzD'(u_\ell) + O(b^2)$, so local variation in D multiplies the cubic remainder by one additional b and contributes $O(b^4)$. The uniform Taylor remainder contributes $o(b^4)$. Hence

$$E_0 H_\ell^0 \overline{r}_{V,\ell} = O_p(b^4 + b^3 n^{-1/2} + b^3 n^{-1} + b^3 T^{-1}). \tag{A.14}$$

Replacing H_ℓ^0 by $\widehat{H}_\ell$ adds at most

$$O_p\{(b + n^{-1/2} + T^{-1} + \rho_T) b^3\} = O_p(b^4 + b^3 n^{-1/2} + \rho_T b^3) + o(b^4).$$

Combining these bounds proves the first equality in (4.9). Under $b \asymp T^{-1/5}$,

$$n^{-1/2} = T^{-2/5}, \qquad \sqrt{n}\, b^4 = T^{-2/5}, \qquad b^3 n^{-1/2} = T^{-1}, \qquad b^3 n^{-1} = T^{-7/5},$$

and $\sqrt{n}\rho_T \to 0$. Every remainder is therefore $o_p(n^{-1/2})$, proving the proposition. □

A.5 Calendar-ordered sample functional central limit theorem

For a step process indexed by the local z-mesh, let a tilde denote its polygonal interpolation between adjacent mesh points, held constant outside the first and last local observations.

Lemma A.5 (Calendar-ordered finite-grid functional central limit theorem)**.** *Define*

$$\mathbb{G}_{T\ell}(z) = n^{-1/2} \sum_{t=1}^{T} k_2(z_{t\ell}) \mathbf{1}\{z_{t\ell} \le z\} H_{D\ell} U_{t,T}, \qquad z \in [-1, 1].$$

Jointly over the fixed grid, $\{\widetilde{\mathbb{G}}_{T\ell}(\cdot)\}_{\ell \le M}$ converges in $\{C[-1,1]^p\}^M$ to $\{G_\ell(\cdot)\}_{\ell \le M}$. The same joint convergence holds after adjoining the finitely many polynomially weighted processes obtained by replacing $k_2(z)$ with $q_4(z)k_2(z)$, coordinatewise.

Proof. We first establish finite-dimensional convergence without reordering any observations. Fix finitely many z values and a scalar Cramér–Wold coefficient for every grid point, coordinate, and z value. The resulting linear combination can be written

$$S_T = \sum_{t=1}^{T} d_{t,T}, \qquad d_{t,T} = n^{-1/2} c_{t,T}' U_{t,T},$$

where $c_{t,T}$ is deterministic, uniformly bounded, supported on the union of the local calendar windows, and contains all local prefix indicators. Because $E(U_{t,T} \mid \mathcal{F}_{t-1,T}) = 0$, $\{d_{t,T}, \mathcal{F}_{t,T}\}$ is a calendar-ordered martingale-difference array.

Its predictable quadratic variation is

$$\sum_{t=1}^{T} E(d_{t,T}^2 \mid \mathcal{F}_{t-1,T}) = n^{-1} \sum_{t=1}^{T} c_{t,T}' \Gamma_{t,T} c_{t,T}. \tag{A.15}$$

The functions multiplying each coordinate of $\Gamma_{t,T}$ belong to the piecewise-Lipschitz class used in Lemma A.3. Applying (A.6) to $\Gamma - E\Gamma$, then using (A.7) and Riemann sums, shows that (A.15) converges in probability to the covariance obtained by integrating $k_2(v)^2 \Sigma_\ell$ over the intersections of the selected prefixes. Cross-products between distinct grid windows vanish exactly for all sufficiently large T. At finite T, any overlap remains inside the single calendar-ordered sum (A.15); no variance-time reordering has been used.

For every $\epsilon > 0$,

$$\begin{aligned} E \sum_t E\{d_{t,T}^2 \mathbf{1}(|d_{t,T}| > \epsilon) \mid \mathcal{F}_{t-1,T}\} &= \sum_t E\{d_{t,T}^2 \mathbf{1}(|d_{t,T}| > \epsilon)\} \\ &\le \epsilon^{-(r-2)} \sum_t E|d_{t,T}|^r \le C\epsilon^{-(r-2)} n^{1-r/2} \longrightarrow 0. \end{aligned}$$

Thus the conditional Lindeberg sum converges to zero in L^1, hence in probability. For a Cramér–Wold combination with positive limiting variance, Corollary 3.1 of Hall and Heyde (1980, pp. 58–59) applies: its two substantive conditions are convergence of predictable quadratic variation and conditional Lindeberg. A combination with zero limiting variance converges to zero by its second moment. This proves joint finite-dimensional convergence.

We next prove tightness directly. Let $I(z_1, z_2]$ be the calendar indices whose local coordinates lie in $(z_1, z_2]$. A fourth-moment Burkholder inequality for martingale differences, the boundedness of k_2, and the uniform fourth moment of $U_{t,T}$ give

$$
\begin{aligned}
E\|\mathbb{G}_{T\ell}(z_2) - \mathbb{G}_{T\ell}(z_1)\|^4 &\le C\left\{\frac{\#I(z_1, z_2]}{n}\right\}^2 + C\frac{\#I(z_1, z_2]}{n^2} \\
&\le C(|z_2 - z_1| + n^{-1})^2. \qquad \text{(A.16)}
\end{aligned}
$$

Partition $[-1, 1]$ into intervals of length between δ and 2δ. Doob's inequality applied within each block and (A.16) imply that the probability that any block has oscillation exceeding η is bounded by $C\eta^{-4}\delta^{-1}(\delta+n^{-1})^2$. Taking $T \to \infty$ and then $\delta \downarrow 0$ proves asymptotic equicontinuity of the interpolated process. Finally,

$$
n^{-1/2} \max_{t\le T} \|U_{t,T}\| = O_p(n^{-1/2}T^{1/r}) = o_p(1),
$$

so the interpolation error is $o_p(1)$. The argument is unchanged after multiplication by the bounded coordinates of q_4. Tightness and the finite-dimensional limit prove the lemma. ☐

A.6 Projection matrices, variance time, and the projected Gaussian path

Lemma A.6 (Sample projection and variance-time index)**.** *Uniformly over the fixed grid,*

$$
\begin{aligned}
\|\widehat{Q}_{4\ell} - Q_4\| &= O(n^{-1}), && \text{(A.17)} \\
\|\widehat{Q}_{2\ell}^{\text{eq}} - Q_2\| &= O(n^{-1} + \rho_T^Q), && \text{(A.18)} \\
\sup_{m\le N_\ell} |\widehat{\mathcal{V}}_\ell(m) - \mathcal{V}(z_{t_{\ell,m},\ell})| &= O(n^{-1} + \rho_T^Q), && \text{(A.19)} \\
\max_m\{\widehat{\mathcal{V}}_\ell(m) - \widehat{\mathcal{V}}_\ell(m-1)\} &= O(n^{-1}), && \text{(A.20)}
\end{aligned}
$$

The observed projected total is exactly zero in exact arithmetic:

$$
\sum_{i=1}^{N_\ell} \widehat{\phi}^{Q4}_{t_{\ell,i},\ell,j} = 0.
$$

The deterministic limit index $\mathcal{V}(z) = c_2^{-1}\int_{-1}^{z} k_2(v)^2\,dv$ is continuous and strictly increasing on $[-1, 1]$.

Proof. Equations (A.17) and (A.18) follow directly from (A.4). Matrix inverse perturbation then gives, uniformly in $|z| \le 1$,

$$
e_{0,3}'(\widehat{Q}_{2\ell}^{\text{eq}})^{-1}p_2(z) = \ell_2(z) + O(n^{-1} + \rho_T^Q).
$$

The numerator and denominator of the variance-time index contain the common factor b^{-2}. After removing it, uniform prefix Riemann sums give

$$
\sum_{i\le m} \widehat{\kappa}^2_{t_{\ell,i},\ell} = b^{-2}\left[n\int_{-1}^{z_{t_{\ell,m},\ell}} k_2(v)^2\,dv + O\{1 + n\rho_T^Q\}\right],
$$

and the full sum has the same expression with the integral over $[-1, 1]$. Because $c_2 > 0$, division proves (A.19). One squared weight is $O(b^{-2})$ whereas the full sum is bounded below by cnb^{-2}, proving (A.20). Taking absolute values of the equivalent weights before squaring does not alter these calculations.

For every nonempty open interval $I \subset (-1, 1)$, $\int_I k_2(v)^2\,dv > 0$: K is positive there and the nonzero quadratic polynomial ℓ_2 has at most two roots. Hence $\mathcal{V}$ is strictly increasing.

Finally, the intercept first-order condition of the quartic weighted least-squares projection is

$$\sum_t w_{t\ell}\{\widehat{a}_{t\ell,j} - q_4(z_{t\ell})'\widehat{\delta}_{\ell j}\} = 0.$$

Multiplication by T^{-1} is exactly the asserted zero-total property. $\square$

Lemma A.7 (Quartic-projected Gaussian limit and positive denominators). *Let $\phi_{t\ell}^{Q4}$ be the sample-Q_4 projection of the ideal raw score. Then the polygonal interpolation of*

$$\sqrt{n}\sum_{t:z_{t\ell}\le z}\phi_{t\ell}^{Q4}$$

converges jointly over the fixed grid in $C[-1,1]^p$ to $Z_\ell(z)$. Under the estimated variance-time map, the limit is $\mathcal{Z}_\ell(s)$. For every coordinate, $\mathcal{R}_{\ell j} > 0$ and $\mathcal{Q}_{\ell j} > 0$ almost surely.

Proof. For the ideal raw score define the $5 \times p$ matrix

$$\delta_\ell^0 = \widehat{Q}_{4\ell}^{-1}T^{-1}\sum_t w_{t\ell}q_4(z_{t\ell})a_{t\ell}'.$$

Lemma A.5 and (A.17) give

$$\sqrt{n}\,\delta_\ell^0 \Rightarrow Q_4^{-1}L_{4\ell}.$$

Uniformly in z,

$$T^{-1}\sum_{t:z_{t\ell}\le z} w_{t\ell}q_4(z_{t\ell}) = \Psi_4(z) + O(n^{-1}).$$

Therefore

$$\begin{aligned}\sqrt{n}\sum_{t:z_{t\ell}\le z}\phi_{t\ell}^{Q4} &= \mathbb{G}_{T\ell}(z) - \left\{T^{-1}\sum_{t:z_{t\ell}\le z} w_{t\ell}q_4(z_{t\ell})'\right\}\sqrt{n}\,\delta_\ell^0\\ &\Rightarrow G_\ell(z) - L_{4\ell}'Q_4^{-1}\Psi_4(z) = Z_\ell(z).\end{aligned}$$

The convergence is joint and uniform by Lemma A.5. The endpoint is zero because $\Psi_4(1) = Q_4e_{0,5}$ and $L_{4\ell}'e_{0,5} = G_\ell(1)$. Equations (A.19)–(A.20), strict monotonicity of $\mathcal{V}$, and asymptotic equicontinuity permit composition with the estimated variance-time index and yield $\mathcal{Z}_\ell$.

To prove positivity, choose fixed $\tau < s_1 < s_2 < 1-\tau$ and put $z_i = \mathcal{V}^{-1}(s_i)$. For coordinate j, the increment $Z_{\ell j}(z_2) - Z_{\ell j}(z_1)$ is a centered Gaussian stochastic integral with scalar integrand

$$k_2(v)\left[\mathbf{1}\{z_1 < v \le z_2\} - q_4(v)'Q_4^{-1}\{\Psi_4(z_2) - \Psi_4(z_1)\}\right].$$

The matrix Σ_ℓ is positive definite because $\Omega(u_\ell)$ is positive definite and $H_D(u_\ell)$ has full row rank. The variance of the increment could therefore be zero only if the bracketed indicator equaled a polynomial of degree at most four almost everywhere on the set where $k_2 \ne 0$. This is impossible: such a polynomial would vanish on one nonempty open interval and equal one on another. Thus the fixed increment is nondegenerate and equals zero with probability zero. A path with zero range on the trimmed interval would make this increment zero, proving $\mathcal{R}_{\ell j} > 0$ almost surely. If $\mathcal{Q}_{\ell j} = 0$, continuity would make the path identically zero on the trimmed interval and again force the same increment to vanish. Hence $\mathcal{Q}_{\ell j} > 0$ almost surely. $\square$

Proof of Lemma 4.6. Let $r_2 \in \mathbb{R}^5$ be the coefficient vector, padded with two zeros, such that $\ell_2(v) = r_2' q_4(v)$, and define

$$W_{0\ell} = \int_{-1}^{1} d\mathbb{W}_\ell(v).$$

Because

$$\int_{-1}^{1} q_4(v) k_2(v)\, dv = \int_{-1}^{1} q_4(v) K(v) q_4(v)'\, dv\, r_2 = Q_4 r_2,$$

we have, for every $z \in [-1, 1]$,

$$\mathrm{Cov}\{W_{0\ell}, Z_\ell(z)\} = \Sigma_\ell \left\{ \int_{-1}^{z} k_2(v)\, dv - r_2' \Psi_4(z) \right\} = 0.$$

Hence $W_{0\ell}$ is independent of the entire projected path $\{Z_\ell(z) : z \in [-1, 1]\}$. Moreover, $\int_{-1}^{1} k_2(v)\, dv = 1$, $\mathrm{Var}(W_{0\ell}) = 2\Sigma_\ell$, and $\mathrm{Cov}\{G_\ell(1), W_{0\ell}\} = \Sigma_\ell$. Therefore

$$G_\ell(1) = \tfrac{1}{2} W_{0\ell} + V_\ell,$$

where $W_{0\ell}$ is jointly independent of (V_ℓ, Z_ℓ). Across the fixed grid, the stacked vector W_0 is Gaussian with block-diagonal covariance whose blocks $2\Sigma_\ell$ are positive definite.

The denominator collection is a measurable function of the projected paths and is strictly positive almost surely by Lemma A.7. Condition on the stacked collection (V, Z). For $x > 0$, the event $S_F(d) \le x$ is a finite rectangle in the nonsingular Gaussian vector W_0:

$$\left| \tfrac{1}{2} W_{0,\ell j} + V_{\ell j} + d_{\ell j} \right| \le x \mathcal{F}_{\ell j} \quad \text{for every } (\ell, j).$$

Its conditional probability is continuous in x by dominated convergence. If $0 < x_1 < x_2$, the larger rectangle contains a set of positive Lebesgue measure outside the smaller one, and the nonsingular Gaussian density is strictly positive there. The conditional probability is therefore strictly increasing in x; integration over (V, Z) preserves continuity and strict increase. This argument applies to every deterministic drift d and proves the lemma. $\square$

A.7 Fitted-score transfer and why quartic projection is enough

Lemma A.8 (Leverage-adjusted score and projected-path equivalence)**.** *The conclusions of Proposition 4.5 hold.*

Proof. Fix ℓ and write

$$\Delta_\ell = \widehat{\beta}_\ell - \beta_\ell^0, \qquad \widehat{L}_\ell = E_0 \widehat{H}_\ell, \qquad L_\ell^0 = E_0 H_\ell^0.$$

Proposition 4.3 and Lemma A.4 give

$$\|\Delta_\ell\| = O_p(n^{-1/2} + b^3 + \rho_T), \qquad \|\widehat{L}_\ell - L_\ell^0\| = O_p(b + n^{-1/2} + T^{-1} + \rho_T). \tag{A.21}$$

The fitted residual has the exact decomposition

$$\widehat{e}_{t\ell} = \varepsilon_{t,T} + R_{t,T}' r_{t\ell} - X_{t\ell}' \Delta_\ell.$$

Leverage and the two regularization bounds. On the kernel support, $p_2(z)$ is bounded and $\|\widehat{H}_\ell\|_{\mathrm{op}} = O_p(1)$. Consequently

$$|\widetilde{h}_{t\ell}| \le Cn^{-1}\|Z_{t,T}R'_{t,T}\|\, O_p(1).$$

The uniform L^r bound and a union bound imply

$$\max_{t\le T}\{\|U_{t,T}\| + \|A_{t,T}\|\} = O_p(T^{1/r}),$$

and hence

$$h_{\max,\ell} := \max_t |\widetilde{h}_{t\ell}| = O_p(T^{1/r}/n) = o_p(1). \tag{A.22}$$

On an event whose probability tends to one, $h_{\max,\ell} < 1/2$. Because $\epsilon_h, \epsilon_d \in (0, 1/2)$, on this event the upper leverage bound and residual-denominator floor are inactive whenever raw leverage is positive; a negative raw leverage is set to zero. Thus $0 \le h^c_{t\ell} \le h_{\max,\ell}$ and

$$|\widehat{e}^{\mathrm{HC3}}_{t\ell} - \widehat{e}_{t\ell}| \le 2h_{\max,\ell}|\widehat{e}_{t\ell}| \quad \text{for all } t. \tag{A.23}$$

This proves why the upper truncation and denominator floor are asymptotically inactive.

Let $\eta^{\mathrm{HC3}}_{t\ell} = \widehat{e}^{\mathrm{HC3}}_{t\ell} - \widehat{e}_{t\ell}$. Since $a_{t\ell} = L^0_\ell V_{t\ell}\varepsilon_{t,T}$, the exact raw-score difference is

$$\widehat{a}_{t\ell} - a_{t\ell} = (\widehat{L}_\ell - L^0_\ell)V_{t\ell}\varepsilon_{t,T} + \widehat{L}_\ell V_{t\ell}R'_{t,T}r_{t\ell} - \widehat{L}_\ell V_{t\ell}X'_{t\ell}\Delta_\ell + \widehat{L}_\ell V_{t\ell}\eta^{\mathrm{HC3}}_{t\ell}. \tag{A.24}$$

Because p_2 is bounded on the support, the four terms are bounded by fixed multiples of

$$\|\widehat{L}_\ell - L^0_\ell\|\|U_{t,T}\|, \quad b^3\|A_{t,T}\|, \quad \|A_{t,T}\|\|\Delta_\ell\|,$$

and

$$h_{\max,\ell}\{\|U_{t,T}\| + b^3\|A_{t,T}\| + \|A_{t,T}\|\|\Delta_\ell\|\},$$

respectively. The weighted laws and the uniform moments give

$$n^{-1}\sum_{t\in\mathcal{I}_\ell}(\|U_{t,T}\|^2 + \|A_{t,T}\|^2) = O_p(1).$$

Using (A.21) and (A.22) in the preceding bounds yields

$$n^{-1}\sum_t K(z_{t\ell})^2\|\widehat{a}_{t\ell} - a_{t\ell}\|^2 = o_p(1). \tag{A.25}$$

The same calculation retains the explicit order

$$n^{-1}\sum_t K(z_{t\ell})^2\|\widehat{a}_{t\ell} - a_{t\ell}\|^2 = O_p\Big[\{b + n^{-1/2} + T^{1/r}/n + \rho_T\}^2\Big]. \tag{A.26}$$

Since

$$\widehat{\phi}_{t\ell} - \phi_{t\ell} = n^{-1}K(z_{t\ell})(\widehat{a}_{t\ell} - a_{t\ell}),$$

(A.25) is exactly $n\sum_t \|\widehat{\phi}_{t\ell} - \phi_{t\ell}\|^2 = o_p(1)$.

For the maximum, (A.24) and the maximal U, A bounds give

$$n^{-1/2}\max_t K(z_{t\ell})\|\widehat{a}_{t\ell}-a_{t\ell}\| \le O_p\left(\frac{T^{1/r}}{\sqrt{n}}[b+n^{-1/2}+b^3+\rho_T]\right) + O_p\left(\frac{T^{2/r}}{n\sqrt{n}}\right) = o_p(1).$$

For example, $T^{1/r}b/\sqrt{n} = T^{1/r-3/5} \to 0$ and $T^{1/r}/n = T^{1/r-4/5} \to 0$; the ridge term is smaller because $\rho_T = o(n^{-1/2})$. This proves (4.11). Multiplication by $z_{t\ell}^d$, $0 \le d \le 4$, changes neither the L^2 nor the maximum bound because $|z_{t\ell}| \le 1$ on the support.

Projected prefix paths. For a scalar raw score sequence x_t, define the sample projection

$$(\mathcal{P}_{4\ell,T}x)_t = q_4(z_{t\ell})'\widehat{Q}_{4\ell}^{-1}T^{-1}\sum_s w_{s\ell}q_4(z_{s\ell})x_s.$$

A normalized projected prefix can be written as a raw prefix minus a bounded partial quartic-moment matrix times the corresponding full-window quartic score. Therefore any component whose raw normalized prefix and five full-window polynomial moments are $o_p(1)$ remains $o_p(1)$ after projection.

The only component of (A.24) that is not immediately negligible under normalized accumulation is the same-window fitting term. Its frozen-design, population-loading component is

$$L_\ell^0\{D_\ell \otimes p_2(z_{t\ell})p_2(z_{t\ell})'\}\Delta_\ell = \{I_p \otimes [\ell_2(z_{t\ell})p_2(z_{t\ell})']\}\Delta_\ell. \tag{A.27}$$

Every coordinate on the right is a polynomial of degree at most four. Because q_4 spans all such polynomials, the sample weighted projection annihilates (A.27) exactly, for every realized Δ_ℓ.

We now bound the remainders. The score-loading term has normalized maximal prefix order

$$\|\widehat{L}_\ell - L_\ell^0\|O_p(1) = o_p(1),$$

and the same order for each full quartic moment. The path Taylor term is bounded, before or after projection, by

$$C\sqrt{n}\,b^3\left\{n^{-1}\sum_{t\in\mathcal{I}_\ell}(1+\|A_{t,T}\|)\right\} = O_p(\sqrt{n}\,b^3) = O_p(T^{-1/5}).$$

For the fitting term, decompose

$$A_{t,T} = D_\ell + \{D(u_t)-D_\ell\} + \{A_{t,T}-EA_{t,T}\} + \{EA_{t,T}-D(u_t)\}.$$

After exact removal of the D_ℓ, L_ℓ^0 polynomial in (A.27), replacing L_ℓ^0 by $\widehat{L}_\ell$ contributes at most

$$\begin{aligned} C\sqrt{n}\,\|\widehat{L}_\ell - L_\ell^0\|\|\Delta_\ell\| &= O_p\{b+n^{-1/2}+\sqrt{n}\,b^4+\sqrt{n}\,b\rho_T+\sqrt{n}\rho_T n^{-1/2}+\sqrt{n}\rho_T b^3\} \\ &= o_p(1). \end{aligned}$$

The smooth local variation contributes

$$C\sqrt{n}\,b\|\Delta_\ell\| = O_p\{b+\sqrt{n}\,b^4+\sqrt{n}\,b\rho_T\} = o_p(1).$$

For the centered component, Lemma A.2 applied to the finite collection of matrix-valued polynomial loadings gives an $O_p(1)$ normalized maximal prefix process; multiplication by $\|\Delta_\ell\| = o_p(1)$ makes it $o_p(1)$. Its full quartic moments obey the same bound. The triangular mean approximation contributes $O_p(\sqrt{n}\,T^{-1}\|\Delta_\ell\|) = o_p(1)$.

Finally, (A.23) gives the normalized HC3-prefix bound

$$C\sqrt{n}\,h_{\max,\ell}\left\{n^{-1}\sum_{t\in\mathcal{I}_\ell}(\|U_{t,T}\| + b^3\|A_{t,T}\| + \|A_{t,T}\|\|\Delta_\ell\|)\right\} = O_p(T^{1/r}/\sqrt{n}) = o_p(1).$$

Projection preserves this order because $\widehat{Q}_{4\ell}^{-1} = O(1)$ and all partial quartic moment matrices are uniformly bounded. Combining the components proves (4.12). The five powers z^d, $0 \le d \le 4$, needed in multiplier projection coefficients are already included among the full-window polynomial moments just controlled. This proves the lemma and Proposition 4.5. ☐

Lemma A.9 (Raw-score centering)**.** *Uniformly over the fixed grid,*

$$\|\bar{a}_\ell\| = O_p(n^{-1/2} + b^3 + \rho_T) = o_p(1). \tag{A.28}$$

Proof. The ideal weighted mean obeys

$$\left\|\frac{\sum_t w_{t\ell}a_{t\ell}}{\sum_t w_{t\ell}}\right\| = O_p(n^{-1/2})$$

by Lemma A.2, because $\sum_t w_{t\ell} \asymp T$ and $\sum_t w_{t\ell}^2 = O(T/b)$. Apply the decomposition (A.24), whose componentwise bounds were established in Lemma A.8, without taking prefix maxima. The score-loading term is $o_p(1)$, the path remainder is $O_p(b^3)$, the fitted term is $O_p(\|\widehat{\beta}_\ell - \beta_\ell^0\|)$, and the leverage-adjustment term is $o_p(1)$. Equation (4.8) then proves (A.28). ☐

A.8 Conditional multiplier functional central limit theorem

Lemma A.10 (Feasible conditional multiplier functional central limit theorem)**.** *Conditionally on the sample, in probability, the joint collection consisting of all multiplier endpoints and all draw-specific quartic-projected prefix paths converges to an independent copy of $\{G_\ell(1), Z_\ell(\cdot)\}_{\ell\le M}$. The result remains true after raw-score centering and replacement of the ideal scores, projection matrices, and variance time by their feasible counterparts.*

Proof. We give the argument in nine steps corresponding to the objects defined in the preceding section.

Step 1: realized covariance arrays for ideal scores. Let f and g be any two bounded polynomial-prefix loadings appearing in the endpoint or quartic-projection process. Within one window, the conditional covariance of the normalized ideal multiplier sums is

$$n^{-1}\sum_t f(z_{t\ell})g(z_{t\ell})H_{D\ell}C_{t,T}H_{D\ell}'. \tag{A.29}$$

The multiplier class contains only a finite number of polynomial coordinates, while the prefix endpoint ranges over a one-jump bounded-variation class. Equation (A.6), applied to C, gives uniform convergence of (A.29) to the covariance kernel obtained by replacing $C_{t,T}$ with $\Omega(u_\ell)$ and taking a Riemann integral. The sample martingale process used $\Gamma_{t,T}$ in its predictable covariance; the same weighted law applied to Γ, together with $E\Gamma_t(u) = EC_t(u) = \Omega(u)$, gives exactly the same limit and establishes the realized-versus-predictable covariance correspondence.

Step 2: joint grid covariance and common calendar multipliers. For two grid points, the conditional covariance has the same calendar sum but contains the product of their two local loadings. The same multiplier ξ_t appears in both terms, so any finite-sample overlap covariance is represented exactly. Since the fixed grid points are distinct and $b \to 0$, the supports are disjoint for all sufficiently large T, and the limiting cross-window covariance is zero. This conclusion is reached after forming the calendar-ordered covariance; no local prefixes are reordered across windows.

Step 3: conditional Lindeberg. For any scalar Cramér–Wold combination, write its normalized ideal multiplier sum as $\sum_t c_{t,T}\xi_t$. Then

$$m_T := \max_t |c_{t,T}| \le Cn^{-1/2} \max_{t\le T} \|U_{t,T}\| = O_p(T^{1/r}/\sqrt{n}) = o_p(1),$$

and $\sum_t c_{t,T}^2 = O_p(1)$ by Step 1. Conditional on the data,

$$\begin{aligned}
&\sum_t E^*\{c_{t,T}^2\xi_t^2\mathbf{1}(|c_{t,T}\xi_t| > \epsilon)\} \\
&\qquad \le \left(\sum_t c_{t,T}^2\right) E\{\xi_1^2\mathbf{1}(|\xi_1| > \epsilon/m_T)\} = o_p(1),
\end{aligned}$$

using uniform integrability of ξ_1^2. This is the conditional Lindeberg condition.

Step 4: conditional finite-dimensional convergence. Conditionally on the data, every finite vector of ideal multiplier sums is exactly Gaussian. Steps 1 and 2 show convergence of its covariance matrix, so its conditional finite-dimensional distributions converge in probability to those of an independent copy of the sample Gaussian limit. Step 3 verifies the Lindeberg condition for the triangular multiplier sequence, although Gaussianity already reduces finite-dimensional convergence to the covariance calculation.

Step 5: conditional tightness on the discrete mesh. For any one of the finitely many ideal multiplier prefix processes, let $\Delta_T^*(z_1, z_2)$ denote its increment between two mesh points with $|z_2 - z_1| \le \delta$. The deterministic part of the realized covariance sum is bounded by $C(\delta + n^{-1})$: the interval contains at most $n\delta + O(1)$ observations, and its polynomial and kernel loadings are bounded. Uniformly over prefix intervals, (A.6) leaves a stochastic covariance remainder

$$r_T = O_p(n^{-1/2} + T^{-1}) = o_p(1).$$

Consequently, with probability approaching one,

$$\sup_{|z_2 - z_1|\le\delta} E^*|\Delta_T^*(z_1, z_2)|^2 \le C(\delta + n^{-1}) + r_T. \tag{A.30}$$

Conditionally on the data, each increment is centered Gaussian, so its fourth moment is three times the square of its variance. Partition $[-1, 1]$ into $O(\delta^{-1})$ blocks of length at most δ. The usual adjacent-block argument and (A.30) give, for every $\eta > 0$,

$$P^*\{\omega(\Delta_T^*, 2\delta) > \eta\} \le C\eta^{-4}\delta^{-1}\{\delta + n^{-1} + r_T\}^2 + o_p(1),$$

where ω is the uniform modulus of continuity of the polygonal interpolant. Hence, for every $\epsilon, \eta > 0$,

$$\lim_{\delta \downarrow 0} \limsup_{T \to \infty} P[P^*\{\omega(\Delta_T^*, 2\delta) > \eta\} > \epsilon] = 0.$$

The finite collection of coordinates obeys the same bound. The largest variance-time mesh increment is $O(n^{-1})$ by (A.20), and the largest normalized ideal summand is $o_p(1)$ by Step 3. Thus polygonal interpolation and transfer to the variance-time mesh are valid.

Step 6: transfer from ideal to feasible leverage-adjusted scores. For an unprojected feasible-minus-ideal multiplier prefix, conditional Doob inequality gives

$$E^* \max_m \left\| \sqrt{n} \sum_{i=1}^{m} \xi_{t_{\ell,i}} (\widehat{\phi}_{t_{\ell,i},\ell} - \phi_{t_{\ell,i},\ell}) \right\|^2 \le 4n \sum_t \|\widehat{\phi}_{t\ell} - \phi_{t\ell}\|^2 = o_p(1).$$

For projected paths, write the projection as a raw prefix minus a bounded partial quartic-moment matrix times the full-sample vector of five quartic moments. The z^d versions of (4.10), $0 \le d \le 4$, control all those terms by the same conditional Doob bound. Thus the feasible and ideal multiplier paths differ by $o_p(1)$ uniformly. Equivalently, the realized covariance arrays computed from feasible scores have the same limits: by Cauchy–Schwarz,

$$\begin{aligned} &\left\| n \sum_t (\widehat{\phi}_{t\ell} \widehat{\phi}_{t\ell}' - \phi_{t\ell}\phi_{t\ell}') \right\| \\ &\quad \le \left(n \sum_t \|\widehat{\phi}_{t\ell} - \phi_{t\ell}\|^2 \right)^{1/2} \left[\left(n \sum_t \|\widehat{\phi}_{t\ell}\|^2 \right)^{1/2} + \left(n \sum_t \|\phi_{t\ell}\|^2 \right)^{1/2} \right] = o_p(1). \end{aligned}$$

Step 7: raw-score centering. Subtracting $\overline{a}_{\ell j}$ changes a normalized multiplier prefix by

$$n^{-1/2} \overline{a}_{\ell j} \sum_{t: z_{t\ell} \le z} K(z_{t\ell}) \xi_t.$$

Its conditional variance is bounded by $C\|\overline{a}_\ell\|^2 = o_p(1)$ by (A.28). The same calculation after multiplication by each coordinate of q_4 controls the draw-specific projection coefficient. Thus kernel-weighted centering of the raw score is asymptotically negligible.

Step 8: sample projection matrices and estimated variance time. Equation (A.17) and the continuous inverse map transfer the quartic projection from its sample Gram matrix to Q_4. Equations (A.19)–(A.20), together with conditional equicontinuity from Step 5, transfer the paths from the deterministic limit index to the estimated discrete variance-time index. The explicit bridge term is also continuous; in exact arithmetic its endpoint is already zero because the projection contains an intercept.

Step 9: reprojection inside each draw and ratio maps. The coefficient $\widehat{\delta}_{\ell j}^*$ in (3.14) is a linear function of the multiplied, centered raw-score sequence. Steps 1–8 give joint convergence of this coefficient and every raw prefix. Continuous linear mapping therefore yields Z_ℓ and verifies the stated order of transformations. In particular, the projection is recomputed after multiplication rather than applied before the multiplier.

The range and matched quadratic maps are continuous on $C[0, 1]$ at paths with positive denominators, and Lemma A.7 supplies that positivity. The fixed finite maximum is continuous. Hence the two conditional ratio maxima converge weakly at every continuity point of their limit laws. Lemma 4.6 gives continuity of those laws and the unique-crossing property in (4.14); the exact conditional quantiles therefore converge to their central limiting quantiles. Finally, conditional on the data, the B_T simulated statistics are iid. The conditional Dvoretzky–Kiefer–Wolfowitz bound makes their empirical distribution uniformly close to the exact conditional distribution when $B_T \to \infty$. Combined with the same unique-crossing property, this proves consistency of the simulated raw quantiles. □

A.9 Proofs of the joint limit and band results

Proof of Theorem 4.7. Proposition 4.3 and Lemma A.5 give the joint endpoint limit. Lemmas A.7, A.8, and A.9 give the joint feasible projected-path and variance-time limits. Lemma A.10 gives the corresponding conditional multiplier limit and empirical-quantile consistency. All mappings are applied jointly to a fixed finite collection, which proves the theorem. □

Proof of Theorem 4.8. By Proposition 4.3, jointly over the fixed grid,

$$\sqrt{n}\{\widehat{\theta}_j(u_\ell) - \theta_{0,j}(u_\ell)\} \Rightarrow G_{\ell j}(1).$$

By Lemmas A.7 and A.8, $\sqrt{n}\widehat{R}_{\ell j}^{Q4} \Rightarrow \mathcal{R}_{\ell j}$ jointly, and every limiting range is positive. Therefore

$$\max_{\ell,j} \frac{|\widehat{\theta}_j(u_\ell) - \theta_{0,j}(u_\ell)|}{\widehat{R}_{\ell j}^{Q4}} \Rightarrow S_R^0.$$

Theorem 4.7 gives $\widehat{c}_R(1-\alpha) \to_p c_R(1-\alpha)$. Continuity and unique crossing at the quantile imply rejection probability $\alpha + o(1)$, and inversion gives simultaneous coverage $1 - \alpha + o(1)$. Moreover, $\epsilon_T^{\mathrm{num}} = o(n^{-1/2})$ while $\sqrt{n}\widehat{R}_{\ell j}^{Q4}$ has a positive limit, so the numerical floor is inactive with probability tending to one. □

Proof of Corollary 4.9. The map

$$f \longmapsto \left\{(1-2\tau)^{-1} \int_\tau^{1-\tau} f(s)^2 \, ds\right\}^{1/2}$$

is continuous on $C[0,1]$ and positive at every limiting projected path by Lemma A.7. Repeat the proof of Theorem 4.8 with this map and use the separately simulated quadratic quantile. □

A.10 Global restricted estimators

Proof of Lemma 4.11. Fix $K_0 \in \{1, 2\}$ and suppress that subscript. Under the full-path null,

$$y_{t,T} = \mathcal{R}_{t,T}'\gamma_0 + \varepsilon_{t,T}, \qquad \mathcal{Z}_{t,T}\varepsilon_{t,T} = U_{t,T} \otimes b_{K_0}(u_t).$$

Lemma A.2 with coefficients T^{-1} and ordinary Riemann sums give

$$\widehat{\mathcal{D}} = \mathcal{D} + O_p(T^{-1/2}), \tag{A.31}$$

$$T^{-1}\sum_t \mathcal{Z}_{t,T}\varepsilon_{t,T} = O_p(T^{-1/2}), \tag{A.32}$$

$$T^{-1}\sum_t \mathcal{Z}_{t,T}\mathcal{Z}_{t,T}'\varepsilon_{t,T}^2 = \mathcal{S} + o_p(1). \tag{A.33}$$

The last equation is the global weighted law for $C_{t,T} \otimes \{b_{K_0}(u_t)b_{K_0}(u_t)'\}$.

The identity-weight preliminary estimator satisfies the exact relation

$$\widehat{\gamma}^{(0)} - \gamma_0 = \{\widehat{\mathcal{D}}'\widehat{\mathcal{D}} + \rho_T^g I\}^{-1}\widehat{\mathcal{D}}' T^{-1} \sum_t \mathcal{Z}_{t,T}\varepsilon_{t,T} \\ - \rho_T^g \{\widehat{\mathcal{D}}'\widehat{\mathcal{D}} + \rho_T^g I\}^{-1}\gamma_0.$$

Full column rank, (A.31), and $\sqrt{T}\rho_T^g \to 0$ imply $\widehat{\gamma}^{(0)} - \gamma_0 = O_p(T^{-1/2})$ and consistency.

Suppose inductively that $\delta^{(j-1)} = \widehat{\gamma}^{(j-1)} - \gamma_0 = O_p(T^{-1/2})$. Then $\widehat{e}_t^{(j-1)} = \varepsilon_{t,T} - \mathcal{R}_{t,T}'\delta^{(j-1)}$ and

$$(\widehat{e}_t^{(j-1)})^2 - \varepsilon_{t,T}^2 \\ = -2\varepsilon_{t,T}\mathcal{R}_{t,T}'\delta^{(j-1)} + (\mathcal{R}_{t,T}'\delta^{(j-1)})^2.$$

The explicit envelope moments in Assumption 4.2 imply

$$T^{-1}\sum_t \|\mathcal{Z}_{t,T}\mathcal{Z}_{t,T}'\||\varepsilon_{t,T}|\|\mathcal{R}_{t,T}\| = O_p(1),$$

$$T^{-1}\sum_t \|\mathcal{Z}_{t,T}\mathcal{Z}_{t,T}'\|\|\mathcal{R}_{t,T}\|^2 = O_p(1).$$

Consequently

$$\widehat{\mathcal{S}}^{(j)} = \mathcal{S} + o_p(1), \qquad \widehat{W}^{(j)} = \mathcal{S}^{-1} + o_p(1),$$

because Assumption 4.2 imposes $\rho_T^S + \widetilde{\rho}_T^S \to 0$ and continuity of matrix inversion applies on a neighborhood of the positive-definite matrix $\mathcal{S}$. The exact weighted linear-GMM identity gives

$$\widehat{\gamma}^{(j)} - \gamma_0 = \{\widehat{\mathcal{D}}'\widehat{W}^{(j)}\widehat{\mathcal{D}} + \rho_T^g I\}^{-1}\widehat{\mathcal{D}}'\widehat{W}^{(j)} T^{-1}\sum_t \mathcal{Z}_{t,T}\varepsilon_{t,T} + o_p(T^{-1/2}) \\ = \mathcal{H}_{K_0} T^{-1}\sum_t \{U_{t,T} \otimes b_{K_0}(u_t)\} + o_p(T^{-1/2}).$$

Induction over the fixed number of updates proves (4.20) and the root-T rate.

For asymptotic normality, apply the same calendar-ordered Cramér–Wold argument as in Lemma A.5 to $T^{-1/2}\sum_t\{U_{t,T} \otimes b_{K_0}(u_t)\}$. Its predictable covariance converges to $\mathcal{S}_{K_0}$ by the global law for Γ, and its conditional Lindeberg expectation is bounded by $CT^{1-r/2} \to 0$. Hall and Heyde's Corollary 3.1 therefore gives a normal limit with covariance $\mathcal{S}_{K_0}$. Premultiplication by $\mathcal{H}_{K_0}$ gives

$$\mathcal{H}_{K_0}\mathcal{S}_{K_0}\mathcal{H}_{K_0}' = \{\mathcal{D}_{K_0}'\mathcal{S}_{K_0}^{-1}\mathcal{D}_{K_0}\}^{-1}.$$

Finally,

$$\sqrt{n}\|\widehat{\gamma}_{K_0} - \gamma_0\| = O_p(\sqrt{Tb}/\sqrt{T}) = O_p(\sqrt{b}) = o_p(1),$$

which proves (4.21). If the null holds only at the finite grid, the identity $y_t = \mathcal{R}_t'\gamma_0 + \varepsilon_t$ fails for the off-grid observations that dominate the global moments. The expansion then contains a generally nonzero $O(1)$ population moment and is not valid. □

A.11 Path tests and local alternatives

Proof of Theorem 4.14. Under the full-path null, Lemma 4.11 gives

$$\sqrt{n}\max_{\ell\le M}\|L_{K_0}(u_\ell)(\widehat{\gamma}_{K_0}-\gamma_0)\| = o_p(1).$$

Hence the endpoint vector centered by the fitted restricted path has the same joint limit as the endpoint vector centered by the true path. Its denominator and multiplier critical value are unchanged, so Theorem 4.8 gives asymptotic rejection probability α.

Under a fixed alternative, the global weighted laws applied recursively to the finite collection of update maps give $\widehat{\gamma}_{K_0}\to_p \gamma^*_{K_0}$ under the uniqueness and nonsingularity conditions stated in the theorem. If one displayed pseudo-true discrepancy is nonzero, the corresponding numerator converges to a nonzero constant whereas its denominator is $O_p(n^{-1/2})$. The maximum therefore diverges and the adjusted-range test is consistent. □

Proof of Corollary 4.15. Use Lemma 4.11 and the quadratic joint limit in exactly the same way. Under a fixed alternative the quadratic denominator is also $O_p(n^{-1/2})$ with a positive normalized limit. The quadratic statistic must be compared with its own conditional multiplier quantile. □

Proof of Corollary 4.16. Because S is fixed, stack all equation-specific endpoint and prefix coordinates before applying the Cramér–Wold argument. The assumed joint martingale-difference restriction makes every scalar combination a martingale transform with respect to the common filtration. The joint physical- dependence and moment conditions give the same weighted laws used in Lemmas A.5 and A.10. In the martingale case the mapped endpoint covariance is $\Sigma_\ell^{\mathrm{sys}}$; under short memory it is $\Sigma_\ell^{\mathrm{sys},D}$. Both retain the relevant within- and cross-equation covariance blocks. Their assumed positive definiteness on the displayed stacked coordinates gives nondegenerate endpoint and projected-path limits.

Within a multiplier draw, the same calendar-time multiplier multiplies every equation-specific score. The realized conditional covariance therefore reproduces both the within-equation and cross-equation covariance blocks. Conditional Lindeberg and tightness follow from the preceding proofs after replacing the finite index set (ℓ, j) by (s, ℓ, j). Applying the adjusted-range or quadratic ratio map and then one maximum over this enlarged finite set proves the joint band statements. Under a system-wide constant or linear full-path null, Lemma 4.11 applies equation by equation; fixed S makes the maximum of the resulting $O_p(T^{-1/2})$ errors remain $O_p(T^{-1/2})$. The proof of Theorem 4.14 then gives the corresponding system path tests. □

Proof of Proposition 4.17. Let $a_T = n^{-1/2}$ and suppose $\theta_T(u) = L_{K_0}(u)\gamma_0 + a_T\Delta(u)$. The global sample moment at γ_0 is

$$\begin{aligned} T^{-1}\sum_t \mathcal{Z}_{K_0,t,T}\{y_{t,T}-\mathcal{R}'_{K_0,t,T}\gamma_0\} &= T^{-1}\sum_t\{U_{t,T}\otimes b_{K_0}(u_t)\} \\ &\quad + a_T T^{-1}\sum_t\{A_{t,T}\Delta(u_t)\}\otimes b_{K_0}(u_t). \end{aligned}$$

The second sample average converges to $h_{K_0,\Delta}$ by the global weighted law. Every updated covariance weight still converges to $\mathcal{S}_{K_0}^{-1}$, because the additional residual component is $O(a_T)=o(1)$ and the fitted parameter remains $O_p(a_T + T^{-1/2})$ from γ_0. Repeating the exact linear-GMM expansion in the proof of Lemma 4.11 gives

$$\widehat{\gamma}_{K_0}-\gamma_0 = a_T\mathcal{H}_{K_0}h_{K_0,\Delta} + O_p(T^{-1/2}) + o_p(a_T).$$

Since $T^{-1/2}/a_T = \sqrt{b} \to 0$, this proves (4.24).

The triangular path θ_T satisfies the derivative bounds in Assumption 4.1 uniformly: the derivatives of its additional component are multiplied by a_T. The proof of Proposition 4.3 therefore applies uniformly and gives

$$\sqrt{n}\{\widehat{\theta}(u_\ell) - \theta_T(u_\ell)\} \Rightarrow G_\ell(1).$$

Subtracting $L_{K_0}(u_\ell)\widehat{\gamma}_{K_0}$ and using (4.24) yields (4.25) with drift $d_{\ell,\Delta}$.

All fitted-score transfer bounds remain valid uniformly along the local alternative. The deterministic local departure is absorbed by the local quadratic fit, and its Taylor remainder is multiplied by a_T, so it cannot alter the central denominator limit. The multiplier is raw-score centered and continues to estimate the central law. Hence its critical value converges to $c_F(1-\alpha)$. Continuous mapping proves (4.26), and quantile continuity gives the exact power expression (4.27); continuity for the drifted ratio law is supplied by Lemma 4.6. If $d_\Delta = 0$, this law is the null law and power equals α.

If $d_\Delta \neq 0$, select a coordinate i with $d_i \neq 0$. For almost every realization, $0 < \mathcal{F}_i < \infty$ and $|G_i| < \infty$, so $|G_i + \lambda d_i|/\mathcal{F}_i \to \infty$ as $|\lambda| \to \infty$. Dominated convergence gives power tending to one and therefore exceeding size for all sufficiently large $|\lambda|$.

□

A.12 Proofs for the serial-dependence extension

This subsection proves Theorem 4.10 and Corollary 4.12. The argument combines the physical-dependence approximation used above with a dependent wild multiplier in the spirit of Shao (2010a) and Bücher and Kojadinovic (2016). We retain the notation introduced at the start of Appendix A.

Lemma A.11 (Uniform covariance summability)**.** *Under Assumption 4.1(ii)–(iii),*

$$\sup_{u\in[0,1]} \sum_{h\in\mathbb{Z}} (1+|h|)\|\Gamma_U(u,h)\| < \infty. \tag{A.34}$$

The series defining $\Lambda(u)$ is therefore absolutely and uniformly convergent. If $E\{U_t(u)\} = 0$ and the L^r local-stationarity condition in Assumption 4.1(ii) holds, Λ is continuous.

Proof. It is enough to consider scalar coordinates. Write $U_t(u) = \sum_{j\geq 0} \mathcal{P}_{t-j} U_t(u)$ in L^2. Orthogonality of martingale projections and the coupling bound used in Lemma A.2 give, for $h \geq 0$,

$$|\operatorname{Cov}\{U_0(u), U_h(u)\}| \leq \sum_{j\geq 0} \|\mathcal{P}_{-j}U_0(u)\|_2 \|\mathcal{P}_{-j}U_h(u)\|_2 \leq \sum_{j\geq 0} \delta_r^U(j)\delta_r^U(j+h).$$

The same bound applies to negative lags after transposition. Multiplying by $1+h$, summing, and using (4.2) proves (A.34). Uniform convergence follows because the dependence coefficients are defined with a supremum over u.

For continuity, split the covariance series at a fixed H. Each of the finitely many terms with $|h| \leq H$ is continuous by the L^r continuity of $U_t(u)$ and Hölder's inequality. The two tails are uniformly small by (A.34). Letting first $v \to u$ and then $H \to \infty$ proves continuity of Λ. □

Lemma A.12 (Local functional central limit theorem under serial dependence)**.** *Under Assumption 4.3,* *jointly over the fixed grid,*

$$\{\widetilde{\mathbb{G}}_{T\ell}(\cdot)\}_{\ell\le M} \Rightarrow \{G_\ell^D(\cdot)\}_{\ell\le M} \quad in\ \{C[-1,1]^p\}^M.$$

The result remains true after adjoining the finitely many processes with loadings $q_4(z)k_2(z)$.

Proof. Fix a finite collection of prefix endpoints and take a Cramér–Wold linear combination. As in the proof of Lemma A.5, it has the form

$$S_T = n^{-1/2}\sum_{t=1}^{T} c_{t,T}' U_{t,T},$$

where the deterministic $c_{t,T}$ are uniformly bounded and supported on a fixed union of local windows. Replace $U_{t,T}$ by its stationary approximation at u_t, at a total L^2 cost $O(n^{1/2}/T) = o(1)$.

For $m \ge 0$, let

$$U_t^{(m)}(u) = E\{U_t(u) \mid \eta_{t-m}, \ldots, \eta_t\} - E\{U_t(u)\}.$$

The sequence $U_t^{(m)}(u_t)$ is m-dependent. Successive innovation replacement and Lemma A.2 yield

$$\left\| n^{-1/2}\sum_t c_{t,T}'\{U_t(u_t) - U_t^{(m)}(u_t)\}\right\|_2 \le C\sum_{k>m}\delta_r^U(k),$$

uniformly in T. The right side tends to zero as $m\to\infty$.

For fixed m, separate the summands into blocks of length exceeding m and insert gaps of length m. The retained blocks are independent. The gaps are negligible as the long-block length subsequently tends to infinity, and the $r > 8$ moment bound gives the Lindeberg condition. The independent-block central limit theorem therefore applies once the covariance has been identified.

For one local window, write the covariance as a lag sum. For every fixed h,

$$\begin{aligned} n^{-1}\sum_t c_{t,T}' E\{U_{t,T}U_{t+h,T}'\}c_{t+h,T} \\ \longrightarrow \int_{-1}^{1} c_\ell(v)'\Gamma_U(u_\ell, h)c_\ell(v)\,dv, \end{aligned}$$

where $c_\ell(v)$ denotes the relevant equivalent-kernel and prefix loading. Indeed, $h/T \to 0$, $h/(Tb_T)\to 0$, and the Riemann approximation is uniform over every fixed finite set of lags. Lemma A.11 permits passage from a finite lag sum to all $h\in\mathbb{Z}$. The resulting covariance replaces $\Omega(u_\ell)$ by $\Lambda(u_\ell)$ and contains $\int k_2(v)^2\,dv$ over the intersection of the selected prefixes. This is the covariance of G_ℓ^D.

For $\ell \ne \ell'$, the two local windows are separated by at least cT calendar observations for some $c > 0$ and all large T. Their covariance is bounded by a constant times the tail of (A.34) beyond $cT/2$, and hence tends to zero. The joint Gaussian limit consequently has independent grid-point blocks.

It remains to establish tightness. For an increment over local coordinates $(z_1, z_2]$, apply Lemma A.2 with the increment weights. Since their squared sum is $O\{(|z_2 - z_1| + n^{-1})\}$ after normalization,

$$E\|\mathbb{G}_{T\ell}(z_2) - \mathbb{G}_{T\ell}(z_1)\|^4 \le C(|z_2 - z_1| + n^{-1})^2.$$

The partition argument following (A.16) proves asymptotic equicontinuity. The largest interpolated jump is $o_p(1)$ by the rth-moment bound. The polynomial loadings are bounded on $[-1,1]$, so the same proof applies to every process needed by the quartic projection. □

Lemma A.13 (Uniform local lag-window covariance)**.** *Let $\mathscr{A}$ be the finite collection of bounded piecewise Lipschitz functions generated by k_2, the coordinates of $q_4 k_2$, and their products. For $a \in \mathscr{A}$ and $x \in [-1, 1]$, write*

$$a_{t\ell}(x) = a(z_{t\ell})\mathbf{1}\{|z_{t\ell}| < 1, z_{t\ell} \le x\}.$$

Set $\kappa_T(h) = \exp(-|h|/L_T)$ and, with terms outside $\{1, \dots, T\}$ omitted, define

$$\widehat{C}_{T\ell}(a, x; \widetilde{a}, \widetilde{x}) = \frac{1}{n} \sum_{h \in \mathbb{Z}} \kappa_T(h) \sum_t a_{t\ell}(x) \widetilde{a}_{t+h,\ell}(\widetilde{x}) H_{D\ell} U_{t,T} U'_{t+h,T} H'_{D\ell}.$$

Also define the local covariance modulus

$$\omega_\Gamma(\delta) = \max_{\ell \le M} \sup_{|u - u_\ell| \le \delta} \sum_{h \in \mathbb{Z}} \|\Gamma_U(u, h) - \Gamma_U(u_\ell, h)\|.$$

Lemma A.11 and continuity at each fixed lag imply $\omega_\Gamma(\delta) \to 0$ as $\delta \downarrow 0$. Under Assumption 4.3, uniformly over the fixed grid, $a, \widetilde{a} \in \mathscr{A}$, and $x, \widetilde{x} \in [-1, 1]$,

$$\begin{aligned} &\widehat{C}_{T\ell}(a, x; \widetilde{a}, \widetilde{x}) \\ &\quad = H_{D\ell} \Lambda(u_\ell) H'_{D\ell} \int_{-1}^{x \wedge \widetilde{x}} a(v) \widetilde{a}(v)\, dv + O_p\left(\frac{L_T}{\sqrt{n}} + \omega_\Gamma(b) + \frac{L_T^2}{T} + \frac{1}{L_T} + \frac{1}{n} \right). \end{aligned} \tag{A.35}$$

The remainder is $o_p(1)$ under (4.17). If $\ell \neq \ell'$, the analogous cross-window covariance converges uniformly to zero. For endpoints $x_1 < x_2$, define $\Delta a_{t\ell}(x_1, x_2) = a(z_{t\ell})\mathbf{1}\{x_1 < z_{t\ell} \le x_2\}$. Let $\widehat{C}^{\Delta}_{T\ell}$ denote the same lag-window sum with both loadings replaced by Δa. Set

$$p_0 = \frac{r}{2}, \qquad \nu = \frac{1}{1/2 + 1/p_0} = \frac{2r}{r+4}, \qquad a_T^D = \frac{1}{n} + \left(\frac{L_T}{\sqrt{n}} \right)^{\nu}.$$

Uniformly in the endpoints,

$$\sup_{x_1 < x_2} \frac{\|\widehat{C}^{\Delta}_{T\ell}(a, x_1, x_2)\|_{\mathrm{op}}}{|x_2 - x_1| + a_T^D} = O_p(1). \tag{A.36}$$

Proof. It is enough to prove the result for one matrix coordinate. Expand the centered part by lags. For $h \ge 0$, reindex by $s = t + h$ and put

$$Y_{s,h}(u) = U_{s-h}(u) U_s(u)' - E\{U_{s-h}(u) U_s(u)'\}.$$

This process is causal with respect to $\mathcal{F}_s$; the reindexing is needed because $U_t U'_{t+h}$ is not $\mathcal{F}_t$-measurable when $h > 0$. Hölder's inequality and innovation replacement give

$$\delta^{Y_h}_{r/2}(k) \le C\{\delta^U_r(k) + \mathbf{1}\{k \ge h\} \delta^U_r(k - h)\}, \qquad k \ge 0. \tag{A.37}$$

The same bound, after transposition, applies to negative lags. Consequently, $\sup_h \sum_{k \ge 0} \delta^{Y_h}_{r/2}(k) < \infty$. After the reindexing, the centered contribution at a positive lag is

$$\left\| \sup_{x,\widetilde{x}} \left| \frac{1}{n} \sum_s a_{s-h,\ell}(x) \widetilde{a}_{s\ell}(\widetilde{x}) \{Y_{s,h}(u_s)\} \right| \right\|_{r/2} \leq C n^{-1/2}.$$

For a fixed h, the product of the two prefix indicators is an indicator of a calendar interval, up to at most $|h|+2$ endpoint observations. Lemma A.2, with coefficients bounded by C/n, is maximal over the remaining partial sum and proves the display. The bound is uniform in h. Replacing u_s by u_{s-h}, or equivalently replacing the second local time u_{t+h} by u_t in the original indexing, costs $C|h|/T$ by (4.1); replacing the triangular array by its stationary approximation costs C/T. Since

$$\sum_{h \in \mathbb{Z}} \kappa_T(h) = O(L_T), \qquad \sum_{h \in \mathbb{Z}} |h| \kappa_T(h) = O(L_T^2),$$

Minkowski's inequality yields the uniform centered bound

$$O_p\left(\frac{L_T}{\sqrt{n}} + \frac{L_T^2}{T} \right). \tag{A.38}$$

This argument also shows explicitly why uniformity over prefix endpoints does not introduce a growing entropy factor: on the observation mesh each pair of prefixes reduces to a partial sum, and Lemma A.2 is already maximal over partial sums. For completeness, consider an increment interval containing m local observations and put $p_0 = r/2 > 4$. At any fixed lag, Lemma A.2 bounds its centered product sum in L^{p_0} by $C\sqrt{m + |h| + 1}$. At a dyadic scale with block length m, there are at most Cn/m blocks. Hence

$$\left\| \max_{J \in \mathcal{D}_m} \left| \sum_{s \in J} \{Y_{s,h}(u_s)\} \right| \right\|_{p_0} \leq C(n/m)^{1/p_0} \sqrt{m + |h| + 1},$$

where $\mathcal{D}_m$ is the collection of dyadic blocks at that scale. Every observation interval is a disjoint union of at most two blocks at each dyadic scale. Summing the resulting geometric series and then the exponentially weighted lag bounds gives, uniformly over all increment intervals,

$$O_p\left[\frac{L_T}{\sqrt{n}} (|x_2 - x_1| + n^{-1})^{1/2 - 1/p_0} + \frac{L_T^{3/2}}{n} \right]$$

for the centered quadratic form. Young's inequality, with conjugate power $\nu = 1/(1/2 + 1/p_0)$, bounds this display by $O_p(|x_2 - x_1| + a_T^D)$. The term $L_T^{3/2}/n$ is smaller than a_T^D under $L_T/\sqrt{n} \to 0$. Covariance summability bounds the mean by $C(|x_2 - x_1| + n^{-1})$; the local-stationarity approximation errors are proportional to the interval length and are absorbed because $L_T^2/T = o(1)$. This proves (A.36) and records the rate needed below.

It remains to identify the mean. For each fixed h, local stationarity and the Riemann-sum calculation in Lemma A.3 give

$$\begin{aligned} &\frac{1}{n} \sum_t a_{t\ell}(x) \widetilde{a}_{t+h,\ell}(\widetilde{x}) H_{D\ell} E(U_{t,T} U'_{t+h,T}) H'_{D\ell} \\ &\quad = H_{D\ell} \Gamma_U(u_\ell, h) H'_{D\ell} \int_{-1}^{x \wedge \widetilde{x}} a(v) \widetilde{a}(v)\, dv + O\left(\frac{|h|+1}{n} \|\Gamma_U(u_\ell, h)\| + \frac{|h|+1}{T} \right), \end{aligned}$$

apart from the aggregate local-covariance error $\omega_\Gamma(b)$. The loading-shift remainder is summable because Lemma A.11 gives $\sum_h (1+|h|)\|\Gamma_U(u_\ell,h)\| < \infty$. The triangular-array remainders accumulated under the exponential taper are $O(L_T^2/T)$. Finally,

$$\sum_h |1-\kappa_T(h)|\,\|\Gamma_U(u_\ell,h)\| \le \frac{1}{L_T}\sum_h |h|\,\|\Gamma_U(u_\ell,h)\| = O(L_T^{-1}),$$

where $1-e^{-y} \le y$. Combining these bounds with the centered result proves (A.35).

For distinct grid points, the calendar distance between the two windows is at least cT. Every nonzero term in the cross-window lag-window form then has multiplier covariance at most $\exp(-cT/L_T)$. Cauchy–Schwarz and the local second-moment laws bound the remaining normalized double sum by $O_p(n)$, so the cross-window form is $O_p\{n\exp(-cT/L_T)\} = o_p(1)$. The corresponding covariance of the unmultiplied sample processes is bounded by the uniformly summable score covariance tail beyond $cT/2$, as used in Lemma A.12. This proves the last assertion. □

Lemma A.14 (Conditional Gaussian equicontinuity)**.** *For any process in the finite collection used by the quartic projection, let $\mathbb{G}^*_{T\ell}$ denote its polygonally interpolated dependent-multiplier prefix process. Under Assumption 4.3, for every $\epsilon,\eta>0$,*

$$\lim_{\delta\downarrow 0}\limsup_{T\to\infty} P\left[P^*\left\{\sup_{|z-z'|\le\delta}\|\mathbb{G}^*_{T\ell}(z)-\mathbb{G}^*_{T\ell}(z')\|>\epsilon\right\}>\eta\right]=0,$$

uniformly over the fixed grid.

Proof. Lemma A.13, applied to an increment loading, gives on events whose probability tends to one

$$d_T(z,z')^2 := E^*\|\mathbb{G}^*_{T\ell}(z)-\mathbb{G}^*_{T\ell}(z')\|^2 \le C\left\{|z-z'|+a_T^D\right\},\qquad a_T^D=\frac{1}{n}+\left(\frac{L_T}{\sqrt{n}}\right)^{2r/(r+4)}=o(1). \quad \text{(A.39)}$$

Here and below C may be replaced by a tight random constant; restricting to events on which it is bounded makes the displayed deterministic notation valid. The stronger rate $a_T^D\log T=o(1)$ follows from $L_T(\log T)^2/\sqrt{n}=o(1)$ and $2r/(r+4)>1$.

To make the chaining step explicit, partition $[-1,1]$ at dyadic mesh widths 2^{-j} and stop at $J_T = \lfloor \log_2\{(a_T^D)^{-1}\}\rfloor$. Conditional Gaussian tail bounds and a union bound over the at most 2^{j+1} adjacent increments give, for every $c_0>0$,

$$P^*\left\{\max_k |\Delta_{j,k}\mathbb{G}^*_{T\ell}|>c_0 2^{-j/4}\right\}\le C2^j\exp(-cc_0^2 2^{j/2}),\qquad j\le J_T.$$

The right side is summable in j, and its tail from any fixed level J tends to zero as $J\to\infty$. Below the last dyadic scale, apply (A.39) directly to all pairs of local mesh points at distance at most $2a_T^D$. There are fewer than T^2 such pairs, so a second Gaussian union bound shows that their largest oscillation is

$$O_p^*\left\{(a_T^D\log T)^{1/2}\right\}=o_p(1).$$

The coarse dyadic chain, the last-scale bound, and linear interpolation prove the stated conditional asymptotic equicontinuity. Polynomial loadings are bounded on $[-1,1]$, and the number of processes and grid points is fixed, so the argument is uniform over every process needed by the quartic projection. □

Lemma A.15 (Conditional dependent-multiplier limit)**.** *Let the ideal local score processes be multiplied by the stationary Gaussian sequence in* (4.16)*. Conditionally on the data, their endpoints and polygonally interpolated prefix paths converge in probability to an independent copy of the limit in Lemma A.12. The conclusion is unchanged after fitted scores replace ideal scores, after kernel-weighted score centering, after centering and standardizing the realized multiplier sequence, and after the quartic projection is recomputed within each draw.*

Proof. For a bounded local loading $g_{t\ell}(z)$ that contains the kernel, a prefix indicator, and any required polynomial coordinate, consider

$$\mathbb{G}^*_{T\ell}(z) = n^{-1/2} \sum_t g_{t\ell}(z) H_{D\ell} U_{t,T} \xi_{t,T}.$$

Conditional on the data, every finite collection of these variables is Gaussian. It is therefore enough to prove convergence of its conditional covariance and asymptotic equicontinuity in the associated conditional semimetric.

The AR(1) construction gives

$$E^*(\xi_{t,T}\xi_{s,T}) = \varrho_T^{|t-s|} = \exp(-|t-s|/L_T).$$

Hence the conditional covariance of two prefix coordinates is

$$\frac{1}{n} \sum_{t,s} g_{t\ell}(z) g_{s\ell}(z') H_{D\ell} U_{t,T} U'_{s,T} H'_{D\ell} \exp(-|t-s|/L_T). \tag{A.40}$$

This is an exponential lag-window estimator of the covariance of the local weighted score process.

Lemma A.13 gives uniform convergence of this conditional covariance to the covariance in Lemma A.12; it also makes cross-window covariance negligible. Lemma A.14 supplies conditional asymptotic equicontinuity. Conditional Gaussianity and these two lemmas establish the ideal-score result.

We next transfer the result to the feasible fitted score. The exponential Toeplitz covariance matrix $R_T = (\exp[-|t-s|/L_T])_{t,s}$ satisfies

$$\|R_T\|_{\text{op}} \le \max_t \sum_s \exp(-|t-s|/L_T) \le C L_T. \tag{A.41}$$

For a normalized fitted-minus-ideal multiplier sum at any one prefix, this gives

$$E^* \left\| \sqrt{n} \sum_t (\widehat{\phi}_{t\ell} - \phi_{t\ell}) \xi_{t,T} \right\|^2 \le C L_T n \sum_t \|\widehat{\phi}_{t\ell} - \phi_{t\ell}\|^2 = o_p(1),$$

where the final step follows from (A.26) and (4.17). More explicitly, the variance bound holds for every one of the at most T prefixes. A conditional Gaussian union bound therefore gives

$$\max_{m \le N_\ell} \left\| \sqrt{n} \sum_{i \le m} (\widehat{\phi}_{t_{\ell,i},\ell} - \phi_{t_{\ell,i},\ell}) \xi_{t_{\ell,i},T} \right\| = O_p^* \left[\left\{ L_T \log T \, n \sum_t \|\widehat{\phi}_{t\ell} - \phi_{t\ell}\|^2 \right\}^{1/2} \right] = o_p(1).$$

The last equality uses $n \sum_t \|\widehat{\phi}_{t\ell} - \phi_{t\ell}\|^2 = O_p\{(r_T^s)^2\}$ and the four rates in (4.17). Multiplication by each of the five bounded powers of the local coordinate obeys the same bound. A projected prefix is a raw prefix minus a bounded partial quartic-moment matrix times these five full-window sums, so the transfer remains uniform after the draw-specific projection.

The kernel-weighted fitted-score mean is $O_p(n^{-1/2} + b^3 + \rho_T)$ after the local polynomial decomposition in Lemma A.8 and the centering bound in Lemma A.9. By (A.41), uniformly over prefixes, subtracting this mean changes a normalized multiplier process by

$$O_p^*\Big[\|\overline{a}_\ell\|\{L_T \log T\}^{1/2}\Big] = o_p(1).$$

Indeed, $L_T \log T\{n^{-1}+b^6+\rho_T^2\} = o(1)$ under (4.17) and the ridge condition. The five polynomially weighted processes obey the same bound. Finally,

$$E^*\left(T^{-1}\sum_t \xi_{t,T}\right)^2 = O(L_T/T),$$

and the sample second moment of the stationary Gaussian sequence is $1 + O_p\{(L_T/T)^{1/2}\}$. Subtracting its realized mean and dividing by its realized standard deviation therefore changes every process by $o_p(1)$. This proves all stated transfers. □

Proof of Theorem 4.10. The local design laws, influence expansion, leverage-adjustment bound, and quartic annihilation in Propositions 4.3 and 4.5 use unconditional score centering and the physical-dependence maximal inequality; the martingale restriction enters only the specialized functional central limit theorem. Lemma A.12 supplies its dependent-score replacement. The deterministic projection and variance-time calculations in Lemmas A.6 and A.7, together with the centering result in Lemma A.9, are unchanged, so the continuous mapping theorem gives the sample limits based on G_ℓ^D and Z_ℓ^D.

Lemma A.15 gives the matching conditional multiplier limit. The denominators are positive almost surely, and the continuity argument in Lemma 4.6 depends only on positive definiteness of the finite-dimensional Gaussian covariance; it therefore applies with (4.19). Conditional weak convergence implies convergence of the method-specific multiplier quantile. The band statement then follows exactly as in the proof of Theorem 4.8. □

Proof of Corollary 4.12. The global design and fitted moment-covariance laws of large numbers use Assumption 4.1(ii)–(iii) and do not require a martingale score. The finite sequence of GMM weight updates therefore converges to the same lag-zero population weight $\mathcal{S}_{K_0}^{-1}$. Expanding the final first-order condition around γ_0 gives

$$\sqrt{T}(\widehat{\gamma}_{K_0} - \gamma_0) = \mathcal{H}_{K_0}^0 T^{-1/2}\sum_{t=1}^T\{U_{t,T} \otimes b_{K_0}(u_t)\} + o_p(1).$$

An m-dependent approximation and blocking argument identical to the one in Lemma A.12, now with full-sample deterministic loading $b_{K_0}(u_t)$, gives a Gaussian limit with covariance $\mathcal{L}_{K_0}$. Premultiplication by $\mathcal{H}_{K_0}^0$ gives the displayed covariance in the corollary. The expansion is $O_p(1)$ at the root-T scale, and hence

$$\sqrt{n_T}\max_{\ell\le M}\|L_{K_0}(u_\ell)(\widehat{\gamma}_{K_0} - \gamma_0)\| = O_p\{(n_T/T)^{1/2}\} = O_p(\sqrt{b_T}) = o_p(1).$$

This is the only rate from the restricted estimator needed in the path-test limit. □

B Additional Monte Carlo evidence

This appendix supplements the main Monte Carlo tables with critical-value and bandwidth sensitivity, denominator behavior, multiplier simulation error, score-transfer comparisons, one-at-a-time projection and ordering comparisons, and finite-sample evidence for serially correlated instrumented scores.

B.1 Projected-score scale benchmark

For reference, define

$$\widehat{V}_{\ell j}^{W} = \left\{ \sum_{t=1}^{T} \left(\widehat{\phi}_{t\ell,j}^{Q4} \right)^2 \right\}^{1/2}.$$

The benchmark replaces the self-normalizer by $\widehat{V}_{\ell j}^{W}$ while retaining the local estimate, fitted score, quartic projection, and common multiplier draws.

Table B.1: Self-normalizers and a projected-score scale benchmark

Design	Method	Coverage	Band width	Constancy reject	Linear-path reject
iid regressor, constant path	Adjusted range	0.940	1.207	0.023	0.017
iid regressor, constant path	Quadratic	0.957	1.258	0.033	0.010
iid regressor, constant path	Projected-score scale	0.933	0.930	0.060	0.030
iid regressor, linear path	Adjusted range	0.933	1.215	0.070	0.057
iid regressor, linear path	Quadratic	0.927	1.260	0.073	0.057
iid regressor, linear path	Projected-score scale	0.927	0.932	0.087	0.050
iid regressor, sinusoidal path	Adjusted range	0.937	1.205	0.753	0.333
iid regressor, sinusoidal path	Quadratic	0.960	1.255	0.757	0.327
iid regressor, sinusoidal path	Projected-score scale	0.893	0.924	0.930	0.557
iid overidentified IV, constant path	Adjusted range	0.963	1.606	0.013	0.013
iid overidentified IV, constant path	Quadratic	0.953	1.665	0.023	0.023
iid overidentified IV, constant path	Projected-score scale	0.953	1.149	0.023	0.020
GARCH errors, constant path	Adjusted range	0.930	2.267	0.060	0.063
GARCH errors, constant path	Quadratic	0.940	2.299	0.063	0.070
GARCH errors, constant path	Projected-score scale	0.890	1.719	0.087	0.077

Notes: All methods use the same simulated samples, one-step identity-weight local-quadratic estimates, leverage-adjusted fitted scores, evaluation grid, and Gaussian multiplier draws. Adjusted-range and quadratic inference differ only in the functional applied to the quartic-projected partial-sum path. The projected-score scale benchmark uses the square root of the sum of squared quartic-projected score contributions and the same multiplier maximum.

Table B.1 shows why band width alone is not an adequate comparison criterion. The projected-score benchmark is narrower than both self-normalizers, but coverage falls to 0.893 for the iid sinusoidal path and 0.890 under GARCH errors. Its larger rejection frequencies in those rows therefore accompany lower simultaneous coverage.

B.2 Finite-sample critical-value sensitivity

Table B.2 uses four separately seeded null designs with 1,000 replications per design. The unscaled asymptotic rule has coverage between 0.926 and 0.966. Multiplying the conditional quantile by 1.05 moves the coverage range to 0.944–0.972 and keeps null rejection at or below 0.045 in these designs. The main tables nevertheless use the unscaled theorem-based rule; all larger factors report finite-sample sensitivity and the corresponding reduction in null and alternative rejection probabilities. The GARCH row is outside the maintained moment conditions and is not used to select a critical value.

Table B.2: High-replication null calibration at $T = 500$

	Coverage				Constancy reject				Linear-path reject			
Design	1.00	1.05	1.10	1.15	1.00	1.05	1.10	1.15	1.00	1.05	1.10	1.15
Exactly identified, iid regressor	0.928	0.947	0.965	0.973	0.050	0.034	0.023	0.016	0.043	0.030	0.020	0.015
Overidentified IV	0.966	0.972	0.985	0.989	0.023	0.012	0.011	0.010	0.022	0.016	0.011	0.010
Exactly identified, AR(1) regressor	0.940	0.963	0.977	0.988	0.035	0.026	0.014	0.009	0.026	0.016	0.011	0.006
GARCH(1,1) error	0.926	0.944	0.960	0.972	0.057	0.045	0.029	0.021	0.054	0.044	0.024	0.016

Notes: Every row uses a constant coefficient path, $N_{\mathrm{MC}} = 1{,}000$ independently generated samples, $B = 499$ common iid Gaussian multiplier draws, and $b_T = (1.5/\sqrt{12})T^{-1/5}$. Column headings are scalar multiples of the same within-sample multiplier critical value; each scalar is applied to coverage and both tests. The samples and seeds are separate from the nine-design experiment in Table 2. The 1.00 columns implement the theorem-based procedure; the larger factors are finite-sample sensitivity calculations rather than selected nominal critical values. The GARCH design is a heavy-tail stress test outside the maintained moment conditions.

Table B.3: Finite-sample sensitivity to scalar multiples of the multiplier critical value

	Coverage				Constancy reject				Linear-path reject			
Design	1.00	1.05	1.10	1.15	1.00	1.05	1.10	1.15	1.00	1.05	1.10	1.15
iid regressor, constant path	0.940	0.973	0.987	0.993	0.023	0.017	0.013	0.007	0.017	0.010	0.003	0.003
iid regressor, linear path	0.933	0.950	0.960	0.967	0.070	0.050	0.043	0.033	0.057	0.043	0.033	0.027
iid regressor, sinusoidal path	0.937	0.957	0.973	0.980	0.753	0.673	0.613	0.493	0.333	0.270	0.230	0.180
iid overidentified IV, constant path	0.963	0.970	0.983	0.987	0.013	0.013	0.010	0.007	0.013	0.013	0.010	0.010
iid overidentified IV, sinusoidal path	0.973	0.973	0.983	0.990	0.430	0.347	0.280	0.233	0.190	0.127	0.097	0.070
AR(1) regressor, constant path	0.933	0.947	0.960	0.970	0.063	0.047	0.030	0.023	0.053	0.027	0.020	0.020
AR(1) regressor, sinusoidal path	0.947	0.960	0.963	0.980	0.803	0.740	0.687	0.563	0.410	0.347	0.270	0.220
GARCH errors, constant path	0.930	0.953	0.957	0.967	0.060	0.043	0.040	0.033	0.063	0.050	0.040	0.030
AR(1) regressor, localized path	0.940	0.960	0.970	0.987	0.083	0.067	0.043	0.027	0.053	0.033	0.017	0.010

Notes: Every column applies the indicated scalar to the same within-sample multiplier critical value for bands and both tests. The raw 1.00 columns implement the asymptotic procedure; the remaining columns are finite-sample sensitivity calculations.

Table B.3 changes only the scalar multiplying the raw conditional multiplier quantile. Within each replication, the indicated scalar is applied to the band, constancy test, and linear-path test. Consequently, a gain in coverage is displayed together with its associated change in null rejection and power. These calculations are not used to redefine the asymptotic critical value.

B.3 Bandwidth constant

Table B.4: Sensitivity to the bandwidth constant at $T = 500$

Path	c_b	b_T	Coverage	Band width	Constancy reject	Linear-path reject
Constant path	1.00	0.083	0.945 (0.016)	1.568	0.045 (0.015)	0.030 (0.012)
Constant path	1.50	0.125	0.950 (0.015)	1.198	0.040 (0.014)	0.030 (0.012)
Constant path	2.00	0.167	0.955 (0.015)	0.998	0.040 (0.014)	0.025 (0.011)
Sinusoidal path	1.00	0.083	0.920 (0.019)	1.570	0.625 (0.034)	0.275 (0.032)
Sinusoidal path	1.50	0.125	0.930 (0.018)	1.207	0.725 (0.032)	0.345 (0.034)
Sinusoidal path	2.00	0.167	0.940 (0.017)	1.007	0.845 (0.026)	0.400 (0.035)

Notes: Results use 200 common samples, $B = 499$ iid Gaussian multiplier draws, and the unscaled multiplier critical value. For a given path and replication, all bandwidth constants use the same observations and multiplier draws. Every bandwidth equals $(c_b/\sqrt{12})T^{-1/5}$, and all 25 evaluation windows remain inside the sample. The same critical value is used for coverage and both path tests. Coverage and rejection entries include binomial Monte Carlo standard errors in parentheses.

Table B.4 varies only c_b; the bandwidth order remains $T^{-1/5}$, and this comparison uses the unscaled multiplier critical value throughout. Under the constant path, coverage stays between 0.945 and 0.955, while constancy rejection stays between 0.040 and 0.045. Under the sinusoidal path, a larger c_b narrows the reported bands and increases power because the diffuse signal is estimated from a larger local window. Coverage rises from 0.920 at $c_b = 1.00$ to 0.940 at $c_b = 2.00$. At the same time, constancy and linear-path power rise from 0.625 and 0.275 to 0.845 and 0.400, respectively. The common-sample design separates these changes from Monte Carlo variation across data sets or multiplier draws.

B.4 Denominators and critical values

Table B.5: Adjusted-range numerical diagnostics

Design	Mean critical value	Denominator p10	Denominator median	Shared critical-value error
iid regressor, constant path	4.565	0.058	0.073	0e+00
iid regressor, linear path	4.567	0.058	0.071	0e+00
iid regressor, sinusoidal path	4.567	0.058	0.073	0e+00
iid overidentified IV, constant path	5.576	0.041	0.055	0e+00
iid overidentified IV, sinusoidal path	5.558	0.041	0.054	0e+00
AR(1) regressor, constant path	4.560	0.051	0.066	0e+00
AR(1) regressor, sinusoidal path	4.558	0.054	0.065	0e+00
GARCH errors, constant path	4.651	0.080	0.104	0e+00
AR(1) regressor, localized path	4.561	0.050	0.064	0e+00

Notes: For each replication, the minimum adjusted-range denominator is taken over all displayed grid points and parameter coordinates. The table reports the tenth percentile and median of that minimum. Shared critical-value error is the largest absolute difference, over all replications, between the band critical value and the critical values used for the constancy and linear-path tests.

The denominator quantiles in Table B.5 are strictly positive in every design. In every replication, the band and the two path tests use the same critical value.

B.5 Multiplier simulation size

Table B.6: Nested multiplier-pool comparison under the constant path

B	Mean critical value	Mean $\lvert\Delta q\rvert$	90th pct. $\lvert\Delta q\rvert$	Coverage	Constancy reject	Linear-path reject
199	4.557	0.139	0.258	0.950 (0.015)	0.030 (0.012)	0.035 (0.013)
499	4.571	0.077	0.160	0.950 (0.015)	0.035 (0.013)	0.030 (0.012)
1,999	4.574	0.036	0.070	0.955 (0.015)	0.035 (0.013)	0.025 (0.011)
4,999	4.577	0.000	0.000	0.955 (0.015)	0.030 (0.012)	0.020 (0.010)

Notes: The experiment uses 200 common iid-regressor constant-path samples at $T = 500$. For each sample, one pool of 4,999 Gaussian multiplier draws is generated; the smaller-B critical values use prefixes of that same pool. $\lvert\Delta q\rvert$ is the absolute difference from the 4,999-draw critical value. Coverage and rejection entries include binomial Monte Carlo standard errors in parentheses. The same critical value is used for the band and both tests.

Table B.6 identifies the effect of B by holding the simulated sample and the multiplier pool fixed. Moving from 499 to 4,999 draws changes the mean critical value from 4.571 to 4.577; the mean absolute within-sample difference is 0.077. Coverage is 0.950 at 499 draws and 0.955 at 4,999 draws, well within Monte Carlo uncertainty. The main outer simulation therefore uses 499 draws, while the empirical calculations use the larger pool.

B.6 True-residual substitution and fitted-score transfer

Table B.7: Finite-sample transfer from true-residual to leverage-adjusted fitted scores

Design	T	Score-array distance	Projected-path distance	Energy ratio	Normalizer ratio	Max leverage
iid, constant	250	0.373	0.276	1.123	1.074	0.389
iid, constant	500	0.190	0.205	1.058	1.041	0.264
iid, constant	1000	0.096	0.150	1.034	1.022	0.181
iid, constant	2000	0.052	0.112	1.021	1.013	0.116
iid, sinusoidal	250	0.381	0.277	1.126	1.083	0.390
iid, sinusoidal	500	0.172	0.200	1.062	1.042	0.266
iid, sinusoidal	1000	0.098	0.152	1.037	1.023	0.182
iid, sinusoidal	2000	0.061	0.120	1.020	1.010	0.115
GARCH, constant	250	1.431	0.590	1.120	1.054	0.370
GARCH, constant	500	0.679	0.470	1.065	1.034	0.238
GARCH, constant	1000	0.364	0.371	1.034	1.012	0.169
GARCH, constant	2000	0.179	0.294	1.020	1.012	0.113

Notes: Results use 100 replications at each sample size and evaluation points $u = 0.3, 0.5, 0.7$. Score-array distance is the mean of $Tb_T \max_{\ell,j} \sum_t (\widehat{\phi}_{t\ell j} - \phi_{t\ell j})^2$. Projected-path distance is the mean of the maximum quartic-projected partial-sum difference multiplied by $\sqrt{Tb_T}$. Energy ratio is $\sum \widehat{\phi}^2 / \sum \phi^2$, and Normalizer ratio is the mean fitted-to-true-residual adjusted-range ratio. By retaining the estimated local loading and equivalent-kernel weights, this comparison isolates residual substitution and leverage correction. Table B.8 makes the separate comparison with the complete population-ideal score. The bandwidth is $b_T = (1.5/\sqrt{12})T^{-1/5}$.

The experiment uses the same bandwidth order and local-quadratic estimator as the main experiment. The benchmark replaces the fitted residual by the true structural innovation while retaining the estimated local loading and estimated equivalent-kernel weights. The first two measures compare these true-residual and leverage-adjusted fitted-score arrays before and after quartic projection. The score-energy and denominator ratios have benchmark values equal to one. Their movement toward one, together with declining leverage, checks the residual-substitution component of Lemma A.8.

Table B.8 then replaces the benchmark by the complete ideal contribution used in the theorem: the true innovation, known population moment matrix, analytic local-quadratic equivalent kernel, and its implied

Table B.8: Finite-sample transfer from the population-ideal score to the leverage-adjusted fitted score

Design	T	Score distance	Projected-path distance	Expansion remainder	Energy ratio	Normalizer ratio	Max leverage
iid, constant	250	0.586	0.361	0.510	1.255	1.146	0.389
iid, constant	500	0.291	0.268	0.424	1.123	1.079	0.264
iid, constant	1000	0.143	0.194	0.260	1.073	1.043	0.181
iid, constant	2000	0.086	0.151	0.195	1.036	1.020	0.116
iid, sinusoidal	250	0.587	0.356	0.536	1.196	1.126	0.390
iid, sinusoidal	500	0.291	0.266	0.375	1.115	1.076	0.266
iid, sinusoidal	1000	0.163	0.209	0.275	1.069	1.042	0.182
iid, sinusoidal	2000	0.092	0.156	0.232	1.033	1.018	0.115
GARCH, constant	250	3.129	0.874	1.401	1.343	1.173	0.370
GARCH, constant	500	1.487	0.659	1.105	1.145	1.083	0.238
GARCH, constant	1000	0.830	0.520	0.784	1.100	1.056	0.169
GARCH, constant	2000	0.388	0.398	0.576	1.039	1.028	0.113

Notes: Results use 100 replications at each sample size and evaluation points $u = 0.3, 0.5, 0.7$. The ideal contribution uses the true innovation, the known population moment matrix D, the analytic Epanechnikov Q_2, and the corresponding ideal equivalent-kernel variance-time index. Score distance is $Tb_T \max_{\ell,j} \sum_t (\widehat{\phi}_{t\ell j} - \phi_{t\ell j})^2$; projected-path distance is the maximum quartic-projected prefix difference multiplied by $\sqrt{Tb_T}$; and expansion remainder is $\sqrt{Tb_T}$ times the maximum absolute difference between the local estimate and its truth-plus-ideal-score expansion. Energy and normalizer ratios divide the leverage-adjusted fitted-score quantities by their population-ideal counterparts. The bandwidth is $b_T = (1.5/\sqrt{12})T^{-1/5}$.

variance-time index. In all three designs, the scaled score distance, projected-prefix distance, and influence-expansion remainder decline as T increases. The fitted-to-ideal energy and adjusted-range ratios move toward one, as does maximum leverage toward zero. Convergence is slower under conditional heteroskedasticity, but each reported measure moves in the theoretically predicted direction between $T = 250$ and $T = 2000$.

B.7 Projection and variance-time comparisons

Table B.9 changes one component at a time on common samples and multiplier draws. In Panel A, the feasible-to-ideal projected-prefix distance under iid errors falls from about 2.63 without projection to 0.65 after removing a constant, 0.50 after quadratic projection, and 0.37–0.38 after quartic projection. Under GARCH errors, the same sequence is 5.03, 1.56, 1.19, and 0.93. This monotone pattern directly supports the quartic score-transfer argument; coverage, width, and rejection frequencies need not be monotone because each projection also changes the finite-sample denominator.

Panel B holds the quartic projection fixed and changes only the ordering used inside the self-normalizer. Calendar time and squared original-kernel weights give nearly identical results in these symmetric interior windows. Replacing them with squared equivalent-kernel variance time raises GARCH simultaneous coverage from 0.890 to 0.935 and moves its constancy rejection frequency from 0.070 to 0.055. The corresponding iid changes are smaller and the power comparisons have mixed signs. These changes show that squared equivalent-kernel variance time improves variance matching in this heavy-tail stress design; the iid and power comparisons are small or mixed.

B.8 Serially correlated instrumented scores

Table B.10 separates three sources of finite-sample variation: serial dependence, replacement of the ideal score by a fitted score, and the multiplier or sieve calibration. No single specification delivers comparable coverage in all exact and overidentified designs at both values of ρ. The oracle score remains sensitive to the dependence calibration. A same-window fitted score introduces additional error when serial dependence is strong, a deleted local fit improves some cells but overcovers others, and the score sieve is conservative in this experiment. The dependent-multiplier theorem is a first-order result and admits a range of diverging correlation lengths. The table shows that their finite-sample behavior depends on score construction and

Table B.9: One-at-a-time comparisons of score projection and variance-time ordering

Design	Specification	Coverage	Width	Constancy	Linearity	Min. denominator	Prefix distance
Panel A: projection, with squared equivalent-kernel variance time							
iid, constant	None	0.920	1.235	0.050	0.045	0.504	2.623
iid, constant	Constant	0.925	1.240	0.055	0.045	0.504	0.648
iid, constant	Quadratic	0.935	1.231	0.040	0.040	0.467	0.505
iid, constant	Quartic	0.935	1.206	0.045	0.045	0.463	0.369
iid, sinusoidal	None	0.940	1.227	0.755	0.370	0.517	2.631
iid, sinusoidal	Constant	0.945	1.235	0.745	0.375	0.515	0.662
iid, sinusoidal	Quadratic	0.950	1.223	0.705	0.310	0.487	0.499
iid, sinusoidal	Quartic	0.970	1.206	0.745	0.330	0.469	0.378
GARCH, constant	None	0.915	2.184	0.085	0.105	0.680	5.029
GARCH, constant	Constant	0.945	2.384	0.055	0.040	0.687	1.565
GARCH, constant	Quadratic	0.925	2.286	0.060	0.075	0.668	1.187
GARCH, constant	Quartic	0.935	2.229	0.055	0.075	0.613	0.933
Panel B: variance-time ordering, with quartic projection							
iid, constant	Calendar time	0.915	1.125	0.040	0.030	0.503	0.369
iid, constant	Squared original-kernel weights	0.915	1.126	0.040	0.030	0.503	0.369
iid, constant	Squared equivalent-kernel weights	0.935	1.206	0.045	0.045	0.463	0.369
iid, sinusoidal	Calendar time	0.940	1.123	0.795	0.405	0.527	0.378
iid, sinusoidal	Squared original-kernel weights	0.940	1.125	0.790	0.405	0.524	0.378
iid, sinusoidal	Squared equivalent-kernel weights	0.970	1.206	0.745	0.330	0.469	0.378
GARCH, constant	Calendar time	0.890	2.025	0.070	0.085	0.683	0.933
GARCH, constant	Squared original-kernel weights	0.890	2.036	0.070	0.085	0.683	0.933
GARCH, constant	Squared equivalent-kernel weights	0.935	2.229	0.055	0.075	0.613	0.933

Notes: Results use $T = 500$, $N_{\mathrm{MC}} = 200$, $B = 199$, and common samples and Gaussian multiplier draws across specifications. All local estimates use the leverage-adjusted fitted score and $b_T = (1.5/\sqrt{12})T^{-1/5}$. Min. denominator is the tenth percentile of the replication-level minimum adjusted-range denominator multiplied by $\sqrt{Tb_T}$. Prefix distance is the mean maximum feasible-minus-population-ideal projected-prefix difference multiplied by $\sqrt{Tb_T}$. For constant designs, both rejection columns estimate size; for the sinusoidal design, both estimate power.

Table B.10: Finite-sample sensitivity under serially correlated instrumented scores

				Exactly identified		Overidentified	
Score and calibration	T	N_{MC}	B	$\rho = 0.3$	$\rho = 0.6$	$\rho = 0.3$	$\rho = 0.6$
Oracle score, exponential $c_\ell = 4$	500	200	199	0.930	0.905	0.960	0.985
Same-window leverage-adjusted score, exponential $c_\ell = 4$	500	100	199	0.880	0.760	0.900	0.930
Wider deleted fit, exponential $c_\ell = 4$	500	100	199	0.940	0.880	0.980	0.970
Oracle score, diagonal AR sieve	500	50	199	0.980	1.000	0.980	0.980
Oracle score, exponential $c_\ell = 4$	1000	100	199	0.950	0.930	0.950	0.980
Same-window leverage-adjusted score, Bartlett $c_\ell = 2$	1000	300	499	0.867	0.823	0.963	0.923
Deleted fit, Bartlett $c_\ell = 2$	1000	300	499	0.933	0.953	0.993	0.963
Oracle score, Bartlett $c_\ell = 2$	1000	300	499	0.920	0.907	0.977	0.960

Notes: Entries are simultaneous coverage probabilities for the constant path on the 25-point grid. All local estimators use $b_T = (1.5/\sqrt{12})T^{-1/5}$. Exponential and Bartlett lengths are $L_T = \lfloor c_\ell T^{1/4} \rfloor$. The wider deleted fit estimates each target residual after deleting $|s - t| \leq \lceil T^{1/4} \rceil$ and uses a pilot bandwidth $1.5b_T$; the coefficient estimator remains unchanged. The first row combines two independently seeded experiments. The exponential and Bartlett rows use growing lengths of theorem-admissible order; the diagonal autoregressive sieve is included as an alternative dependent-score calibration.

Table B.11: Comparison with the local blockwise wild bootstrap

Path	Method	Bootstrap draws	Coverage	Width
Constant	Adjusted range	499	0.903 (0.017)	1.311 (0.007)
Constant	LBWB	1299	0.953 (0.012)	0.726 (0.004)
Sinusoidal	Adjusted range	499	0.860 (0.020)	1.306 (0.007)
Sinusoidal	LBWB	1299	0.920 (0.016)	0.776 (0.004)

Notes: Both methods target simultaneous coverage over 25 dates and two coefficients in the serially correlated regression design with $T = 500$, $b_T = (1.5/\sqrt{12})T^{-1/5}$, and AR(1) error coefficient 0.3. Each row uses $N_{\mathrm{MC}} = 300$. The adjusted-range procedure uses the local-quadratic estimator and a dependent Gaussian multiplier with $L_T = 4$. The local blockwise wild bootstrap (LBWB) follows the public implementation of Lin et al. (2025): it uses local-linear estimation, pilot bandwidth $2b_T^{5/9}$, local residual blocks of length 12, and equal-tail calibration over the same 50-coordinate target. This is a comparison of complete procedures, not the matched self-normalizer comparison in Table 3. Parentheses contain Monte Carlo standard errors.

dependence strength. Accordingly, the main numerical analysis concerns martingale instrumented scores, while these results describe the supplementary serial-dependence extension.

Table B.11 compares the adjusted-range implementation with the local blockwise wild bootstrap of Lin et al. (2025) in the design with $\rho = 0.3$. The public blockwise procedure has higher coverage and narrower bands for both paths in this experiment. It also uses a different local-polynomial order, pilot bandwidth, block rule, and equal-tail joint calibration, so the comparison concerns complete procedures rather than the self-normalizing functional alone. The paper's finite-sample benchmark is the martingale-score implementation; the blockwise comparison evaluates the separate problem of inference with serially dependent scores.

C Additional empirical analysis

C.1 Coordinates driving the industry tests

Table C.1: Coordinates attaining the industry-level path-test statistics

	Constancy			Linear path		
Industry	Component	Date	Statistic	Component	Date	Statistic
NoDur	SMB beta	1978-07	7.687	HML beta	1991-02	6.832
Durbl	HML beta	1985-06	7.127	HML beta	1985-06	9.189
Manuf	Market beta	2011-04	3.707	Market beta	2011-04	4.562
Enrgy	HML beta	2011-04	6.679	HML beta	2011-04	9.031
HiTec	Market beta	2005-08	10.353	Market beta	2005-08	8.641
Telcm	Market beta	1977-04	7.878	Market beta	1977-04	6.612
Shops	HML beta	2002-06	7.094	HML beta	2002-06	6.991
Hlth	HML beta	1990-07	6.750	HML beta	2016-04	6.203
Utils	HML beta	2011-04	6.363	HML beta	2011-04	6.130
Other	Market beta	2005-08	11.059	Market beta	2005-08	8.391

Notes: For each industry and restriction, the table reports the coefficient and date attaining the maximum absolute fitted-path deviation divided by its adjusted-range denominator. Table 7 compares these statistics with industry-specific and system-wide critical values.

Table C.1 decomposes each omnibus industry statistic. Its maximum may be attained by alpha or by any of the three factor loadings, and the maximizing coordinate can differ between constancy and linearity. Reporting the coefficient and date prevents an industry-wide exceedance from being interpreted as uniform instability in every exposure. HML attains five of the ten constancy maxima and six of the ten linear-path maxima. Market beta attains four maxima for each restriction; the remaining constancy maximum is the NoDur SMB loading. Several of the strongest market-beta deviations occur in August 2005, while the HML deviations are spread across the 1980s, 1990s, and post-2000 period.

C.2 Score dependence and multiplier sensitivity

The dependent-multiplier calculation draws the stationary Gaussian AR(1) sequence in (4.16). Table 6 uses $L_T = \lfloor 4T^{1/4} \rfloor = 20$ for the main dependent comparison. The additional calculations below use $L_T = \lfloor T^{1/4} \rfloor = 5$ and fixed lengths. The block comparison assigns one independent Gaussian draw to every nonoverlapping block of L consecutive months. Each dependent or block multiplier sequence is centered and rescaled to unit sample variance, and the same sequence is used at every evaluation date and for every coefficient. These calculations alter only the multiplier dependence; the local estimates, fitted scores, variance-time indices, and observed self-normalizers are unchanged.

Table C.2: HiTec: self-normalizer and multiplier sensitivity

Specification	Critical value	Band width	Constancy	Linear path	B
Adjusted range, dependent multipliers, $L_T = 5$	5.868	1.867	10.353*	8.641*	1,999
Quadratic, dependent multipliers, $L_T = 5$	23.147	1.964	38.701*	32.303*	1,999
Adjusted range with iid multipliers	5.693	1.811	10.353*	8.641*	1,999
Quadratic with iid multipliers	22.313	1.893	38.701*	32.303*	1,999
Adjusted range, dependent multipliers, $L = 4$	5.777	1.838	10.353*	8.641*	1,999
Adjusted range, block multipliers, $L = 4$	5.783	1.840	10.353*	8.641*	1,999
Adjusted range, dependent multipliers, $L = 8$	6.106	1.943	10.353*	8.641*	1,999
Adjusted range, block multipliers, $L = 8$	6.203	1.973	10.353*	8.641*	1,999
Adjusted range, dependent multipliers, $L = 12$	6.152	1.957	10.353*	8.641*	1,999
Adjusted range, block multipliers, $L = 12$	6.251	1.989	10.353*	8.641*	1,999

Notes: The first two rows compare the adjusted-range and matched quadratic self-normalizers under identical dependent multiplier draws. The next two rows give the iid comparison; remaining rows vary the fixed dependent-multiplier or nonoverlapping block length L. An asterisk marks a statistic exceeding the corresponding unscaled critical value. All rows use $B = 1{,}999$ so that the self-normalizer and multiplier comparisons have the same draw budget. These calculations are supplementary to the $B = 4{,}999$ joint comparison in Table 6.

The first two rows of Table C.2 use identical dependent multiplier draws and therefore isolate the adjusted-range and quadratic functionals. The next two rows provide the iid comparison. The remaining rows vary fixed values of L for the dependent and block constructions. These unscaled calculations supplement the $B = 4{,}999$ iid and growing-length comparison in the main text and show how the critical value changes with the dependence span.

D Additional empirical industry-portfolio paths

This appendix reports the componentwise local-GMM paths for the nine Fama–French industry portfolios not displayed in the main text. The construction is identical to the HiTec analysis: the dependent variable is the industry excess return; the regressors and instruments are a constant, market excess return, SMB, and HML; and the gray dashed curves are simultaneous finite-grid adjusted-range bands based on the unscaled iid system-wide critical value. Dotted curves with open circles and long-dashed curves with crosses show the fitted constant and linear paths. The evaluation grid contains 71 points from 0.15 to 0.85 of the full sample.

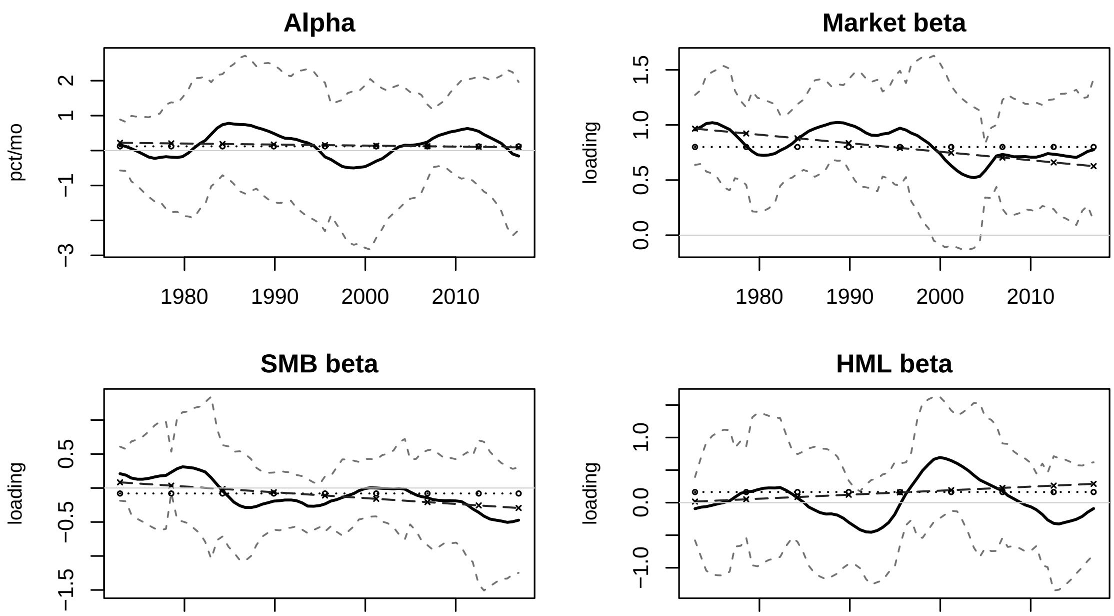

Figure 3: Time-varying local-GMM factor exposures for the NoDur industry portfolio

Notes: Solid curves connect fixed-grid local-GMM estimates; gray dashed curves connect simultaneous finite-grid adjusted-range band endpoints. Dotted curves with open circles show the fitted constant paths; long-dashed curves with crosses show the fitted linear paths. The bands are not continuum confidence bands. The sample is July 1963 to May 2026. Returns are measured in percent per month.

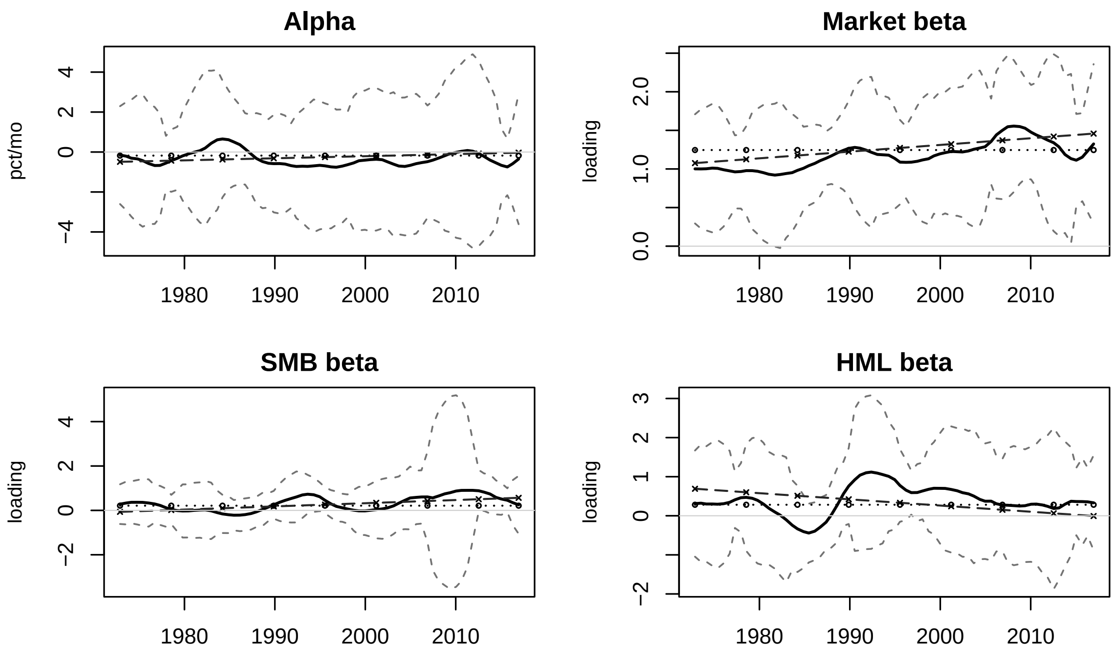

Figure 4: Time-varying local-GMM factor exposures for the Durbl industry portfolio

Notes: Solid curves connect fixed-grid local-GMM estimates; gray dashed curves connect simultaneous finite-grid adjusted-range band endpoints. Dotted curves with open circles show the fitted constant paths; long-dashed curves with crosses show the fitted linear paths. The bands are not continuum confidence bands. The sample is July 1963 to May 2026. Returns are measured in percent per month.

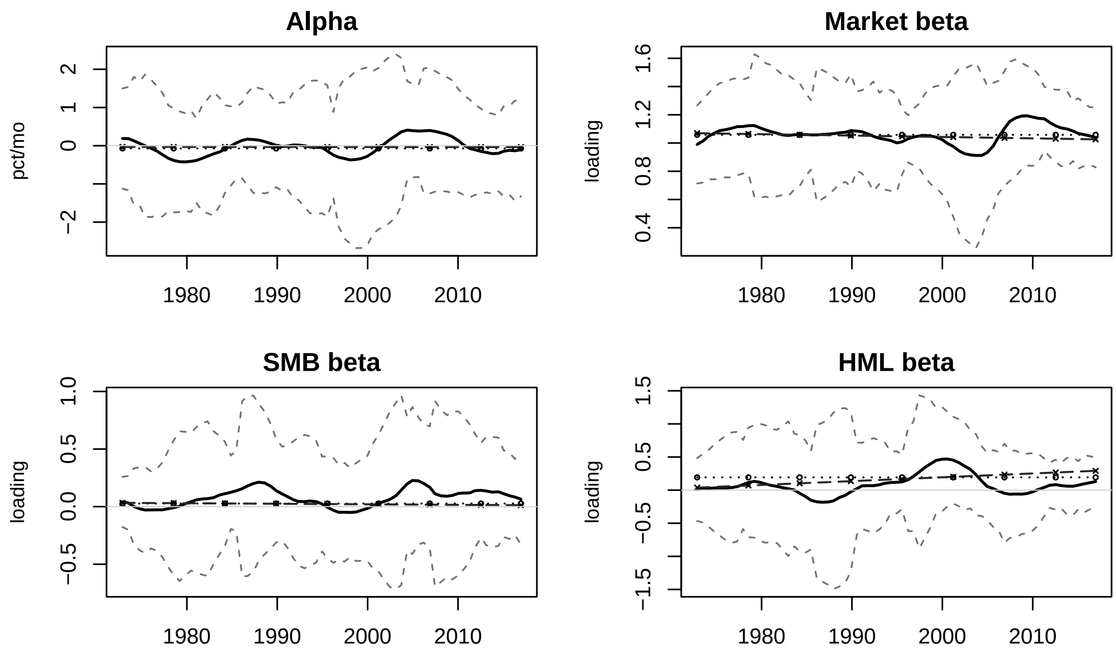

Figure 5: Time-varying local-GMM factor exposures for the Manuf industry portfolio

Notes: Solid curves connect fixed-grid local-GMM estimates; gray dashed curves connect simultaneous finite-grid adjusted-range band endpoints. Dotted curves with open circles show the fitted constant paths; long-dashed curves with crosses show the fitted linear paths. The bands are not continuum confidence bands. The sample is July 1963 to May 2026. Returns are measured in percent per month.

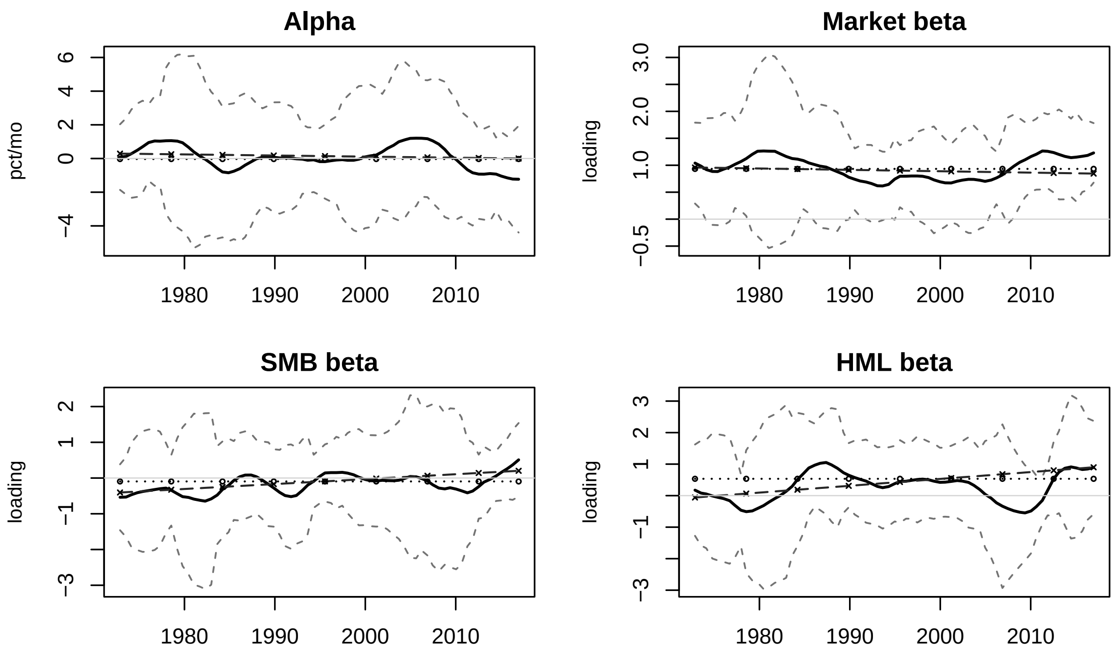

Figure 6: Time-varying local-GMM factor exposures for the Enrgy industry portfolio

Notes: Solid curves connect fixed-grid local-GMM estimates; gray dashed curves connect simultaneous finite-grid adjusted-range band endpoints. Dotted curves with open circles show the fitted constant paths; long-dashed curves with crosses show the fitted linear paths. The bands are not continuum confidence bands. The sample is July 1963 to May 2026. Returns are measured in percent per month.

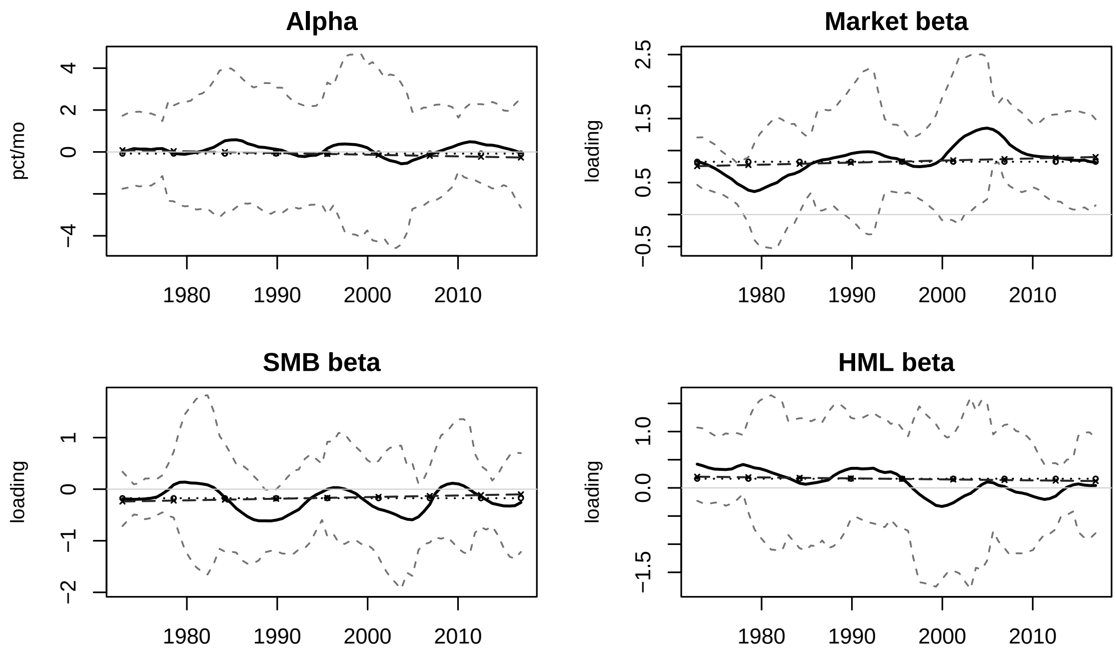

Figure 7: Time-varying local-GMM factor exposures for the Telcm industry portfolio

Notes: Solid curves connect fixed-grid local-GMM estimates; gray dashed curves connect simultaneous finite-grid adjusted-range band endpoints. Dotted curves with open circles show the fitted constant paths; long-dashed curves with crosses show the fitted linear paths. The bands are not continuum confidence bands. The sample is July 1963 to May 2026. Returns are measured in percent per month.

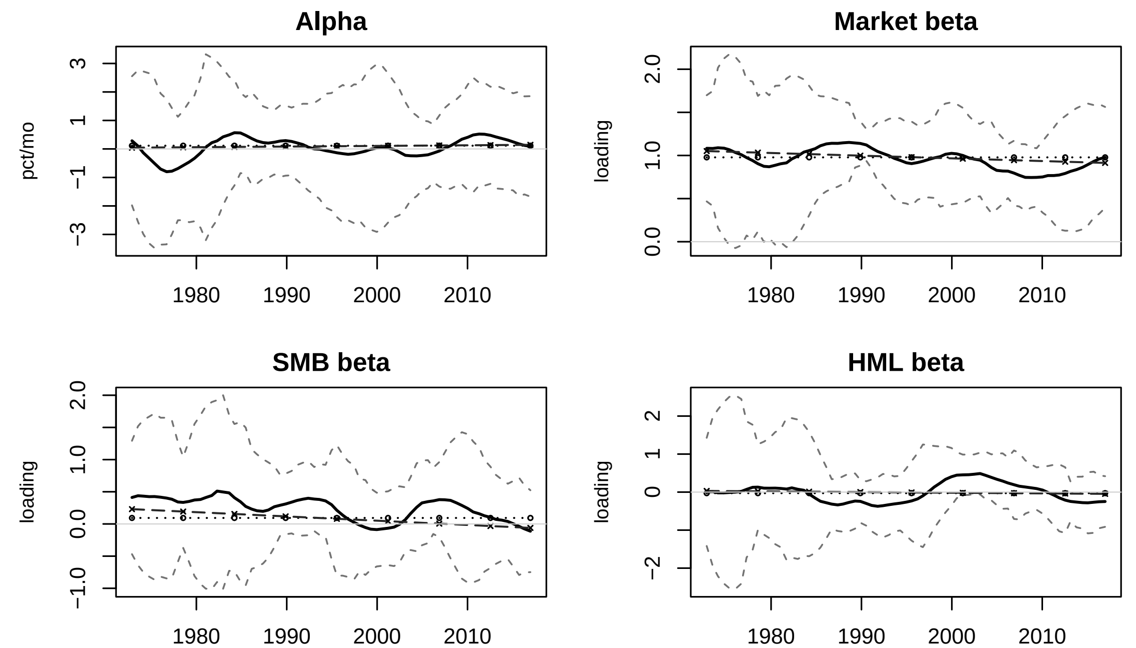

Figure 8: Time-varying local-GMM factor exposures for the Shops industry portfolio

Notes: Solid curves connect fixed-grid local-GMM estimates; gray dashed curves connect simultaneous finite-grid adjusted-range band endpoints. Dotted curves with open circles show the fitted constant paths; long-dashed curves with crosses show the fitted linear paths. The bands are not continuum confidence bands. The sample is July 1963 to May 2026. Returns are measured in percent per month.

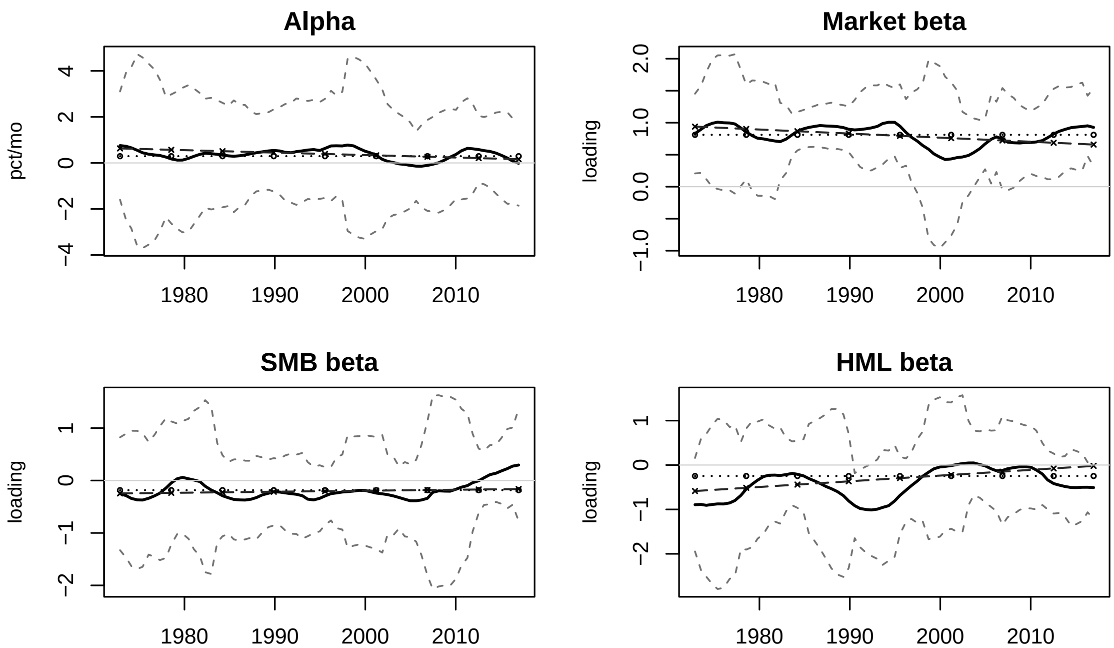

Figure 9: Time-varying local-GMM factor exposures for the Hlth industry portfolio

Notes: Solid curves connect fixed-grid local-GMM estimates; gray dashed curves connect simultaneous finite-grid adjusted-range band endpoints. Dotted curves with open circles show the fitted constant paths; long-dashed curves with crosses show the fitted linear paths. The bands are not continuum confidence bands. The sample is July 1963 to May 2026. Returns are measured in percent per month.

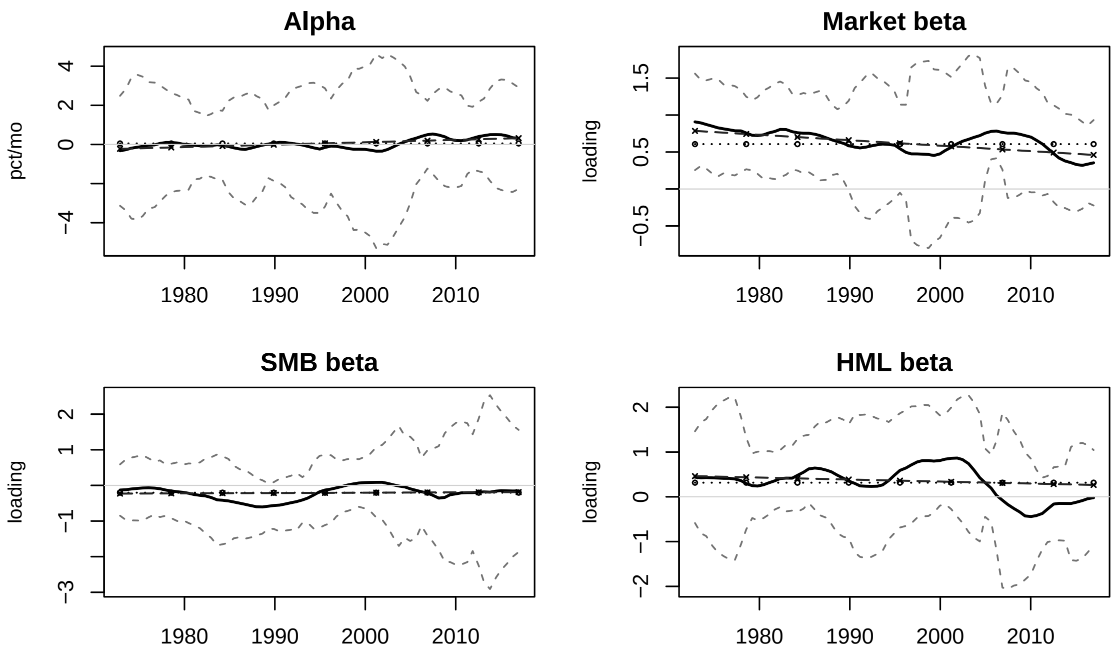

Figure 10: Time-varying local-GMM factor exposures for the Utils industry portfolio

Notes: Solid curves connect fixed-grid local-GMM estimates; gray dashed curves connect simultaneous finite-grid adjusted-range band endpoints. Dotted curves with open circles show the fitted constant paths; long-dashed curves with crosses show the fitted linear paths. The bands are not continuum confidence bands. The sample is July 1963 to May 2026. Returns are measured in percent per month.

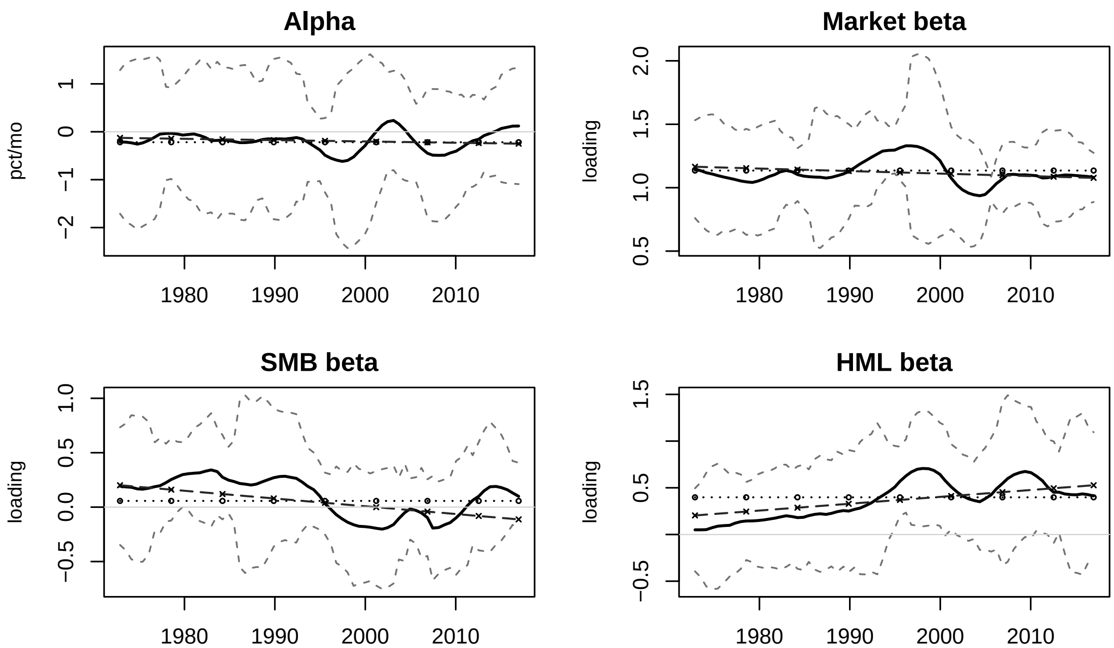

Figure 11: Time-varying local-GMM factor exposures for the Other industry portfolio

Notes: Solid curves connect fixed-grid local-GMM estimates; gray dashed curves connect simultaneous finite-grid adjusted-range band endpoints. Dotted curves with open circles show the fitted constant paths; long-dashed curves with crosses show the fitted linear paths. The bands are not continuum confidence bands. The sample is July 1963 to May 2026. Returns are measured in percent per month.