Sector Correspondence Signals for US–Japan Industry Lead-Lag

Taro Kanatani* — Shiga University
Yoichi Yoshida — Shiga University

July 23, 2026

Abstract

We show that a simple Sector Correspondence Signal (SCS), built only on a fixed US–Japan industry mapping, outperforms signals based on principal component analysis (PCA) in gross open-to-close paper portfolios for cross-border industry rotation between non-overlapping equity markets. Starting from the subspace-regularized PCA of Nakagawa et al. (2026), we separate two ingredients of that framework: covariance-based factor extraction and taxonomy-based industry information. A long-window plain PCA keeps the covariance ingredient only, dropping the taxonomy prior; SCS keeps the taxonomy ingredient only, mapping US sector returns directly into Japanese industries without any covariance estimation. The long-window plain PCA already improves on the benchmark; SCS improves further on every metric we consider, including risk-adjusted return, drawdown, and the Carhart α. A Model Confidence Set test (MCS; Hansen et al. 2011) retains the SCS family in the superior set across all risk-aversion levels and test statistics we examine. Net-of-cost evaluation is left for future work.

Keywords: lead-lag, industry rotation, sector correspondence, US–Japan equity, PCA, open-to-close.
JEL classification: G14, G15, G11.

1. Introduction

The US and Japanese equity markets—both among the world's largest and core allocations for many investors—have non-overlapping trading hours. Industry-specific information incorporated into US sector prices by the NYSE close (a sharp decline in IT, a surge in energy, an outperformance in banks, and so on) becomes observable to Tokyo investors before the next Tokyo open and tradable in Japan from that session. How should this US industry information be translated into the Japanese 17-industry space, and which industries should be longed or shorted in Tokyo that day? This is a concrete decision problem faced daily by investors engaged in cross-border allocation.

We answer this question with two cross-market industry lead-lag signals built on a universe of 11 US industry ETFs and 17 Japanese industry ETFs (28 names in total).

*Email: `t-kanatani@biwako.shiga-u.ac.jp`

Both signals take the subspace-regularized principal component analysis (PCA) framework of Nakagawa et al. (2026)—hereafter PCA_SUB—as a starting point. PCA_SUB combines short- and long-window cross-market industry covariance with a coarse cyclical/defensive labeling carried in its prior subspace, and we isolate each role separately. The first signal, a long-window plain PCA, drops the taxonomy prior and uses only a cross-market industry covariance estimate. The second signal, the *Sector Correspondence Signal* (SCS), uses no PCA at all—no eigenvectors, no covariance matrix—and instead maps each US sector ETF's return directly to its matched Japanese industry through a fixed Global Industry Classification Standard (GICS)-based correspondence between the 11 US sectors and the 17 Japanese industries. SCS is therefore not a PCA variant; it asks whether taxonomy information, supplied directly as a finer US–Japan industry correspondence rather than as the coarse cyclical/defensive labeling used in PCA_SUB, can serve as a signal on its own. Both signals target the cross-sectional ranking of Japanese industry open-to-close returns and are executed intraday at the morning auction and the afternoon close. Relative to prior work on index-level spillovers or within-market industry lead-lag, our setting simultaneously satisfies (i) *cross-border, cross-market* scope, (ii) a *daily* horizon, and (iii) *cross-sectional dispersion at the industry level*.

Our setting builds on three strands of prior work. First, within-market lead-lag arises from two complementary mechanisms: Lo and MacKinlay (1990) model how nonsynchronous trading mechanically generates cross-autocorrelations (originally at the size dimension in the US), while Hong and Stein (1999) model gradual information diffusion as a behavioral source; the industry and economic-link evidence of Hou (2007); Cohen and Frazzini (2008); Hong et al. (2007); Menzly and Ozbas (2010) builds on both. Second, international market-level return and volatility spillovers (Hamao et al. 1990; Lin et al. 1994; Karolyi and Stulz 1996) extend these mechanisms across borders and document transmission across Tokyo, London, and New York, with the non-overlapping trading hours making the first mechanism especially direct. At the country-index level, Eun and Shim (1989) first document via VAR that innovations in the U.S. market transmit to other national markets, and Rapach et al. (2013) extend this line by showing that lagged U.S. returns predict non-U.S. country returns, attributing the finding to the gradual-diffusion mechanism above. We extend this line to cross-sectional industry-level dispersion within a single non-U.S. market. Third, low-dimensional linear prediction of high-dimensional return systems— factor-augmented forecasting (Bai and Ng 2008), shrinkage covariance estimation (Ledoit and Wolf 2004), and subspace-regularized PCA (Nakagawa et al. 2026)—provides methods for extracting cross-sectional signals from many-asset panels; our plain-PCA variants build directly on these tools, and the industry-momentum benchmark of Moskowitz and Grinblatt (1999) serves as a control (§ 4.6).

Our contribution is twofold. First, the two-signal framing isolates two complementary information sources: a long-window plain PCA uses only covariance information, while SCS uses only the GICS-based industry taxonomy. Each is independently useful (positive gross R/R), and their gross return streams have low mutual R^2 ($< 5\%$). Second, and more surprising, the simple taxonomy signal outperforms the covariance-based PCA family in gross paper-portfolio performance under the same evaluation convention as the PCA_SUB benchmark of Nakagawa et al. (2026): SCS attains gross R/R ≈ 3.60 (MDD $\approx 6.62\%$) and an annualized Carhart $\alpha = 32.4\%$ ($t = 10.20$), against R/R ≈ 2.16 for the best plain PCA variant (PCA1750) and ≈ 1.91 for the PCA_SUB reference. A Model Confidence Set analysis (Hansen et al. 2011) over a seven-strategy candidate set (the plain-PCA window-length sweep plus the SCS family) retains the SCS family under every loss choice, the long-window PCA1750 under the $T_{\max}$ statistic at $\eta \leq 5$ (consistent with the window-length finding of § 4.2), and eliminates the PCA family entirely under the more restrictive T_R statistic. The two families also remain stable across regime splits, consistent with a taxonomy-based prior in V_0 that does not require covariance re-estimation.

The signal is robust across subperiods and retains incremental alpha against alternative benchmarks (§ 4.5, § 4.6): the SCS-family lead is preserved across COVID and interest-rate-regime splits, and SCS retains a large risk-adjusted α that is not absorbed by PCA1750, industry momentum, and short-term reversal jointly. The present paper focuses on signal construction and gross-return properties; net-of-cost evaluation is beyond the scope of this Note.

The remainder is organized as follows: § 2 and § 3 present the data and methods, § 4 reports the empirical results (including subperiod stability and incremental-alpha tests), and § 5 offers concluding remarks.

2. Data and Universe

The main universe consists of $N = 28$ ETFs: 11 US-side Select Sector SPDR ETFs tracking the S&P 500 sector indices (XLB, XLC, XLE, XLF, XLI, XLK, XLP, XLRE, XLU, XLV, XLY) and 17 Japan-side NEXT FUNDS TOPIX-17 industry ETFs (tickers 1617–1633); the GICS–TOPIX-17 correspondence is in Table I. Because XLC (listed 2018-06) and XLRE (listed 2015-10) begin midway through the sample, the filter requires ≥ 6 valid US ETFs, with missing returns zero-imputed and pairwise correlation for C_t (variance floor 10^{-10}). The Japanese Fama–French factors (Kenneth R. French data library) used for the FF3/Carhart regressions of § 4.3 are supplementary inputs.

Throughout the paper, t indexes a Japanese trading day. We use two return definitions: the close-to-close return $r^{cc}_{i,t} = P^{\text{close}}_{i,t}/P^{\text{close}}_{i,t-1} - 1$ on all 28 names (used for PCA factor estimation and benchmark construction), and the open-to-close return $r^{oc}_{j,t} = P^{\text{close}}_{j,t}/P^{\text{open}}_{j,t} - 1$ on Japanese names ($j \in J$) only (the strategy execution return). For US instruments, $r^{cc}_{i,t-1}$ denotes the most recent US session that has closed by the morning auction of Japanese day t (formalized in § 3). Strategies are executed as open-to-close intraday trades on day t in Tokyo—opened at the morning auction (*itayose*) and closed at the afternoon close—conditional on US sector information available before that Tokyo open. Confining execution to r^{oc} excludes the overnight close-to-open jump r^{co}, which is realized before the strategy can act on the latest US signal and is non-executable on Tokyo quotes; it also removes overnight-gap risk and carry components (*gyaku-nichibu*, stock-lending fees, margin interest). Reported gross OC returns are measured at official auction/close prints, not at executable bid/ask quotes; net-of-cost evaluation is beyond the scope of this Note.

The raw data span 2010-01-01 to 2025-12-31, restricted to US–Japan common trading days with valid Japanese OC returns. The backtest start depends on the estimation window L: 2014 for $L = 60$, 2015 for $L = 1{,}250$, and 2018 for $L = 1{,}750$. Window-length comparisons are aligned to a *common evaluation period* (2021-05-27 to 2025-12-30, $n = 987$ trading days), the longest range over which all L values can be evaluated. The construction of the PCA_SUB prior subspace V_0 and prior correlation matrix C_0 (Nakagawa et al. (2026)) is detailed in Appendix A.

3. Methodology

This section presents two cross-market industry lead-lag signals that draw on two distinct types of information.

A long-window plain PCA uses the covariance channel: from a long rolling window of returns, we estimate the joint US–Japan correlation matrix, extract the leading K principal components, and project the latest US-side factor scores onto the Japanese-side loadings to rank Japanese industries (§ 3.1). The two variants PCA1750 and PCA1250 use $L = 1{,}750$ and $L = 1{,}250$ trading days, respectively.

The Sector Correspondence Signal (SCS) uses the taxonomy channel: it ignores covariance entirely and assigns each Japanese TOPIX-17 industry the most recent return of its matched US Select Sector SPDR through a fixed GICS-based correspondence between the 11 US sectors and the 17 Japanese industries (§ 3.2, Table I).

The subspace-regularized PCA (PCA_SUB) benchmark of Nakagawa et al. (2026) combines both channels: it shrinks a short-window joint correlation matrix C_t toward a static prior $C_t^{\mathrm{reg}} = (1-\lambda)C_t + \lambda C_0$, where C_0 is direction-labeled by a hand-built subspace V_0 (global, country-spread, and cyclical/defensive factors); the original hyperparameters ($L = 60$, $\lambda = 0.9$, $K = 3$) and the full construction of V_0 and C_0 are in Appendix A. PCA1750 corresponds to the no-prior limit ($\lambda = 0$) with a longer estimation window. SCS is not a special case of PCA_SUB: it replaces the coarse cyclical/defensive labels carried by V_0 with the finer US–Japan industry correspondence, used directly rather than as a prior. The portfolio construction is shared across all signals (§ 3.3).

Calendar convention. Throughout this section t indexes a Japanese trading day. The US close-to-close return $r^{cc}_{i,t-1}$ denotes the most recent US session that has closed by the morning auction of Japanese day t. Cross-market holidays are handled by matching each Japanese trading day to the latest preceding US trading day, so consecutive signal indices may skip a US calendar day.

3.1. PCA1750: long-window plain PCA (covariance side)

At each t we estimate the pairwise correlation matrix $C_t \in \mathbb{R}^{N\times N}$ from the standardized returns $z_{i,\tau} = (r^{cc}_{i,\tau} - \hat{\mu}_{i,t})/\hat{\sigma}_{i,t}$ within the window $W_t = \{t-L, \ldots, t-1\}$, $L \in \{1{,}250,\ 1{,}750\}$ (approximately 5–7 years; pairwise estimation accommodates the mid-sample listings of XLC in 2018 and XLRE in 2015). The top $K = 3$ eigenvectors of $C_t = V_t \Lambda_t V_t^\top$ are split into US and JP blocks $V^{(K)}_{U,t} \in \mathbb{R}^{N_U\times K}$ and $V^{(K)}_{J,t} \in \mathbb{R}^{N_J\times K}$. Projecting the US-side standardized return vector $z_{U,t-1}$ available at the open of Japanese day t and reconstructing through the JP-side loadings yields the signal

$$s_t^{\mathrm{PCA}} = V^{(K)}_{J,t} V^{(K)\top}_{U,t} z_{U,t-1}.$$

The instability of short-window C_t that motivates the shrinkage term in PCA_SUB is absorbed instead by lengthening the estimation window; the resulting eigenvector stability is documented in § 4.2.

3.2. SCS: sector-taxonomy side

SCS uses a fixed GICS-to-TOPIX-17 mapping (Table I) to translate each Japanese industry to the corresponding US sector ETF. The SCS score for Japanese industry j on day t is the close-to-close return of the matched US ETF $m(j)$ for the latest US session that has closed by the open of Japanese day t:

$$s_{j,t}^{\mathrm{SCS}} = r_{m(j),\,t-1}^{cc}.$$

SCS uses no covariance information at all and lets the GICS taxonomy alone determine the US-to-JP propagation. The mapping is fixed as a one-shot taxonomy-based correspondence before evaluating strategy performance and is not estimated from returns.

Table I: GICS-to-TOPIX-17 mapping table (used to construct SCS). Note: $*$ marks industries inferred from their major constituents (with some room for debate); unmarked rows are uniquely determined within GICS.

JP code	JP industry name	US ETF	GICS sector	Note
1617	Foods	XLP	Consumer Staples	
1618	Energy Resources	XLE	Energy	
1619	Construction & Materials	XLI	Industrials	$*$
1620	Raw Materials & Chemicals	XLB	Materials	
1621	Pharmaceutical	XLV	Health Care	
1622	Automobiles & Transportation Equipment	XLY	Consumer Discretionary	
1623	Steel & Nonferrous Metals	XLB	Materials	
1624	Machinery	XLI	Industrials	
1625	Electric Appliances & Precision Instruments	XLK	Information Technology	$*$
1626	IT & Services, Others	XLC	Communication Services	$*$
1627	Electric Power & Gas	XLU	Utilities	
1628	Transportation & Logistics	XLI	Industrials	
1629	Commercial & Wholesale Trade	XLI	Industrials	$*$
1630	Retail Trade	XLY	Consumer Discretionary	
1631	Banks	XLF	Financials	
1632	Financials (ex Banks)	XLF	Financials	
1633	Real Estate	XLRE	Real Estate	

For 13 of the 17 Japanese industries the correspondence is essentially determined by the industrial definitions; the four remaining industries (1619, 1625, 1626, 1629) are economically non-unique under the coarser GICS 11-sector taxonomy and are assigned based on the dominant business activities of major constituents (marked with $*$ in Table I). A robustness check that entirely neutralizes these four ambiguous cells is reported in Appendix D and confirms that the gross return level is not driven by the four ambiguous cells, although including them materially improves the risk-adjusted and drawdown profile; both the full and neutralized SCS variants are retained in the Model Confidence Set superior set of § 4.4 under every loss choice.

3.3. Portfolio construction

Both families share the same portfolio rule. The 17 Japanese industry scores $\{s_{j,t}\}$ are cross-sectionally ranked, and an equal-weighted long–short portfolio is formed by going long the top q quantile and short the bottom q quantile. We adopt $q = 0.3$, which selects $n_q = \max(1, \text{round}(qN_J)) = 5$ industries on each side under $N_J = 17$. Because SCS maps the 17 Japanese industries to only 11 US sector ETFs, all Japanese industries sharing the same US sector receive identical scores on any given day; we break the resulting ties at the cutoff in ascending TOPIX-17 ticker order ($1617 < 1618 < \cdots < 1633$), a deterministic rule that is invariant across days and independent of the realized SCS values. Trading is open-to-close intraday on day t: positions are opened at the morning auction and closed at the afternoon close. The strategy return is $R_t^{\text{strat}} = \sum_{j \in J} w_{j,t}\, r_{j,t}^{oc}$ with $w_{j,t} = n_q^{-1}[\mathbf{1}\{\text{rank}(s_{j,t}) > N_J - n_q\} - \mathbf{1}\{\text{rank}(s_{j,t}) \leq n_q\}]$. The signal $s_{j,t}$ is either $s_{j,t}^{\text{SCS}}$ or the j-th component of s_t^{PCA}.

4. Empirical Results

Unless otherwise noted, all analyses are aligned to the common evaluation period (2021-05-27 to 2025-12-30, $n = 987$ trading days)—the longest range over which all L values can be evaluated—with the baseline specification $K = 3$ and $q = 0.3$. Annualization uses the *actual day count* $n_{\text{eff}} = N/\text{yrs}$ to be consistent with the cross-market intersection constraint (only days that are trading days in both the US and Japanese markets and on which all ETF data are available): for the common evaluation period $n_{\text{eff}} = 214.8$ days/year, and for the extended evaluation period (2018-01-04 to 2025-12-30, $N = 1{,}737$) $n_{\text{eff}} = 217.5$ days/year. We define $\text{AR} = \mu \cdot n_{\text{eff}} \cdot 100$, $\text{RISK} = \sigma \cdot \sqrt{n_{\text{eff}}} \cdot 100$, $\text{R/R} = \text{AR/RISK}$, $\text{MDD} = \max_{t \leq N}\bigl(1 - W_t / \max_{s \leq t} W_s\bigr) \cdot 100$ with $W_t = \prod_{s \leq t}(1 + r_s)$ the cumulative wealth process, and $\text{Calmar} = \text{AR/MDD}$. The R/R ratio is the annualized return-to-risk ratio, equivalent to a Sharpe-type ratio of gross strategy return (no risk-free rate is subtracted because the long–short construction is dollar-neutral and overnight-free). All performance figures in this paper are reported gross of transaction costs.

4.1. Overview of main strategies

We adopt two main strategies that span the two-signal framing described in § 3: PCA1750 (the covariance-side signal) and SCS (the taxonomy-side signal). Table II reports their gross paper-portfolio performance against the PCA_SUB reference of Nakagawa et al. (2026) under the same evaluation convention used in that prior work.

Each step of the progression improves on this sample: the covariance-side PCA1750 lifts R/R from the PCA_SUB reference of 1.91 to 2.16, and the taxonomy-side SCS further lifts it to 3.60 while lowering MDD from 9.12% (PCA_SUB) and 8.28% (PCA1750) to 6.62%. The terminal wealth over the common evaluation period ($n = 987$ trading days) follows the same ordering, with SCS reaching the highest level (Figure 1).

Beyond gross performance, SCS and PCA1750 are also structurally distinct at the position level: on any given day they share on average 3.0 of their 10 directional selections (1.6 on the long side, 1.4 on the short side), close to the 2.9 (50/17) expected when two strategies pick size-5 subsets from the 17 Japanese industries independently at random. By contrast, the two plain-PCA variants PCA1750 and PCA1250 share 7.1 of 10. The

Table II: Gross paper-portfolio performance of the two main strategies and the PCA_SUB reference of Nakagawa et al. (2026) over the extended evaluation period 2018-01-04 to 2025-12-30 (US–Japan common trading days N = 1,737; n_{eff} = 217.5 days/year). The best value in each column (high AR, low RISK, high R/R, low MDD, high Calmar) among the two main strategies below the midrule is in bold.

Strategy	AR	RISK	R/R	MDD	Calmar
PCA_SUB ($L = 60$, $\lambda = 0.9$; ref.)	17.97%	9.43%	1.91	9.12%	1.97
PCA1750 (covariance-side)	18.15%	**8.41%**	2.16	8.28%	2.19
SCS (taxonomy-side)	**30.57%**	8.49%	**3.60**	**6.62%**	**4.62**

near-zero return correlation between SCS and PCA1750 documented in § 4.6 thus reflects this structural independence at the selection stage rather than a coincidence in returns.

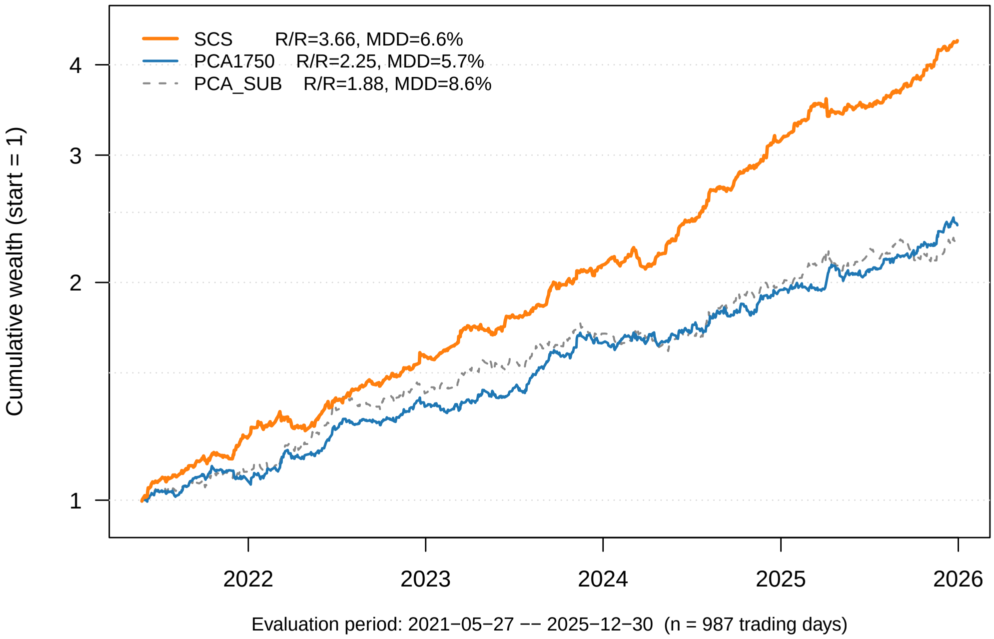

Figure 1: Cumulative wealth of the main strategies and the PCA_SUB reference (common evaluation period 2021-05-27 to 2025-12-30, n = 987 trading days; start = 1; log vertical axis). The R/R and MDD figures in the legend are computed on the actual-day basis n_{eff} = 214.8 days/year.

4.2. Robustness within the PCA family

We check the PCA1750 specification along three choices: the estimation window L, the shrinkage intensity λ, and the factor count K. Supporting tables and figures are in Appendix A–B.

The window length matters most. An 11-level sweep $L \in \{60, \ldots, 2{,}500\}$ places the maximum at L = 1,250 (R/R = 2.40) with L = 1,750 close behind (R/R = 2.34). Both are roughly threefold above the short-window PCA at L = 60 and outperform PCA_SUB on the same common period. The long-window plain PCA thus does not need the taxonomy-informed prior of PCA_SUB once the covariance estimate is stabilized by extending the window.

Shrinkage and the factor count line up with this reading. A (L, λ) grid and the $\lambda = 1$ limit confirm that pinning the loadings to the prior V_0 produces strictly inferior R/R (1.32–1.67). At fixed L, $K = 3$ is the parsimony / tail-risk best operating point. We adopt $L \in \{1{,}250, 1{,}750\}$ with $K = 3$ throughout.

The estimated eigenvectors are stable except around the COVID trough, when PC3 undergoes a single discontinuous reallocation (PC1 and PC2 hold steady throughout with $|\cos\theta(V_t, V_{t-60})| \approx 0.99$). The strategy works in both the PC3-stable and the PC3-unstable phases (Appendix A.7).

4.3. Fama–French and Carhart alpha

We regress strategy returns on the Japanese Fama–French three-factor model (MKT, SMB, HML; Ken French Data Library, Japan) and the Carhart four-factor model (adding WML), with Newey–West HAC standard errors (lag = 5). Daily strategy returns are compounded to monthly returns by calendar month (incomplete months with fewer than 15 trading days are excluded), the monthly excess return is the dependent variable, and the evaluation window is the longest sample available for each strategy (n = 90–180 months). The alpha is annualized by multiplying the monthly estimate by 12.

Table III: Alphas from Fama–French three-factor and Carhart four-factor regressions (longest sample available for each strategy; monthly excess returns as the dependent variable; Newey–West HAC, lag = 5). The best value in each column (high α, high $t(\alpha)$) is in bold. Note that n differs across strategies.

			FF3			Carhart 4		
Strategy	n	Period	α (%)	$t(\alpha)$	adj. R^2	α (%)	$t(\alpha)$	adj. R^2
SCS	96	2018-01–2025-12	**32.37**	**10.20**	0.009	**32.42**	**10.22**	0.030
MOM	180	2010-05–2025-12	2.53	0.79	0.013	2.10	0.65	0.046
PCA_SUB	180	2010-05–2025-12	16.50	5.77	−0.014	16.92	5.91	0.012
PCA1250	117	2015-10–2025-12	21.35	7.20	−0.004	21.48	7.14	−0.011
PCA1750	90	2018-01–2025-12	19.93	5.97	0.010	19.86	6.01	0.000

Excluding MOM, all four strategies exhibit highly significant FF3 alpha ($t(\alpha) \geq 5.77$); adding WML for the Carhart specification leaves both the alpha and its significance essentially unchanged. The t-statistic for SCS (10.20) is the largest across all strategies, exceeding the highest PCA-family value (PCA1250, 7.20), and the annualized alpha for SCS (32.37%) likewise exceeds every PCA-family strategy under the same intraday ETF execution layer. The adjusted R^2 is near zero throughout: at the monthly horizon the co-movement with FF/Carhart factors is extremely weak. In the Carhart β estimates (reported in Appendix C), the three PCA-family strategies show no significant loading, while SCS exhibits a significantly negative $\beta_{\mathrm{WML}} = -0.169$ ($t = -2.24$); the alpha remains essentially unchanged in moving from FF3 to Carhart (from 32.37% to 32.42%), so it is not absorbed by WML.

4.4. Model Confidence Set analysis

The SCS-vs-PCA comparison and the plain-PCA window-length selection ($L = 1{,}750$ vs other windows in the sweep of § 4.2) together raise a post-selection concern. We address it with a Model Confidence Set (MCS; Hansen et al. 2011), which identifies the subset of strategies that cannot be statistically distinguished from the best under a chosen loss function. The candidate set is seven strategies: the plain-PCA window-length sweep (PCA_SUB and plain PCA at four long windows) and the SCS family (SCS and the ambiguous-4 neutralized variant SCS_neut). We use the mean-variance utility loss across a range of risk-aversion levels and report both the $T_{\max}$ and the more restrictive T_R statistic. Full implementation details (loss formula, bootstrap, η grid, software) are in Appendix E.

Table IV: Retained set of the 90% Model Confidence Set across eight loss–statistic combinations (four risk-aversion levels $\eta \in \{0, 1, 5, 10\}$ under mean-variance utility loss, crossed with the two MCS test statistics $T_{\max}$ and T_R). Candidate set: seven strategies (PCA_SUB, plain PCA at four long windows, SCS, and SCS_neut). Sample: extended common period 2018-01-04 to 2025-12-30, $N = 1{,}737$ trading days. The SCS family is retained under every combination; PCA elimination is total under T_R at every η and four-of-five under each $T_{\max}$ variant. See Appendix E for full implementation.

Loss / Statistic	PCA strategies eliminated	SCS strategies retained	Marginal PCA
$\eta = 0$ (negative-return) / $T_{\max}$	4/5	2/2	PCA1750 ($p = 0.114$)
$\eta = 1$ / $T_{\max}$	4/5	2/2	PCA1750 ($p = 0.125$)
$\eta = 5$ / $T_{\max}$	4/5	2/2	PCA1750 ($p = 0.153$)
$\eta = 10$ / $T_{\max}$	4/5	2/2	PCA_SUB ($p = 0.163$)
$\eta \in \{0, 1, 5, 10\}$ / T_R	**5/5**	2/2	—

Table IV supports a layered reading. First, the SCS family is retained throughout: both SCS and SCS_neut survive under every $T_{\max}$ variant, and they are the only two strategies retained under the more restrictive T_R statistic. Second, the headline long-window choice PCA1750 is *also* in the superior set under $T_{\max}$ for $\eta \leq 5$, with PCA_SUB taking its place at $\eta = 10$. The marginal retention of PCA1750 by $T_{\max}$ supports the window-length finding of § 4.2: extending the estimation window beyond the short-window PCA_SUB baseline does provide a statistically meaningful improvement that survives post-selection scrutiny under the negative-return and low-η mean-variance losses. The SCS family's lead over the PCA family is invariant to loss choice, risk aversion, and test statistic: SCS achieves the smallest mean loss in every column.

4.5. Subperiod stability

We evaluate risk-adjusted returns under two economically meaningful regime splits: pre-COVID (2018-01 to 2020-02, $n = 482$) versus post-COVID (2020-04 to 2025-12, $n = 1{,}235$; excluding the 2020-03 gap), and pre-interest-rate-shift (2018-01 to 2022-11, $n = 1{,}076$) versus post-shift (2022-12 to 2025-12, $n = 661$). Both strategies achieve R/R ≥ 1.88 in every period, so the result is not a regime-dependent backfit. SCS in particular stays in 3.51–4.08 across the four subperiods, with no breakdown under either split. This stability is consistent with the two-signal framing of § 3: SCS uses a fixed taxonomy mapping in V_0, while PCA1750 relies on a covariance structure estimated from past returns.

Table V: R/R by regime split × strategy (evaluation period 2018-01-04 to 2025-12-30, $L = 1{,}750$). Within each strategy row, the best value across the four subperiod columns is in bold.

Strategy	Full	Pre-COVID	Post-COVID	Low rates	Rising rates
SCS	3.60	**4.08**	3.70	3.51	3.75
PCA1750	2.16	**2.55**	1.88	2.20	2.12

4.6. Incremental alpha against alternative benchmarks

Although SCS and PCA1750 share the same economic source—the US-to-Japan sector lead-lag—the position-level near-independence documented in § 4.1 carries over to the return space: the joint R^2 is below 5%, and each retains a large incremental α not absorbed by the other or by Japanese-domestic benchmarks. Over the extended evaluation period 2018-01-04 to 2025-12-30 ($N = 1{,}737$ daily returns), we run two joint Newey–West HAC regressions (lag = 5) in which SCS and PCA1750 are each regressed on the sibling strategy together with the traditional industry 12-1 momentum (IM; cumulative return from month $t-12$ to month $t-2$, following Moskowitz and Grinblatt 1999) and 1-month reversal (REV; return in month $t-1$). SCS retains a large incremental annualized $\alpha = 27.3\%$ ($t(\alpha) = 9.9$); PCA1750 retains $\alpha = 12.6\%$ ($t(\alpha) = 3.4$) in the reverse direction. The mutual slope on the sibling strategy is $\beta \approx 0.19$ in both directions, small but statistically significant, so the two families are complementary rather than fully orthogonal. The slopes on IM and REV are not statistically significant in either regression ($|t| < 1.6$), so Japanese-domestic momentum and reversal do not explain either strategy (Table F.1).

5. Concluding Remarks

We constructed two daily cross-market industry lead-lag signals that draw on two distinct types of information combined in the subspace-regularized PCA of Nakagawa et al. (2026): a covariance-derived plain PCA at the no-prior limit ($\lambda = 0$, no V_0) with a long estimation window, and a Sector Correspondence Signal (SCS) that uses no PCA at all and instead maps each US sector ETF's return directly to its matched Japanese industry through a fixed GICS-based correspondence. Under the same gross paper-portfolio evaluation convention as the PCA_SUB benchmark, SCS outperforms every PCA variant we consider on every metric we report, including risk-adjusted return, drawdown, and the Carhart α. A Model Confidence Set test retains the SCS family in the superior set across all risk-aversion levels and test statistics we examine. Both signals remain stable across regime splits, consistent with a sector taxonomy that is regime-invariant by construction.

One interpretation is that SCS captures, more directly than PCA, the intuitive sector-based associations through which investors trade across markets—for instance, treating US semiconductor returns as a cue for Japanese semiconductor returns. A finer correspondence, such as a firm-level matching between US and Japanese counterparts, might extend this idea further. We leave a formal mechanism analysis for future work.

We evaluate gross intraday returns only; a full implementation analysis, including bid-ask spreads, liquidity, short-sale constraints, and borrowing costs, is left for future work. Other natural extensions include other market pairs (EU–JP, EU–US) and real-time detection of PC3 structural changes. Two further directions concern the scope of the signal itself: aggregating the 17 cross-sectional signals into a market-wide forecast as a market-timing input, and extending the universe beyond industry peers (e.g., a cross-asset-class universe spanning equities, rates, FX, and commodities). Both lie outside the cross-sectional industry-peer scope on which the present results rest.

Acknowledgements

The authors used Claude (Anthropic) to assist with drafting, code review, and English-language editing during the preparation of this manuscript. All analyses, results, and final wording are the authors' own responsibility.

Appendix

This Appendix complements the main text by providing the detailed specifications, tables, figures, and sensitivity analyses that the main text compresses. It is organized into six sections aligned with the cross-references in the main text: A (PCA_SUB specification, V_0 construction, (L, λ) grid, $\lambda = 1$ limit, and eigenvector temporal stability), B (window length and factor count sensitivity), C (Fama–French / Carhart β for all strategies), D (SCS mapping robustness: ambiguous-4 neutralization), E (Model Confidence Set details: candidate set, per-strategy p-values, and implementation), and F (incremental α from mutual and benchmark regressions).

A. Subspace-regularized PCA (PCA_SUB) specification and sensitivities

This appendix specifies the subspace-regularized PCA of Nakagawa et al. (2026) ("PCA_SUB") used as the reference comparison in the main text, and provides the supporting sensitivity analyses: the two-dimensional (L, λ) grid, the $\lambda = 1$ projection limit together with the rolling long-window C_0 experiment, and the time series of adjacent-window eigenvector angles underlying the PC3 regime reallocation discussed in the main text.

A.1. Overview and parameter alignment

PCA_SUB shrinks a short-window rolling correlation matrix C_t toward a static prior correlation matrix C_0 that has been direction-labeled by an economic prior subspace $V_0 \in \mathbb{R}^{N\times 3}$:

$$C_t^{\text{reg}} = (1-\lambda)\, C_t + \lambda\, C_0, \qquad \lambda \in [0,1].$$

The three columns of $V_0 = [v_1, v_2, v_3]$ encode a global factor, a country-spread factor, and a cyclical/defensive factor, each constructed by hand from the 28 industry ETFs in the universe. The static prior C_0 is built from V_0 together with magnitudes read off a long-window training-period correlation matrix; the explicit construction is given in § A.2 below. The eigendecomposition of C_t^{reg} then provides the directions onto which the US-side return cross-section is projected to forecast the JP-side cross-section, and a long–short portfolio is formed at the cross-sectional quantile $q = 0.3$.

Table A.1 compares the hyperparameters of PCA_SUB with those of the plain long-window PCA used as the main PCA benchmark in the paper.

Table A.1: Hyperparameters of PCA_SUB (Nakagawa et al., 2026) and of the main PCA benchmark of this paper.

Parameter	PCA_SUB	Main PCA benchmark
Estimation window L	60	1,250 or 1,750
Shrinkage coefficient λ	0.9	0
Prior subspace V_0	$[v_1, v_2, v_3]$	not used
Number of factors K	3	3
Cross-sectional quantile q	0.3	0.3

The plain long-window PCA of the main text corresponds to the special case $\lambda = 0$ of PCA_SUB. In this limit the second term of the regularization equation vanishes, the prior subspace V_0 and the auxiliary matrices D_0, Δ, C_0 are not used, and the eigendecomposition is taken directly on the raw rolling correlation matrix C_t. The short-window instability that motivates the shrinkage term of Nakagawa et al. (2026) is absorbed by enlarging L to 1,250 or 1,750, and the effective number of free hyperparameters is therefore reduced from five to two (L and K). The (L, λ) grid in § A.5 below shows that this simplification dominates PCA_SUB on both R/R and maximum drawdown over the industry ETF universe.

A.2. Construction of the prior subspace V_0

The prior subspace $V_0 = [v_1, v_2, v_3] \in \mathbb{R}^{N \times 3}$ consists of three unit-orthogonal vectors corresponding, respectively, to a global factor, a country-spread factor, and a cyclical/defensive factor.

Global factor v_1. An equal-weight vector across all N names,

$$v_1 = \frac{1}{\sqrt{N}}(1, 1, \dots, 1)^\top \in \mathbb{R}^N.$$

Country-spread factor v_2. The raw vector

$$\tilde{v}_2 = \Big(\underbrace{1/\sqrt{N_U}, \dots, 1/\sqrt{N_U}}_{N_U},\ \underbrace{-1/\sqrt{N_J}, \dots, -1/\sqrt{N_J}}_{N_J}\Big)^\top$$

is orthogonalized against v_1 and normalized to unit length.

Cyclical/defensive factor v_3. A sign-coded vector $\tilde{v}_3$ is constructed from the economic classification in Table A.2, then orthogonalized against v_1, v_2 and normalized to unit length. The remaining 12 names receive a zero coefficient. The three directions are constructed by hand and held fixed over time.

Table A.2: Economic classification used to build v_3. The remaining 12 names receive a zero coefficient.

Category	Tickers	Sign
US cyclical	XLB, XLE, XLF, XLRE	+1
US defensive	XLK, XLP, XLU, XLV	−1
Japan cyclical	1618, 1625, 1629, 1631	+1
Japan defensive	1617, 1621, 1627, 1630	−1
Remaining 12	—	0

A.3. Regularized correlation matrix and the static prior C_0

The core of PCA_SUB is to shrink the rolling correlation matrix C_t toward a static direction-labeled matrix C_0:

$$C_t^{\text{reg}} = (1 - \lambda)\, C_t + \lambda\, C_0, \qquad \lambda \in [0, 1].$$

The matrix C_0 is not the rank-three outer product $V_0 V_0^\top$; it is constructed in three steps so that magnitudes are read from a training-period correlation matrix while V_0 supplies only direction labels.

Step 1. Compute the long-window correlation matrix C_{full} over the training period (2010–2014 in Nakagawa et al. 2026).

Step 2. Read the magnitude on each prior direction from C_{full}, forming the diagonal matrix

$$D_0 = \operatorname{diag}\bigl(V_0^\top C_{\text{full}} V_0\bigr) \in \mathbb{R}^{3\times3}.$$

Step 3. Rescale so that the result has unit diagonal, using $\Delta = \operatorname{diag}(V_0 D_0 V_0^\top)$:

$$C_0 = \Delta^{-1/2} V_0 D_0 V_0^\top \Delta^{-1/2}.$$

Two-sided scaling by $\Delta^{-1/2}$ ensures $\operatorname{diag}(C_0) = \mathbf{1}$ so that C_0 is a valid correlation matrix. The $\lambda = 1$ limit isolates the projector $\operatorname{eig}(C_0)$ rather than V_0 itself; the role separation between direction (V_0) and magnitude (C_{full}) is diagnosed directly in Section A.6 below.

A.4. Signal construction and numerical stability

The top $K = 3$ eigenvectors $V_t^{\text{reg},(K)}$ of C_t^{reg} are split into a US block $V_{U,t}^{\text{reg},(K)}$ and a JP block $V_{J,t}^{\text{reg},(K)}$, and the cross-market projection signal is

$$s_t = V_{J,t}^{\text{reg},(K)} \bigl(V_{U,t}^{\text{reg},(K)}\bigr)^\top z_{U,t} \in \mathbb{R}^{N_J}.$$

Portfolio construction (cross-sectional quantile $q = 0.3$, long top and short bottom, open-to-close intraday execution) is identical to the main-text specification.

Two numerical points are worth noting. First, because XLC starts in 2018 and XLRE in 2015, the rolling window contains missing observations for some names, so C_t is estimated pairwise; pairwise sample correlation matrices are not guaranteed to be positive semi-definite (Higham 2002), and we apply a positive-semi-definite projection that clips negative eigenvalues to zero and resets the diagonal to one. Because C_t^{reg} is a convex combination of two PSD matrices, C_t^{reg} is itself PSD. Second, the third prior direction v_3 is supported on only $4 + 4 + 4 + 4 = 16$ names, so Δ takes at most eight discrete values (four cyclical categories times two countries); this discreteness propagates into $\Delta^{-1/2}$ and explains the negative R/R reported at $K = 2$ in Section A.6, where the discrete defensive/cyclical labels move in the opposite direction to realized correlations over 2021–2025.

A.5. (L, λ) two-dimensional grid

We embed the PCA_SUB specification ($L = 60, \lambda = 0.9$) and our specification ($L = 1750, \lambda = 0$) in a common grid to visualize the interaction between estimation window length and the shrinkage coefficient. We run $\lambda \in \{0, 0.1, 0.3, 0.5, 0.7, 0.9, 0.95, 1.0\}$ and $L \in \{60, 1750\}$ at the common evaluation period (sixteen points, $K = 3$, $q = 0.3$, script `lambda_sweep.R`).

At $L = 1750$ the best point is $\lambda = 0$ and increasing λ generally lowers R/R: the long window provides richer information than the prior. At $L = 60$ the dependence is inverse U-shaped and peaks at $\lambda \approx 0.95$: with a short window C_t is noisy and shrinkage toward the prior helps. The two columns intersect near $\lambda \in \{0.9, 0.95\}$, where the short-window specification slightly exceeds the long-window one (R/R higher by +0.10 to +0.16). The overall optimum is $(L, \lambda) = (1750, 0)$ at R/R = 2.34 with maximum drawdown −5.5%; the PCA_SUB point $(60, 0.9)$ is dominated on both metrics (R/R = 1.88, MDD = −8.6%).

Table A.3: R/R on the (L, λ) grid ($K = 3$, common evaluation period). Column-wise maximum in bold. The cell $L = 60, \lambda = 0.9$ corresponds to the PCA_SUB specification.

λ	$L = 60$	$L = 1750$
0.00	0.67	**2.34**
0.10	0.59	1.97
0.30	0.46	2.14
0.50	0.83	2.19
0.70	1.76	1.97
0.90	1.88	1.78
0.95	**1.94**	1.78
1.00	1.68	1.79

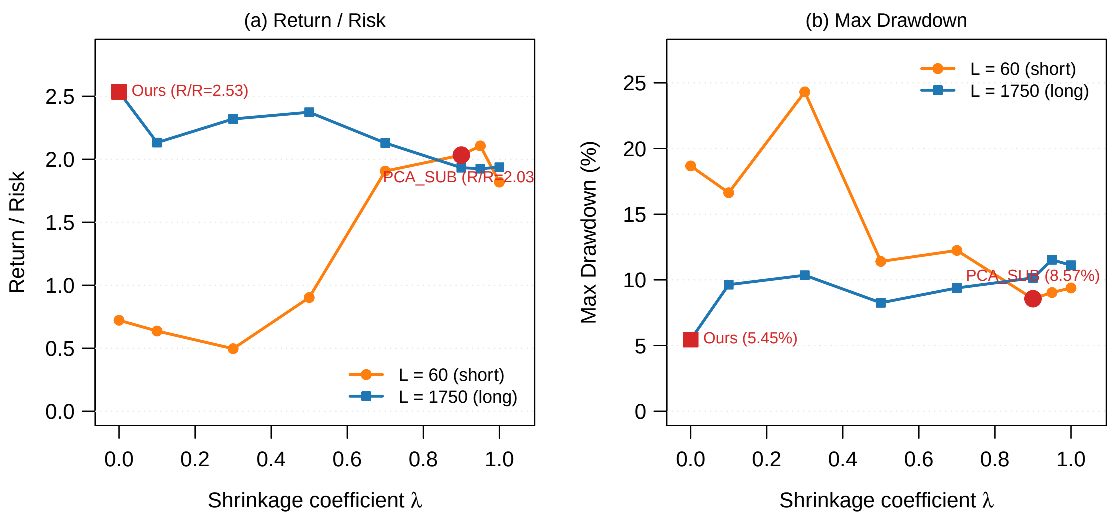

Figure A.1: Two-dimensional (L, λ) grid ($K = 3$, common evaluation period). Left: Return / Risk. Right: maximum drawdown. The orange line is $L = 60$ (short window) and the blue line is $L = 1750$ (long window). The short window peaks near $\lambda \approx 0.95$ and the long window is best at $\lambda = 0$; the two cross near $\lambda \approx 0.9$. Red markers indicate our specification (Ours, $L = 1750, \lambda = 0$) and PCA_SUB ($L = 60, \lambda = 0.9$). Our specification dominates PCA_SUB on both R/R and MDD.

A short window with strong shrinkage and a long window with no shrinkage are mutually substitutable stabilization devices. However, the long-window route is preferred on three grounds: it attains the highest R/R, the lowest MDD, and requires no hand-constructed prior subspace.

A.6. $\lambda = 1$ limit and long-window C_0 experiment

The (L, λ) grid above sweeps $\lambda \in [0, 1]$ continuously. To isolate the predictive content of the prior itself we examine the $\lambda \to 1$ limit. Note that the projection basis at $\lambda = 1$

is $\mathrm{eig}(C_0)$, not V_0 directly: V_0 specifies directions while magnitudes flow through the training-period correlation C_{full} (Section A.3).

The backtest run by `pca_sub_lambda1.R` yields R/R $= 1.32$ to 1.51 at $K = 3$, which is 56% to 65% of the corresponding PCA1750 value (R/R $= 2.34$). At $K = 1$ predictive content is very weak; at $K = 2$ the strategy R/R is *negative* (-0.4 to -0.5), because the discrete defensive/cyclical labels carried by $\Delta^{-1/2}$ move opposite to realized correlations over 2021–2025 (Section A.4).

To relax the binding of the prior to the 2010–2014 training period, we replace C_{full} by a rolling long-window correlation matrix with $L_{\mathrm{full}} \in \{1{,}250, 1{,}750\}$ (script `pca_sub_long_c0.R`). The long-window C_0 version improves but remains 60% to 71% of PCA1750 on R/R, and MDD is more than twice as deep (Table A.4).

Table A.4: Long-window C_0 experiment ($K = 3$, common evaluation period). Within the two long-window rows, column-wise best in bold (R/R larger, MDD smaller, ratio larger). Static C_0 and PCA1750 are shown for reference.

Specification	R/R	MDD	Ratio to PCA1750
Static C_0 (PCA_SUB as-is)	1.32–1.51	—	56%–65%
Long-window C_0, $L_{\mathrm{full}} = 1{,}250$	1.40	−13.1%	60%
Long-window C_0, $L_{\mathrm{full}} = 1{,}750$	**1.67**	**−12.6%**	**71%**
Plain long PCA1750	2.34	−5.5%	100%

As long as direction is pinned to V_0, updating only magnitude does not catch up to the plain long-window PCA. Section A.7 below shows that the direction of PC3 itself reorganizes around March 2020, with the PC3 carrier migrating from US sectors to JP industries; once the direction is hard-coded into V_0, a ceiling is created. This justifies our choice to recover both direction and magnitude directly from a long-window realized correlation matrix, without passing through V_0.

A.7. Adjacent-window eigenvector stability over time

We report the time series of adjacent-window angles $|\cos\theta(V_t, V_{t-60})|$ for the daily rolling PCA ($L = 1{,}750$, $K = 3$) that underlies the PC3 regime reallocation summarized in § 4.2 of the main text. The first and second principal components are extremely stable through time, with adjacent-window $|\cos\theta| \approx 0.99$ and overlap with the prior vectors v_1, v_2 of approximately 0.97. The third component, in contrast, exhibits a discrete drop to $|\cos\theta| = 0.44$ around the COVID trough of March 2020, and its overlap with v_3 (the cyclical/defensive prior) varies between 0.37 and 0.68 depending on the period. Known candidate structural events (the GICS communication-services split of 2018-09 and the inclusion of Tesla in the S&P 500 in 2020-12) are not the primary cause of the PC3 jump.

Comparing strategy performance across the PC3-stable and PC3-unstable phases (Table A.5), PCA1750 attains R/R $= 2.26$ in the stable phase and 1.35 in the unstable phase, a clear gap. PCA1250 shows the opposite pattern, with the unstable phase (R/R $= 2.46$) outperforming the stable phase (R/R $= 2.06$); the extended natural evaluation period (from 2015 onward) lets the shorter window track the direction flip in $V_t^{(3)}$ more nimbly. The strategy works in both windows.

The pattern is consistent with PC1 and PC2 carrying a stable joint US–Japan common factor structure, with PC3 supplying the marginal cross-sectional ranking information.

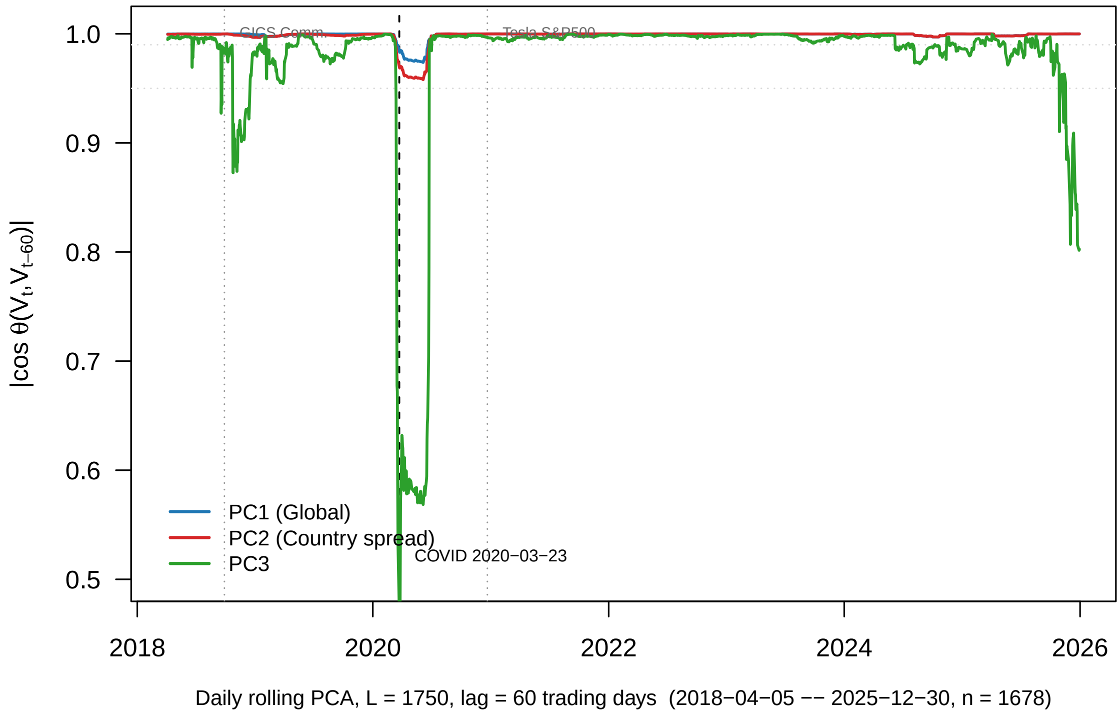

Figure A.2: Adjacent-window $|\cos\theta(V_t, V_{t-60})|$ for daily rolling PCA ($L = 1{,}750$, $K = 3$, lag $= 60$ trading days $\approx$ one quarter). PC1 (global) and PC2 (country spread) are extremely stable near 0.99. Only PC3 shows a discrete drop to 0.44 around the COVID trough of March 2020. The candidate structural events (GICS communication-services split 2018-09, Tesla S&P 500 inclusion 2020-12) are not the primary driver of the PC3 jump. Evaluation period 2018-04-05 to 2025-12-30 ($n = 1{,}678$ trading days).

Table A.5: Strategy performance in PC3-stable and PC3-unstable phases (split at the median of the adjacent-window $|\cos\theta|$). Within each row, the larger R/R is in bold.

Strategy	Stable phase R/R	Unstable phase R/R
PCA1250 (evaluation from 2015)	2.06	**2.46**
PCA1750 (evaluation from 2018)	**2.26**	1.35

Even though the PC3 carrier itself migrates from US to JP, both phases work because the PC3 direction is re-estimated daily and tracks the macro regime.

B. Window length and factor count sensitivity

This appendix supplies the sensitivity tables of plain PCA over the estimation window L and the factor count K that are summarized in § 4.2 of the main text.

B.1. Sensitivity to the estimation window L

We run plain PCA ($\lambda = 0$, $K = 3$, $q = 0.3$, all other choices as in the main text) for eleven window lengths $L \in \{60, 120, 250, 500, 750, 1000, 1250, 1500, 1750, 2000, 2500\}$ over the common evaluation period (Table B.1). The optimum is L = 1,250 (R/R = 2.40), with L = 1,750 a close second (R/R = 2.34). The short-window plain PCA $L = 60$ attains only R/R = 0.67, and our best point exceeds it by more than a factor of three. PCA_SUB rerun over the same period ($L = 60$, $\lambda = 0.9$) attains R/R = 1.88 and is dominated by the entire band $L \in \{1000, 1250, 1500, 1750, 2000\}$. Performance is short-noise dominated below $L \leq 500$ (R/R below 2.0) and deteriorates by L = 2,500 (R/R = 1.44) as the window incorporates older structural breaks. The R/R local dip at L = 1,500 is offset by the MDD optimum at L = 1,750 (5.45%), which together support our main specification $L \in \{1{,}250, 1{,}750\}$.

Table B.1: Window-length sensitivity of plain PCA ($K = 3$, $\lambda = 0$, $q = 0.3$, common evaluation period 2021-05-27 to 2025-12-30, $n = 987$). Best value per column in bold (AR larger, RISK smaller, R/R larger, MDD smaller); the PCA_SUB row is excluded from the bold competition and is reported for reference using its own parameter setting.

L	AR	RISK	R/R	MDD
60	6.09%	9.14%	0.67	18.68%
120	3.71%	9.61%	0.39	25.30%
250	12.93%	9.85%	1.31	21.48%
500	16.19%	9.80%	1.65	19.04%
750	19.11%	9.97%	1.92	10.72%
1000	19.01%	9.15%	2.08	8.46%
1250	20.95%	**8.74**%	**2.40**	6.91%
1500	18.42%	9.20%	2.00	7.36%
1750	**21.54**%	9.21%	2.34	**5.45**%
2000	18.46%	9.45%	1.95	6.35%
2500	14.04%	9.76%	1.44	8.65%
PCA_SUB ($L = 60, \lambda = 0.9$)	18.39%	9.80%	1.88	8.57%

B.2. Justification for fixing $K = 3$

Holding $L \in \{1{,}250, 1{,}750\}$ fixed and varying $K \in \{1, 2, 3, 4, 5\}$, the predictive content is negligible at $K = 1$ (R/R = −0.06 at L = 1,250 and 0.15 at L = 1,750) and remains non-operational at $K = 2$ (R/R = 0.72 at L = 1,250 and 0.80 at L = 1,750). The jump from $K = 2$ to $K = 3$ raises R/R by 1.68 at L = 1,250 and 1.54 at L = 1,750, demonstrating that the third principal component is indispensable. R/R rises marginally at $K = 4$ (L = 1,250: 2.46 vs 2.40; L = 1,750: 2.56 vs 2.34), but maximum drawdown is clearly shallower at $K = 3$ (L = 1,250: 6.91% vs 8.29%; L = 1,750: 5.45% vs 7.19%), so the Calmar ratio favors $K = 3$. By $K = 5$ the R/R falls back to 2.04 at L = 1,250 and 1.90 at L = 1,750 as short-window noise re-enters. We therefore fix $K = 3$ as the main specification on the joint grounds of tail-risk-adjusted efficiency and parsimony. As noted in § 4.2 of the main text, PC3 is itself temporally unstable, yet the decisive R/R jump from $K = 2$ to $K = 3$ suggests that the strategy dynamically picks up whichever macro factor is active at each point on the third axis.

C. Fama–French / Carhart β estimates for all strategies

This appendix provides the full Carhart four-factor (Mkt-RF, SMB, HML, WML) β estimates for the five strategies summarized in § 4.5 of the main text (Newey–West HAC standard errors with lag = 5; monthly excess returns as the dependent variable; the evaluation window is the longest sample available for each strategy). The main text reports only the alpha and adjusted R^2; Table C.1 below provides the underlying 20 β estimates with their t-statistics.

Table C.1: Carhart four-factor β estimates (Newey–West HAC, lag = 5).

Strategy	β_{Mkt}	t	β_{SMB}	t	β_{HML}	t	β_{WML}	t
SCS	+0.063	+1.10	+0.144	+1.28	+0.071	+0.80	−0.169	−2.24
MOM	−0.113	−1.73	−0.036	−0.38	+0.006	+0.07	+0.197	+2.12
PCA_SUB	−0.057	−0.87	+0.002	+0.03	−0.075	−1.13	−0.190	−1.97
PCA1250	−0.088	−0.96	−0.188	−1.98	−0.088	−1.09	−0.046	−0.34
PCA1750	−0.075	−0.79	+0.216	+1.56	−0.079	−0.95	+0.043	+0.52

Three observations emerge. First, the industry momentum benchmark MOM loads positively on WML ($\beta_{\mathrm{WML}} = +0.197$, $t = +2.12$), consistent with the conventional momentum factor and confirming that our regression apparatus correctly identifies momentum-related dependence as a sanity check. Second, the three PCA-family strategies have insignificant loadings on all four factors, so mainstream style factors do not explain their returns. Third, SCS shows a significantly negative $\beta_{\mathrm{WML}} = -0.169$ ($t = -2.24$), yet the alpha barely moves when moving from FF3 to Carhart (from 32.37% to 32.42%), so adding WML does not absorb the SCS alpha.

D. SCS mapping robustness: ambiguous-4 neutralization

Four of the seventeen Japanese industries are economically non-unique under the coarser GICS 11-sector taxonomy and are flagged with $*$ in Table I of the main text: 1619 (Construction & Materials, mapped to XLI), 1625 (Electric Appliances & Precision Instruments, mapped to XLK), 1626 (IT & Services, mapped to XLC), and 1629 (Commercial & Wholesale Trade, mapped to XLI). Although the mapping is fixed before any performance evaluation (the "one-shot" protocol described in § 3.2 of the main text), one might still worry that these four discretionary assignments are responsible for the SCS alpha.

To address this concern directly, we re-evaluate SCS after *neutralizing* the four ambiguous industries—i.e. dropping them from the signal ranking and using only the remaining thirteen industries for long/short selection (with $q = 0.30$ applied to the active count). The neutralized version is evaluated on the same extended common 1,737 trading days as Table II of the main text.

Table D.1: Performance of SCS with the four ambiguous industries (1619, 1625, 1626, 1629) included (Full) versus neutralized (Neut: removed from signal ranking, $q = 0.30$ applied to the remaining 13 industries). Same extended 1,737 trading days as Table II of the main text. Calmar = AR/MDD.

Strategy	Variant	AR	RISK	R/R	MDD	Calmar
SCS	Full	**30.57%**	**8.49%**	**3.60**	**6.62%**	**4.62**
SCS	Neut	29.14%	9.91%	2.94	7.63%	3.82

Two observations emerge. First, the gross annualized return barely moves: SCS declines only from 30.57% to 29.14% (-1.43 pp). The *level* of the SCS signal's predictive content is therefore not generated by the four ambiguous mapping cells; the unambiguous thirteen mapping cells account for nearly all of the gross alpha. Second, the risk and drawdown metrics deteriorate materially under neutralization: RISK rises from 8.49% to 9.91%, R/R falls from 3.60 to 2.94, and MDD widens from 6.62% to 7.63%. The mechanical cause is the smaller active universe (the long and short legs hold roughly four assets each instead of five), which concentrates position weight and reduces cross-sectional diversification. We therefore weaken the conclusion: the gross level of SCS alpha is not the product of discretionary choices among the four ambiguous industries, but we do not claim that those four cells are irrelevant to the risk-adjusted profile of the strategy; they materially improve diversification and drawdown control.

E. Model Confidence Set details

This appendix complements § 4.4 of the main text by reporting the full candidate-set composition, per-strategy MCS p-values across all five loss specifications, and the implementation details (bootstrap, statistics, software).

E.1. Candidate set

The seven candidates target the post-selection / data-snooping concern along two axes simultaneously: the choice of plain-PCA window length L and the SCS-vs-PCA family comparison. Five strategies span the plain-PCA window-length sweep summarized in § 4.2 of the main text—PCA_SUB (the Nakagawa et al. 2026 reference at $L = 60$, $\lambda = 0.9$, $K = 3$) and plain PCA ($\lambda = 0$, $K = 3$) at $L \in \{1{,}000, 1{,}250, 1{,}500, 1{,}750\}$—and two cover the SCS family: the headline SCS strategy and SCS_neut, the variant with the four ambiguous mapping cells neutralized (§ 3.3 of the main text; Appendix D). All strategies are evaluated on the same extended common 1,737 trading days (2018-01-04 to 2025-12-30) of the main text. Sample-period restriction to the common 1,737-day window keeps the candidates on a strictly aligned daily loss-differential basis as required by Hansen et al. (2011).

E.2. Per-strategy MCS p-values

Table E.1 reports the per-strategy MCS p-values across the five loss specifications: $T_{\max}$ at $\eta \in \{0, 1, 5, 10\}$ plus T_R at $\eta = 5$. Under the mean-variance utility loss $L_{i,t}(\eta) = -r_{i,t} + (\eta/2)r_{i,t}^2$, $\eta = 0$ recovers the negative-return criterion; $\eta > 0$ progressively penalizes variance. The MCS p-value of a strategy i is interpreted as the largest α at which i is included in the $(1 - \alpha)$ superior set; strategies with $p \geq 0.10$ are retained in the 90% MCS superior set.

Table E.1: MCS p-values (Hansen et al. 2011) across four risk-aversion levels ($\eta \in \{0, 1, 5, 10\}$) for the two MCS test statistics: $T_{\max}$ (Panel A) and the more restrictive T_R (Panel B). Loss function is the mean-variance utility $L_{i,t}(\eta) = -r_{i,t} + (\eta/2)r_{i,t}^2$. Elimination-step order is shown as superscripts in brackets (worst first); bold marks strategies in the 90% MCS superior set (MCS $p \geq \alpha = 0.10$). Candidate set: seven strategies (PCA_SUB, plain PCA at four window lengths, SCS, and SCS_neut). Sample: extended common period N = 1,737 trading days. B = 10,000 bootstrap replications; full implementation in Appendix E.3.

Strategy	$\eta = 0$	$\eta = 1$	$\eta = 5$	$\eta = 10$
Panel A: $T_{\max}$ statistic				
PCA_SUB	0.020[2]	0.020[2]	0.071[2]	**0.163**
PCA1000	0.010[1]	0.008[1]	0.011[1]	0.011[1]
PCA1250	0.020[3]	0.093[3]	0.071[3]	0.016[2]
PCA1500	0.077[4]	0.093[4]	0.071[4]	0.016[3]
PCA1750	**0.114**	**0.125**	**0.153**	0.049[4]
SCS_neut	**0.544**	**0.503**	**0.374**	**0.228**
SCS	**1.000**	**1.000**	**1.000**	**1.000**
Panel B: T_R statistic				
PCA_SUB	0.008[2]	0.007[2]	0.007[2]	0.006[2]
PCA1000	0.000[1]	0.000[1]	0.000[1]	0.000[1]
PCA1250	0.008[3]	0.011[3]	0.010[3]	0.008[3]
PCA1500	0.010[4]	0.011[4]	0.010[4]	0.008[4]
PCA1750	0.013[5]	0.013[5]	0.014[5]	0.012[5]
SCS_neut	**0.544**	**0.503**	**0.377**	**0.229**
SCS	**1.000**	**1.000**	**1.000**	**1.000**

Three patterns emerge. First, the SCS family is unambiguously retained: SCS achieves the smallest mean loss in every specification ($p = 1$ by construction), and SCS_neut achieves MCS $p \in [0.228, 0.544]$ across the four $T_{\max}$ specifications and MCS $p \in [0.229, 0.544]$ across the four T_R specifications. Second, under the more restrictive T_R statistic all five PCA candidates are eliminated uniformly across the four η values, and the elimination order (PCA1000, PCA_SUB, PCA1250, PCA1500, PCA1750) is invariant to η—a stability not present under $T_{\max}$, where $\eta = 10$ reorders the PCA family. Third, the marginal $T_{\max}$ survivor is consistent with the window-length finding of § 4.2 of the main text: PCA1750 (the main paper's chosen long window) is the lone PCA candidate retained under $T_{\max}$ for $\eta \in \{0, 1, 5\}$ at MCS $p \in [0.114, 0.153]$, and PCA_SUB takes its place only at $\eta = 10$ ($p = 0.163$). PCA1750's $T_{\max}$ retention indicates that extending the estimation window beyond the short-window PCA_SUB baseline does provide a statistically meaningful improvement that survives post-selection scrutiny under the negative-return and low-η mean-variance losses; the eliminated long-window variants ($L \in \{1{,}250, 1{,}500\}$) carry similar mean losses to PCA1750 but are eliminated earlier in the iteration sequence because the $T_{\max}$ test is sensitive to pairwise differentials, not point-estimate mean loss alone. The eliminated PCA1000 is the genuinely short long-window variant that the sweep of § 4.2 already identifies as inferior.

E.3. Implementation details

The MCS test is computed with the R package `MCS` (Bernardi and Catania 2018), which implements the procedure of Hansen et al. (2011). We use the stationary bootstrap with $B = 10{,}000$ replications and the package default for the average block length. The $T_{\max}$ statistic is the maximum standardized deviation from the average loss; the T_R statistic is the maximum pairwise relative-loss differential. The implementation seed is `42`, set before the call to `MCSprocedure` (each MCS call resamples internally, so the per-run bootstrap seeds in the package output differ across the five specifications).

F. Incremental α from mutual and benchmark regressions

This appendix reports the regression estimates underlying the incremental-α statements in § 4.5 of the main paper.

Table F.1: Incremental α from joint Newey–West HAC regressions (lag = 5; extended evaluation period 2018-01-04 to 2025-12-30, N = 1,737; α annualized with $n_{\text{eff}} = 217.5$ days per year). Each strategy (SCS or PCA1750) is regressed on the sibling strategy together with the traditional industry 12-1 momentum (IM) and 1-month reversal (REV). Newey–West t-statistics in parentheses. The mutual slope on the sibling strategy is small ($\beta \approx 0.19$) but statistically significant, so we interpret the incremental α as "not fully absorbed by the other family" rather than as full orthogonality. Slopes on IM and REV are not statistically significant ($|t| < 1.6$).

Regression	α	$t(\alpha)$	β_{sibling}	β_{IM}	β_{REV}	R^2
SCS on PCA1750, IM, REV	27.25%	9.89	0.188 (4.57)	−0.129 (−1.53)	0.054 (0.57)	0.037
PCA1750 on SCS, IM, REV	12.65%	3.39	0.184 (4.03)	−0.147 (−1.51)	0.100 (0.85)	0.037

References

Bai, J., and S. Ng, "Large Dimensional Factor Analysis," *Foundations and Trends in Econometrics*, Vol. 3, No. 2, pp. 89–163, 2008.

Bernardi, M., and L. Catania, "The Model Confidence Set Package for R," *International Journal of Computational Economics and Econometrics*, Vol. 8, No. 2, pp. 144–158, 2018.

Cohen, L., and A. Frazzini, "Economic Links and Predictable Returns," *Journal of Finance*, Vol. 63, No. 4, pp. 1977–2011, 2008.

Eun, C. S., and S. Shim, "International Transmission of Stock Market Movements," *Journal of Financial and Quantitative Analysis*, Vol. 24, No. 2, pp. 241–256, 1989.

Hamao, Y., R. W. Masulis, and V. Ng, "Correlations in Price Changes and Volatility across International Stock Markets," *Review of Financial Studies*, Vol. 3, No. 2, pp. 281–307, 1990.

Hansen, P. R., A. Lunde, and J. M. Nason, "The Model Confidence Set," *Econometrica*, Vol. 79, No. 2, pp. 453–497, 2011.

Higham, N. J., "Computing the Nearest Correlation Matrix—A Problem from Finance," *IMA Journal of Numerical Analysis*, Vol. 22, No. 3, pp. 329–343, 2002.

Hong, H., and J. C. Stein, "A Unified Theory of Underreaction, Momentum Trading, and Overreaction in Asset Markets," *Journal of Finance*, Vol. 54, No. 6, pp. 2143–2184, 1999.

Hong, H., W. Torous, and R. Valkanov, "Do Industries Lead Stock Markets?," *Journal of Financial Economics*, Vol. 83, No. 2, pp. 367–396, 2007.

Hou, K., "Industry Information Diffusion and the Lead-Lag Effect in Stock Returns," *Review of Financial Studies*, Vol. 20, No. 4, pp. 1113–1138, 2007.

Karolyi, G. A., and R. M. Stulz, "Why Do Markets Move Together? An Investigation of U.S.–Japan Stock Return Comovements," *Journal of Finance*, Vol. 51, No. 3, pp. 951–986, 1996.

Ledoit, O., and M. Wolf, "A Well-Conditioned Estimator for Large-Dimensional Covariance Matrices," *Journal of Multivariate Analysis*, Vol. 88, No. 2, pp. 365–411, 2004.

Lin, W.-L., R. F. Engle, and T. Ito, "Do Bulls and Bears Move across Borders? International Transmission of Stock Returns and Volatility," *Review of Financial Studies*, Vol. 7, No. 3, pp. 507–538, 1994.

Lo, A. W., and A. C. MacKinlay, "An Econometric Analysis of Nonsynchronous Trading," *Journal of Econometrics*, Vol. 45, No. 1-2, pp. 181–211, 1990.

Menzly, L., and O. Ozbas, "Market Segmentation and Cross-Predictability of Returns," *Journal of Finance*, Vol. 65, No. 4, pp. 1555–1580, 2010.

Moskowitz, T. J., and M. Grinblatt, "Do Industries Explain Momentum?," *Journal of Finance*, Vol. 54, No. 4, pp. 1249–1290, 1999.

Nakagawa, K., Y. Takemoto, K. Kubo, and M. Kato, "Subspace-Regularized Principal Component Analysis for US–Japan Industry Lead-Lag Investment Strategies," *Proceedings of the SIG-FIN of the Japanese Society for Artificial Intelligence*, No. FIN-036, pp. 76–83, 2026. (In Japanese.)

Rapach, D. E., J. K. Strauss, and G. Zhou, "International Stock Return Predictability: What Is the Role of the United States?," *Journal of Finance*, Vol. 68, No. 4, pp. 1633–1662, 2013.