

Geopolitical Supply-Chain Vulnerability and Relative Sectoral Equity Returns: Evidence from U.S. Strategic Dependence on China[★]

Dr Arsene Oka^{a,c,*,1}

^aHoward University, 3rd Floor 2400 Sixth Street, N.W., Washington, DC, 20059, United States of America

ARTICLE INFO

Keywords:

Strategic import vulnerability
supply-chain risk
sectoral equity returns
critical minerals
China

asset pricing

JEL classifications: F14; F51; G12;
G14; L14.

ABSTRACT

Against the background of rising geopolitical fragmentation, this paper examines whether financial markets price changes in U.S. strategic import vulnerability to China across upstream and downstream industries in the United States and China. We combine monthly Structural Strategic Import Risk Index (SSIRI) measures for 2004–2025 with publicly available U.S. and Chinese sectoral equity data. The analysis uses standardized monthly changes, Newey–West inference, extensive financial and global-risk controls, and false-discovery-rate corrections. A one-standard-deviation increase in exposure-weighted critical-minerals and advanced-materials SSIRI is associated with a 1.10 percentage-point increase in the one-month U.S. upstream-minus-downstream return spread observed two months after the signal. A related 1.07 percentage-point three-month response appears in China. Decomposition results indicate an upstream-led U.S. response and a combination of stronger upstream and weaker downstream performance in China. Exposure-weighted SSIRI is more precisely related to equity returns than the equal-family structural measure. U.S. battery-sector results are generally null, while Chinese battery results are sensitive to portfolio construction. The findings reveal a sector-specific association between strategic import vulnerability and relative equity returns.

1. Introduction

Strategic supply-chain vulnerability can affect firms differently depending on their position in the production network. When access to an imported input becomes more vulnerable, upstream producers may benefit from greater scarcity rents, stronger pricing power, increased demand for alternative supply, or expectations of policy support. Downstream firms may instead face higher input costs, narrower margins, inventory requirements, and greater uncertainty about production continuity. Equity prices provide a forward-looking measure of how investors distribute these risks and potential rents across the supply chain. Yet there is limited evidence on whether changes in observed import vulnerability are systematically associated with relative valuations across upstream and downstream industries.

This question has become increasingly important as governments and multinational firms reassess their dependence on geographically concentrated suppliers of critical minerals, battery inputs, and advanced materials. Dependence on China is especially salient because China is a major supplier and processor of inputs used in energy, transportation, electronics, and defense. A high Chinese import share does not, however, fully characterize vulnerability. The economic consequences of a disruption also depend on whether alternative suppliers are numerous and diversified or concentrated among a small number of countries. Evaluating strategic import exposure therefore requires information about both dependence on China and the structure of the non-China supplier base.

This paper examines whether changes in this form of strategic import vulnerability are reflected in sectoral equity returns. It applies the monthly Structural Strategic Import Risk Index (SSIRI) developed and validated in the companion study by Oka (2026a) to publicly available U.S. and Chinese equity data from 2004 through 2025. The companion study develops the theoretical foundations, product-level construction, historical behavior, robustness properties, and policy applications of SSIRI. The present paper takes the index as an empirical input and addresses a distinct financial question: whether changes in measured vulnerability are associated with relative equity returns along multinational supply chains.

[★] The author conducted the majority of this research while affiliated with Howard University.

 arsene.oka@bison.howard.edu (A. Oka)

 <https://www.linkedin.com/in/arsene-oka/> (A. Oka)

At the product-family level, SSIRI combines the current Chinese share of U.S. imports with the squared concentration of non-China suppliers measured over a predetermined twelve-month window. The index is constructed for battery and energy-storage materials (BAT) and critical minerals and advanced materials (CMA). This paper examines two sectoral aggregations. Structural SSIRI gives equal weight to each strategic product family, preserving the influence of low-value but potentially critical bottlenecks. Exposure-weighted SSIRI uses fixed reference-period import weights and therefore places greater emphasis on vulnerabilities associated with larger trade flows.¹

Because exposure weights emphasize product families associated with larger import flows, they may align more closely with the aggregate economic exposure represented by broad sectoral portfolios. Structural SSIRI may instead capture low-value bottlenecks whose economic importance is disproportionate to their trade value. Which measure is more closely associated with equity returns is therefore an empirical question; the analysis does not impose that one measure must have a larger coefficient than the other.

The empirical design relates standardized monthly first differences in SSIRI to upstream-minus-downstream equity return spreads. The U.S. portfolios represent industries positioned at different stages of the CMA and BAT value chains. The Chinese analysis uses publicly available sectoral indices to construct comparable relative-return measures. The spread design removes part of the common market and sector-wide return variation and focuses attention on how valuations differ across positions in the production chain.

For the U.S. portfolios, the central directional expectation is a positive upstream-minus-downstream coefficient. Greater strategic import vulnerability may increase the relative scarcity value of upstream capacity while raising expected input and adjustment costs for downstream users. A positive spread does not require downstream returns to decline; it may arise because upstream firms appreciate more strongly or fall less than downstream firms. The prediction is less sharp for China. Greater U.S. dependence may increase the perceived market power and strategic importance of Chinese upstream suppliers, but expectations of export restrictions, foreign diversification, or lower future export volumes may generate offsetting valuation effects.

The empirical design imposes a two-month signal-to-return gap.² For a SSIRI change dated t , the one-month outcome is the return spread in $t+2$; an h -month outcome cumulates spread returns from $t+2$ through $t+h+1$. The return spread in t is included as a control. This timing separates the full-month customs signal from the equity-return window, removes same-month overlap, and leaves an intervening month during which the underlying trade information may be processed. Because the index levels are highly persistent and fail standard stationarity tests, the principal specifications use standardized monthly first differences. The regressions include market, financial, commodity, volatility, policy-uncertainty, and geopolitical-risk controls. Newey–West standard errors (Newey and West, 1987) address serial correlation and overlapping return windows, while Benjamini–Hochberg corrections (Benjamini and Hochberg, 1995) account for multiple testing.

The strongest evidence concerns CMA. A one-standard-deviation increase in the exposure-weighted CMA change is associated with a 1.10-percentage-point increase in the one-month U.S. upstream-minus-downstream return spread observed two months later. The estimate remains significant after a conservative false-discovery-rate correction across the full 32-test baseline family. Portfolio-leg regressions indicate that this short-horizon U.S. result is primarily an upstream-led revaluation: both portfolio legs rise, but the upstream response is substantially larger.

China provides narrower supporting evidence. A one-standard-deviation increase in the exposure-weighted CMA change is associated with an approximately 1.07-percentage-point increase in the Chinese materials-minus-downstream spread over the three-month return window from $t+2$ through $t+4$.³

The relationship is not universal across strategic sectors. U.S. BAT coefficients are generally insignificant after multiple-testing correction. Chinese BAT results are also sensitive to the choice of public equity indices: a thematic new-energy-minus-new-energy-vehicles spread and an alternative battery-industry-minus-automobile spread produce weakly correlated outcomes and coefficients with opposite signs.⁴ Supplementary timing placebos and recursive forecasting exercises support the informational content of the principal CMA relationship.

¹For each BAT and CMA sector, the companion study provides two SSIRI measures: structural SSIRI, which equally weights strategic product families, and exposure-weighted SSIRI, which uses fixed import-exposure weights.

²The gap is specified ex ante to reflect the customs-data release calendar rather than selected according to the estimated coefficients. A reference-month SSIRI observation requires trade flows accumulated through the end of month t , and the corresponding Census trade statistics are generally released during month $t+2$. The return in $t+2$ is therefore the first calendar-month return interval that ordinarily contains the public release of the reference-month trade information. See <https://www.census.gov/foreign-trade/schedule.html>.

³The similarity in coefficient magnitudes is suggestive, but the two estimates refer to different portfolios, samples, return horizons, and market structures. The Chinese result is therefore interpreted as supporting cross-market evidence rather than as a direct replication of the U.S. finding.

⁴The Chinese BAT evidence is therefore treated as a transparent proxy-sensitivity exercise rather than as confirmatory evidence.

This paper relates to research on production networks, network asset pricing, and the financial effects of international trade and political risk. Production-network models show that input-output linkages can propagate and amplify firm- and sector-level shocks (Acemoglu et al., 2012), while empirical studies document real and financial effects of disruptions to specialized suppliers (Barrot and Sauvagnat, 2016; Carvalho et al., 2021). Supplier–customer relationships also transmit information and return variation across connected firms, and network position can affect expected returns (Cohen and Frazzini, 2008; Gofman et al., 2020). Related research shows that trade-policy and political risks affect investment, economic activity, and financial valuations (Caldara et al., 2020; Hassan et al., 2019; Amiti et al., 2021).

The analysis also complements Oka (2026b). Whereas that study examines how economic policy uncertainty originating in the United States, China, and the euro area propagates across broad U.S. equity sectors, the present paper studies how product-level trade dependence and alternative-supplier concentration are associated with relative valuations across sector-specific upstream and downstream portfolios in the United States and China. The two studies therefore examine different sources of international risk, use different empirical designs, and address distinct transmission mechanisms.

The paper makes three contributions. First, it documents the financial-market relevance of a product-level, trade-based measure of strategic import vulnerability.⁵ Second, it identifies where relative repricing occurs along the supply chain. The spread and portfolio-leg regressions distinguish upstream-led appreciation from downstream underperformance and connect strategic import exposure to multinational valuation, sourcing risk, and portfolio allocation. Third, the paper demonstrates that both index aggregation and equity-proxy construction matter. The contrast between structural and exposure-weighted SSIRI, together with the BAT sensitivity analysis, documents which findings are robust and which depend on how trade exposure is mapped into public equity portfolios.

The remainder of the paper proceeds as follows. Section 2 reviews the related literature. Section 3 describes the data and empirical specifications. Section 4 presents the main results. Section 5 discusses the economic interpretation, implications for multinational financial management, and limitations. Section 6 concludes.

2. Related Literature and Positioning

This paper relates to three strands of research: the propagation of shocks through production networks, the pricing of supplier–customer linkages, and the financial effects of trade and political risk. It connects these literatures by examining whether monthly changes in a trade-based measure of strategic import vulnerability are followed by relative equity repricing across upstream and downstream industries. The Structural Strategic Import Risk Index (SSIRI) itself is developed and validated in the companion study by Oka (2026a); the contribution of the present paper is its application to multinational financial markets.

2.1. Production Networks, Shock Propagation, and Asset Prices

Production networks facilitate specialization and comparative advantage, but they can also transmit disruptions across firms and sectors. In network models, the aggregate consequences of a microeconomic shock depend on the position of the affected firm or industry and the structure of its input-output linkages (Acemoglu et al., 2012). Disruptions involving central or difficult-to-replace suppliers can therefore affect firms well beyond the directly exposed industry.

Empirical evidence confirms that supplier characteristics shape shock propagation. Barrot and Sauvagnat (2016) document substantial output and market-value losses among customers following natural disasters affecting specialized suppliers. Using the 2011 Great East Japan Earthquake, Carvalho et al. (2021) show that disruptions propagate both upstream and downstream through direct and indirect production linkages. Firms also adjust their networks in response to perceived exposure. Pankratz and Schiller (2024) find that climate risk at supplier locations affects suppliers and customers and increases the probability of relationship termination, while Ersahin et al. (2024) show that firms respond to supply-chain risk by changing supplier composition and, in some cases, vertically integrating. Together, these studies emphasize that supply-chain exposure depends not only on current input flows but also on the availability and cost of substitution.

Production-network relationships are also reflected in asset prices. Cohen and Frazzini (2008) document return predictability across economically linked firms and show that publicly available customer information is not immediately incorporated into supplier prices. Network position may also affect systematic risk and expected returns. Gofman

⁵This is a financial application of SSIRI; the development and validation of the index remain the contribution of the companion study.

et al. (2020) find that firms located farther upstream have higher average returns and greater exposure to aggregate productivity shocks. Their model emphasizes vertical creative destruction, through which supplier innovations can reduce the value of downstream firms' assets in place.

The present paper connects these literatures by examining whether changes in import-network vulnerability are followed by relative equity returns. Unlike studies organized around realized disasters, discrete disruptions, or directly observed supplier–customer pairs, the analysis uses SSIRI as a monthly sector-level measure of U.S. dependence on China and the concentration of alternative non-China suppliers. It therefore combines two dimensions of production-network risk: an industry's position within a strategic value chain and time variation in the vulnerability of the underlying import network.

2.2. Trade Policy, Political Risk, and Equity Valuation

A third strand of research studies the real and financial effects of trade policy and political uncertainty. Caldara et al. (2020) show that increases in trade policy uncertainty reduce economic activity and investment. At the firm level, Hassan et al. (2019) construct a text-based measure of political risk and show that exposed firms reduce investment and hiring, increase political activity, and experience greater stock-return volatility.

Equity-market reactions have also been used to measure the incidence of trade shocks. Amiti et al. (2021) examine tariff announcements during the U.S.–China trade war and document heterogeneous stock-price responses according to firms' import, export, and revenue exposure to China. These market reactions predict subsequent firm outcomes and provide information about the incidence of trade protection. Using textual measures derived from earnings conference calls, Ersahin et al. (2024) show that firms respond to greater supply-chain risk by changing supplier composition and, when financially able to do so, increasing vertical integration.

Related research examines how broader uncertainty measures affect sector-level financial outcomes. Choi (2024) documents heterogeneous effects of economic policy uncertainty and geopolitical risk on U.S. sectoral equity volatility. Oka (2026b) uses a Bayesian global vector autoregression to distinguish economic policy uncertainty shocks originating in the United States, China, and the euro area and traces their effects across nine U.S. equity sectors. That study shows that the geographic origin of uncertainty matters for both sectoral responses and the channels through which the shocks propagate.

The present paper is distinct from these studies in its measured source of risk, empirical design, and geographic scope. Trade-policy, economic-policy, and geopolitical-risk indices measure uncertain news or political discussion. SSIRI measures an observed structural condition: the extent to which current U.S. dependence on China coincides with concentration among recorded alternative suppliers. It can vary in the absence of a policy announcement and can capture gradual changes in trade and sourcing patterns that event-based or text-based measures may miss.

The empirical design also differs from that of Oka (2026b). Rather than identifying macroeconomic uncertainty shocks and tracing their propagation across broad U.S. sectors, the present paper compares sector-specific upstream and downstream portfolios in both the United States and China. It uses stationary changes in product-based import vulnerability to study relative valuation along strategic supply chains. The economic object, data, methodology, and transmission mechanism are therefore distinct from those of the earlier study, notwithstanding their shared interest in equity markets and international risk.

3. Data and Empirical Methodology

3.1. Strategic Import-Risk Measures

This paper applies the monthly Structural Strategic Import Risk Index (SSIRI) developed and validated in the companion study by Oka (2026a). That study provides the theoretical foundations, product-family classifications, customs-data construction, historical analysis, robustness tests, and policy applications of the index. The present paper uses SSIRI to examine its financial-market relevance.

The analysis covers battery and energy-storage materials (BAT) and critical minerals and advanced materials (CMA). The headline sample extends from January 2004 through December 2025 and contains 264 monthly observations. The underlying trade data begin in January 2003, providing the reference period for the fixed import weights and the observations needed to initialize the lagged supplier-concentration measure.

For product family f in sector S , let $D_{f,t}$ denote China's share of U.S. imports during month t .⁶ Let $H_{f,t-1;12}^{NC}$ denote the Herfindahl–Hirschman concentration index among non-China suppliers, calculated over the 12 months ending in $t - 1$. The product-family risk contribution is

$$\rho_{f,t} = D_{f,t} \left(H_{f,t-1;12}^{NC} \right)^2. \quad (1)$$

This construction assigns greater risk to products for which the United States is highly dependent on China and the remaining supplier base is concentrated. Lagging the supplier window ensures that the concentration term is predetermined relative to month t . Squaring the HHI places greater emphasis on product families with particularly limited diversification among alternative suppliers.

The structural measure gives equal weight to the product families in each sector:

$$\text{SSIRI}_{S,t}^{\text{Str}} = \frac{1}{N_S} \sum_{f \in S} \rho_{f,t}, \quad (2)$$

where N_S is the number of product families assigned to sector S . The exposure-weighted measure instead uses fixed reference-period import weights:

$$\text{SSIRI}_{S,t}^{\text{Exp}} = \sum_{f \in S} \omega_{f,0} \rho_{f,t}, \quad \sum_{f \in S} \omega_{f,0} = 1, \quad (3)$$

where $\omega_{f,0}$ is product family f 's share of sectoral imports during the 2003 reference period.⁷ Structural SSIRI measures the average vulnerability of the products in a sector, whereas exposure-weighted SSIRI emphasizes vulnerabilities associated with larger import flows.

Table 1 reports descriptive statistics for the principal independent and dependent variables. The aggregation schemes produce different sectoral rankings. Average structural SSIRI is higher for CMA than for BAT, whereas average fixed exposure-weighted SSIRI is higher for BAT. This distinction is consistent with the structural measure preserving the influence of narrow product-level vulnerabilities while exposure weighting emphasizes larger import flows. The monthly equity spreads have means close to zero but substantial dispersion, with standard deviations ranging from approximately 3.1 to 6.1 percentage points.

Before estimation, we assess the time-series properties of the SSIRI measures using augmented Dickey–Fuller (ADF) tests (Dickey and Fuller, 1979; Said and Dickey, 1984) and Kwiatkowski–Phillips–Schmidt–Shin (KPSS) tests (Kwiatkowski et al., 1992). As reported in Appendix Table 6, the results support using monthly first differences rather than index levels in the main specifications.

The principal specifications therefore use standardized monthly first differences:

$$x_{S,t}^m = \frac{\Delta \text{SSIRI}_{S,t}^m - \overline{\Delta \text{SSIRI}_S^m}}{s(\Delta \text{SSIRI}_S^m)}, \quad m \in \{\text{Str}, \text{Exp}\}. \quad (4)$$

Standardization permits comparison across sectors and aggregation methods: the reported coefficient measures the return associated with a one-standard-deviation monthly increase in strategic import vulnerability. In the recursive forecasting exercise, all transformations are estimated using only the corresponding training sample. Figure 1 plots the standardized monthly first differences used in the empirical analysis.

3.2. Sectoral Equity Returns

All equity data are obtained from publicly available sources. The U.S. analysis uses monthly value-weighted industry portfolio returns. For the Chinese indices, daily closing prices are sampled on the final available trading

⁶Mainland China, Hong Kong, and Macao are included in the China total.

⁷Fixed weights prevent subsequent changes in import composition from mechanically altering the economic meaning of the aggregation.

Table 1
Descriptive Statistics for SSIRI and Sectoral Equity Return Spreads

Variable	N	Mean	SD	Minimum	P25	Median	P75	Maximum
<i>Panel A: SSIRI levels</i>								
CMA structural SSIRI	264	0.0608	0.0344	0.0143	0.0335	0.0545	0.0773	0.1905
CMA fixed exposure-weighted SSIRI	264	0.0205	0.0108	0.0071	0.0119	0.0184	0.0262	0.0570
BAT structural SSIRI	264	0.0349	0.0129	0.0109	0.0259	0.0324	0.0407	0.0774
BAT fixed exposure-weighted SSIRI	264	0.0426	0.0245	0.0088	0.0241	0.0367	0.0525	0.1266
<i>Panel B: Monthly SSIRI changes</i>								
CMA structural change	263	0.0004	0.0255	-0.1007	-0.0099	-0.0006	0.0115	0.0960
CMA fixed exposure-weighted change	263	-0.0001	0.0071	-0.0254	-0.0032	-0.0001	0.0030	0.0243
BAT structural change	263	0.0000	0.0126	-0.0367	-0.0066	0.0001	0.0068	0.0455
BAT fixed exposure-weighted change	263	-0.0001	0.0060	-0.0290	-0.0032	-0.0003	0.0026	0.0219
<i>Panel C: Monthly upstream-minus-downstream return spreads</i>								
U.S. CMA spread	264	-0.127	4.267	-13.389	-2.637	-0.107	2.738	10.725
China CMA spread	251	0.186	5.114	-15.710	-2.927	0.083	3.051	16.711
U.S. BAT spread	264	-0.069	6.052	-22.650	-3.236	0.247	3.326	17.235
China BAT spread	124	-0.312	3.142	-8.994	-2.329	-0.161	1.564	9.816

Notes: Panel A reports the levels of structural and fixed exposure-weighted SSIRI. Panel B reports unstandardized monthly first differences in those indices; the empirical regressions use standardized versions of these changes over the corresponding estimation samples. Panel C reports monthly upstream-minus-downstream equity return spreads in percentage points. U.S. portfolio-leg returns are measured relative to the U.S. market, and Chinese portfolio-leg returns are measured relative to the CSI 300. The SSIRI and U.S. samples run from January 2004 through December 2025; the China CMA spread begins in February 2005, and the China BAT spread begins in September 2015. Multi-month regression outcomes are overlapping cumulative returns constructed from the monthly spreads in Panel C.

day of each calendar month, and monthly simple returns are calculated from consecutive month-end prices. All return series are aligned with SSIRI using calendar month-end dates.

For the United States, predetermined upstream and downstream portfolios are constructed from publicly available value-weighted industry returns. The upstream portfolios contain industries positioned closer to the extraction, processing, or supply of the relevant strategic inputs. The downstream portfolios contain industries that use those inputs in manufacturing or final production. The classifications and weights are held fixed throughout the sample. The exact constituent portfolios and weights are reported in Appendix A.

For China, the CMA upstream portfolio is represented by the CSI 300 Materials Index. The downstream leg is a predetermined composite of the CSI 300 Industrials and CSI 300 Information Technology indices. The exact weights used to construct the composite are reported in Appendix A.

The initial China BAT measure compares the CSI New Energy Index with the CSI New Energy Vehicles Index. These thematic indices do not provide as clean an upstream–downstream classification as the CMA indices. We therefore conduct an additional proxy-validation exercise using the CNI Battery Industry Index relative to the Shenwan Automobile Industry Index.⁸

For country c , sector S , and month t , the principal outcome is the upstream-minus-downstream return spread:

$$r_{c,S,t}^{Spread} = r_{c,S,t}^{Up} - r_{c,S,t}^{Down}. \quad (5)$$

Figures 2 and 3 plot the monthly portfolio-legs returns and the resulting upstream-minus-downstream spreads.

The spread removes part of the market-wide and sector-wide return variation shared by the two portfolio legs and focuses on upstream–downstream return differentials across positions in the production chain.

We examine cumulative arithmetic spread returns over horizons $h \in \{1, 3, 6, 12\}$, beginning two months after the signal month:

⁸Because the two China BAT constructions produce weakly correlated return spreads and materially different coefficient estimates, the China BAT evidence is reported as exploratory.

$$R_{c,S,t+2:t+h+1}^{Spread,(h)} = \sum_{j=0}^{h-1} r_{c,S,t+2+j}^{Spread} \quad (6)$$

Thus, for a SSIRI change dated t , the one-month outcome is $r_{c,S,t+2}^{Spread}$. An h -month outcome runs from $t + 2$ through $t + h + 1$; month $t + 1$ is not part of the return window. This two-month signal-to-return gap creates a clean temporal separation between the full-month customs observation and the subsequent equity outcome. It also reduces the possibility that the estimates reflect mechanical same-month comovement.⁹

3.3. Control Variables

The control set is designed to separate relative sectoral returns from broad financial, commodity, and geopolitical conditions. U.S. specifications include the market excess return and the size, value, profitability, investment, and momentum factors. China specifications include the CSI 300 market return. Expanded specifications additionally control for relevant commodity-market conditions, equity-market volatility, economic policy uncertainty, and geopolitical risk. Persistent control variables are expressed in changes or other stationary transformations where appropriate.

Each specification includes the current return spread $r_{c,S,t}^{Spread}$ to absorb short-run persistence. The remaining financial, market, commodity, and uncertainty controls are also dated t . No control is drawn from the intervening month $t + 1$ or from the subsequent return window. Because the dependent variable is an upstream-minus-downstream spread, common market shocks are already partially differenced out. The additional controls address residual exposure to aggregate financial and macroeconomic conditions.

Table 2: Data Sources and Variable Construction

Variable group	Public source	Series and construction	Coverage used
SSIRI inputs	U.S. Census Bureau international trade data	Monthly U.S. import values by country of origin and HS6 product. Mainland China, Hong Kong, and Macao are classified as China. The 2003 observations determine the fixed exposure weights and initialize the lagged supplier-concentration window.	Jan. 2003–Dec. 2025
Structural and exposure-weighted SSIRI	Oka (2026a); calculations from Census trade data	The current China import share is multiplied by the squared HHI of non-China suppliers over the preceding 12 months. Product-family contributions are aggregated using equal weights or fixed 2003 import-exposure weights.	Jan. 2004–Dec. 2025
U.S. industry returns	Kenneth R. French Data Library	Monthly value-weighted industry portfolio returns used to construct the predetermined U.S. CMA and BAT upstream and downstream portfolios.	Jan. 2004–Dec. 2025
U.S. financial factors	Kenneth R. French Data Library	Market excess return, size, value, profitability, investment, and momentum factors. Percentage returns are converted to decimal units before estimation.	Jan. 2004–Dec. 2025
China broad market	China Securities Index Company; public Tencent data endpoint	CSI 300 Index (000300). Daily closing prices are converted to month-end simple returns.	Feb. 2005–Dec. 2025

Continued on next page

⁹The gap improves temporal ordering but does not by itself establish causality or real-time availability of the finalized trade data.

Table 3
Construction of Sectoral Equity Outcomes

Country	Sector	Upstream leg	Downstream leg	Status
United States	CMA	Fixed portfolio of value-weighted upstream critical-material industries	Fixed portfolio of value-weighted downstream industries using critical-material inputs	Main
China	CMA	CSI 300 Materials Index (000909)	Fixed composite of CSI 300 Industrials (000910) and CSI 300 Information Technology (000915)	Main
United States	BAT	Fixed portfolio of value-weighted upstream battery-material industries	Fixed portfolio of value-weighted downstream battery-using industries	Main
China	BAT	CSI New Energy Index (399808)	CSI New Energy Vehicles Index (399976)	Exploratory
China	BAT	CNI Battery Industry Index (980032)	Shenwan Automobile Industry Index (801880)	Alternative

Notes: The principal dependent variable is the upstream return minus the downstream return. A SSIRI change dated t is related to an h -month arithmetic sum of spread returns from $t + 2$ through $t + h + 1$. Portfolio classifications are fixed over time and are not selected according to the regression results. Exact U.S. portfolio constituents and weights, together with the weights of the China CMA downstream composite, are reported in Appendix A. Both China BAT proxy constructions are disclosed to avoid selective reporting.

Table 2 — *Continued from previous page*

Variable group	Public source	Series and construction	Coverage used
China CMA indices	China Securities Index Company; public Tencent data endpoint	CSI 300 Materials (000909), CSI 300 Industrials (000910), and CSI 300 Information Technology (000915). Daily closing prices are converted to month-end returns.	Feb. 2005–Dec. 2025
China BAT thematic indices	China Securities Index Company; public Tencent data endpoint	CSI New Energy (399808) and CSI New Energy Vehicles (399976). Their return spread is treated as an exploratory China BAT proxy.	Sep. 2015–Dec. 2025
Alternative China battery index	CNI Index; public Sina data through AKShare	CNI Battery Industry Index (980032). Daily closing prices are converted to month-end returns.	Feb. 2013–Dec. 2025
Alternative China automobile index	Shenwan Research; public data through AKShare	Shenwan Automobile Industry Index (801880). Daily closing prices are converted to month-end returns.	Mar. 2014–Dec. 2025
Financial and geopolitical controls	Public financial and policy-risk data providers	Market volatility, commodity-market conditions, economic policy uncertainty, and geopolitical risk. Exact series identifiers, transformations, and retrieval sources are reported in Appendix A and the replication package.	As available through Dec. 2025

Notes: SSIRI denotes the Structural Strategic Import Risk Index; BAT denotes Battery and Energy Storage Materials; and CMA denotes Critical Minerals and Advanced Materials. Public APIs and market-data endpoints are retrieval channels rather than the original producers of the corresponding indices. All daily equity series are sampled on the final available trading day of each calendar month. The estimation sample for each specification is determined by the common availability of the dependent variable, SSIRI measure, and controls.

3.4. Baseline Regression

For measure $m \in \{Str, Exp\}$, the baseline two-month-ahead relationship is

$$R_{c,S,t+2:t+h+1}^{Spread,(h)} = \alpha_{c,S,h} + \beta_{c,S,h}^m x_{S,t}^m + \phi_{c,S,h} r_{c,S,t}^{Spread} + \gamma'_{c,S,h} \mathbf{Z}_{c,S,t} + \varepsilon_{c,S,t}^{(h)}, \quad (7)$$

where $x_{S,t}^m$ is the standardized structural or exposure-weighted SSIRI change, $r_{c,S,t}^{Spread}$ is the current spread return, and $\mathbf{Z}_{c,S,t}$ contains the remaining controls. Coefficients are estimated using returns in decimal units and reported in percentage points. Thus, $\beta = 0.010$ corresponds to a one-percentage-point change in the subsequent upstream-minus-downstream return spread associated with a one-standard-deviation increase in SSIRI.

Inference uses heteroskedasticity- and autocorrelation-consistent standard errors (Newey and West, 1987). Serial-correlation corrections are especially important for the overlapping multi-month outcomes in Equation (6). The Newey–West bandwidth is $6 + h$ monthly lags, where h is the return horizon.

Because the analysis considers multiple markets, sectors, horizons, and SSIRI measures, conventional p -values may overstate the strength of the evidence. We therefore report false-discovery-rate-adjusted q -values using the Benjamini–Hochberg procedure (Benjamini and Hochberg, 1995). Unless otherwise stated, the reported global q -values are calculated across the full family of tests reported in the corresponding analysis.

Equation (7) estimates conditional associations rather than causal effects. SSIRI changes and equity returns may respond jointly to trade-policy announcements, geopolitical developments, commodity-market news, or other information. The estimates are therefore interpreted as two-month-ahead relative equity returns associated with changes in measured strategic import vulnerability.

3.5. Comparing the Two SSIRI Measures

To determine whether structural and exposure-weighted changes contain incremental information, we estimate

$$R_{c,S,t+2:t+h+1}^{Spread,(h)} = \alpha_{c,S,h} + \beta_{c,S,h}^{Str} x_{S,t}^{Str} + \beta_{c,S,h}^{Exp} x_{S,t}^{Exp} + \phi_{c,S,h} r_{c,S,t}^{Spread} + \gamma'_{c,S,h} \mathbf{Z}_{c,S,t} + \varepsilon_{c,S,t}^{(h)}. \quad (8)$$

We test the individual coefficients, their joint significance, and the restriction

$$H_0 : \beta_{c,S,h}^{Str} = \beta_{c,S,h}^{Exp}. \quad (9)$$

This specification evaluates whether the structural measure contains information beyond economically weighted exposure and whether the exposure-weighted measure remains informative after conditioning on the breadth of product-level vulnerability. Failure to reject Equation (9) does not establish that the two measures are equivalent; it indicates that their coefficient difference is estimated imprecisely.

3.6. Portfolio-Leg Decomposition

For the principal systems, Equation (7) is re-estimated separately for the upstream return, the downstream return, and their spread. All three regressions use the same observations, controls, and variable transformations. Consequently,

$$\hat{\beta}_{c,S,h}^{Spread} = \hat{\beta}_{c,S,h}^{Up} - \hat{\beta}_{c,S,h}^{Down}. \quad (10)$$

The decomposition identifies whether a positive estimated return differential reflects upstream appreciation, downstream underperformance, or both. This exercise is descriptive of the portfolio components and is not interpreted as separately identifying scarcity rents, input-cost effects, or policy expectations.

3.7. Multiple Testing, Timing Placebos, and Proxy Sensitivity

The global baseline family contains four country-sector outcomes, two SSIRI measures, and four return horizons, for a total of 32 tests. The outcomes are U.S. CMA, China CMA, U.S. BAT, and the original thematic China BAT spread. We report conventional p -values and Benjamini–Hochberg false-discovery-rate-adjusted q -values (Benjamini and Hochberg, 1995). The correction across all 32 tests provides the principal and most conservative multiplicity

adjustment. China BAT is interpreted as exploratory because of its proxy instability, but it remains in the global correction so that the principal results do not benefit from its exclusion. The alternative China BAT construction is evaluated in a separate proxy-sensitivity family.

Timing placebos replace the current SSIRI change with changes observed six or twelve months in the future while retaining the current signal as a control. A future realization should not explain the return window from $t + 2$ onward after conditioning on the correctly dated measure.

China BAT is evaluated separately because its public equity proxies do not provide a stable upstream–downstream mapping. We compare the new-energy thematic spread with the battery-industry-minus-automobile spread over their common sample and report their correlation, coefficient signs, and formal coefficient-difference tests. China BAT is classified as exploratory on measurement grounds, and the complete results from both proxies are disclosed.

3.8. Supplementary Out-of-Sample Exercise

The supplementary forecasting analysis focuses on the U.S. CMA upstream-minus-downstream spread. The baseline exercise uses an initial training window of 120 months and generates one-month-ahead forecasts from an expanding sample. Historical-mean and autoregressive return models serve as benchmarks. Competing models add the structural change, the exposure-weighted change, or both changes. For a target return r_t , the autoregressive benchmark uses r_{t-1} , while the SSIRI models add x_{t-2}^{Str} , x_{t-2}^{Exp} , or both. This preserves the two-month signal-to-return separation used in the in-sample regressions.

Forecast accuracy is evaluated using the root-mean-square error, mean absolute error, and out-of-sample R^2 :

$$R_{OOS}^2 = 1 - \frac{\sum_t \left(r_t - \hat{r}_{t|t-1}^{Model} \right)^2}{\sum_t \left(r_t - \hat{r}_{t|t-1}^{Benchmark} \right)^2}. \quad (11)$$

Nested forecast comparisons use the one-sided Clark–West adjustment (Clark and West, 2007), with false-discovery-rate correction across the forecast comparisons. Robustness exercises vary the initial training window and use moving-block bootstrap confidence intervals to preserve serial dependence in forecast losses.

This exercise evaluates whether SSIRI contains incremental statistical information beyond simple return benchmarks.

4. Empirical Results

This section reports the two-month-ahead relationships defined in Section 3. Coefficients are expressed in percentage points for a one-standard-deviation monthly change in SSIRI. Unless otherwise stated, the reported q -values use the 32-test baseline family defined in Section 3.7. Detailed stationarity diagnostics are reported in Appendix B.

4.1. CMA relative equity returns

Table 4 presents the principal results. The strongest relationship is between exposure-weighted CMA SSIRI and the U.S. upstream-minus-downstream return spread. For the one-month return observed two months after the signal, a one-standard-deviation increase in the exposure-weighted CMA SSIRI change is associated with a 1.100-percentage-point increase in the spread ($n = 261$). The HAC test statistic is 4.40, the unadjusted p -value is below 0.001, and the relationship remains significant after correction across all 32 baseline tests ($q_{all} = 0.0003$).

The structural¹⁰ CMA SSIRI change is also positively associated with the two-month-ahead one-month U.S. spread. A one-standard-deviation structural SSIRI change is associated with a 0.579-percentage-point increase in the spread ($p = 0.0080$; $q_{all} = 0.0490$). The larger exposure-weighted estimate is consistent with import-exposure weights capturing economically important variation that is attenuated under equal product-family weights. The separate regressions do not, however, establish that the two coefficients differ statistically.

The U.S. relationship remains evident over the three-month cumulative window from $t + 2$ through $t + 4$. The three-month coefficients are 0.995 percentage points for the exposure-weighted measure ($p = 0.0040$; $q_{all} = 0.0358$) and 0.785 percentage points for the structural measure ($p = 0.0092$; $q_{all} = 0.0490$). The exposure-weighted coefficient

¹⁰Structural SSIRI is the equal-weight SSIRI.

Table 4
SSIRI Changes and Upstream-Minus-Downstream Equity Returns

Market and horizon	SSIRI measure	Specification	Coefficient (pp)	HAC SE (pp)	Test statistic	p-value	FDR q-value	N	Adjusted R ²
U.S. CMA: 1 month	Structural first difference	Baseline	0.6074	0.2014	3.0166	0.0026	0.0205	261	0.0000
U.S. CMA: 1 month	Exposure-weighted first difference	Baseline	1.0087	0.2398	4.2067	0.0000	0.0008	261	0.0362
U.S. CMA: 1 month	Structural first difference	Commodity controls	0.6120	0.2042	2.9974	0.0027	0.0218	261	-0.0033
U.S. CMA: 1 month	Exposure-weighted first difference	Commodity controls	1.0071	0.2394	4.2065	0.0000	0.0008	261	0.0324
U.S. CMA: 1 month	Structural first difference	Full controls	0.5789	0.2184	2.6506	0.0080	0.0490	261	0.0021
U.S. CMA: 1 month	Exposure-weighted first difference	Full controls	1.0993	0.2497	4.4030	0.0000	0.0003	261	0.0498
U.S. CMA: 3 months	Structural first difference	Baseline	0.7913	0.2866	2.7612	0.0058	0.0307	259	-0.0131
U.S. CMA: 3 months	Exposure-weighted first difference	Baseline	0.9376	0.3370	2.7818	0.0054	0.0307	259	-0.0085
U.S. CMA: 3 months	Structural first difference	Commodity controls	0.7935	0.2881	2.7539	0.0059	0.0314	259	-0.0171
U.S. CMA: 3 months	Exposure-weighted first difference	Commodity controls	0.9383	0.3347	2.8033	0.0051	0.0314	259	-0.0125
U.S. CMA: 3 months	Structural first difference	Full controls	0.7855	0.3016	2.6048	0.0092	0.0490	259	-0.0292
U.S. CMA: 3 months	Exposure-weighted first difference	Full controls	0.9947	0.3454	2.8794	0.0040	0.0358	259	-0.0227
China CMA: 3 months	Structural first difference	Baseline	0.7529	0.4057	1.8559	0.0635	0.1354	247	-0.0014
China CMA: 3 months	Exposure-weighted first difference	Baseline	1.0846	0.3222	3.3657	0.0008	0.0081	247	0.0055
China CMA: 3 months	Structural first difference	Commodity controls	0.7483	0.4109	1.8214	0.0686	0.1716	247	-0.0049
China CMA: 3 months	Exposure-weighted first difference	Commodity controls	1.0975	0.3211	3.4174	0.0006	0.0067	247	0.0024
China CMA: 3 months	Structural first difference	Full controls	0.8341	0.4083	2.0429	0.0411	0.1314	247	-0.0010
China CMA: 3 months	Exposure-weighted first difference	Full controls	1.0713	0.3769	2.8425	0.0045	0.0358	247	0.0038
China BAT: 3 months	Structural first difference	Baseline	0.0981	0.3458	0.2837	0.7767	0.8876	120	-0.0131
China BAT: 3 months	Exposure-weighted first difference	Baseline	3.4277	0.9810	3.4942	0.0005	0.0076	120	0.0719
China BAT: 3 months	Structural first difference	Commodity controls	0.1065	0.3485	0.3056	0.7599	0.8685	120	-0.0295
China BAT: 3 months	Exposure-weighted first difference	Commodity controls	3.4352	0.9996	3.4366	0.0006	0.0067	120	0.0570
China BAT: 3 months	Structural first difference	Full controls	0.0912	0.3522	0.2590	0.7957	0.8804	120	-0.0461
China BAT: 3 months	Exposure-weighted first difference	Full controls	3.5521	1.0198	3.4832	0.0005	0.0079	120	0.0479

Notes: The dependent variable is the upstream-minus-downstream equity return spread. A standardized monthly SSIRI first difference dated t is related to an h -month cumulative arithmetic return beginning in $t + 2$; thus, the return window extends from $t + 2$ through $t + h + 1$. All regressors are dated t , and every specification includes the current return spread.

For the United States, the baseline specification includes the market, size, value, profitability, investment, and momentum factors. For China, the baseline includes the CSI 300 market return. Commodity specifications add global metals returns and, for BAT, global energy returns. Full specifications additionally include changes in the VIX, global economic policy uncertainty, and geopolitical risk. Newey–West HAC standard errors use 6 + h monthly lags.

Coefficients and standard errors are reported in percentage points for a one-standard-deviation increase in the corresponding monthly SSIRI change. Structural SSIRI equally weights strategic product families, whereas exposure-weighted SSIRI uses fixed 2003 import-exposure weights. Reported p -values are two-sided. For each control specification, FDR q -values apply the Benjamini–Hochberg correction across the complete 32-test family: four country-sector outcomes, two SSIRI measures, and four return horizons. The table displays the principal short-horizon estimates; the complete family is reported in the replication files. Values reported as 0.0000 are positive values rounded to four decimal places. Negative adjusted R^2 values may arise from the degrees-of-freedom adjustment and do not indicate negative unadjusted explanatory power. China BAT is exploratory because its estimates are sensitive to the public equity proxy. CMA denotes Critical Minerals and Advanced Materials; BAT denotes Battery and Energy Storage Materials; pp denotes percentage points; and N is the regression sample size.

remains positive at six months, with a global q -value of 0.0653, but its twelve-month q -value rises to 0.1124. The evidence is therefore strongest over short cumulative windows and does not support an indefinitely persistent return differential.

The adjusted R^2 of the one-month U.S. exposure-weighted specification is approximately 0.050. Thus, the SSIRI change explains an economically relevant but still limited portion of monthly spread-return variation. This magnitude is unsurprising given the volatility of monthly equity returns and the large number of other factors affecting industry valuations.

China provides supporting evidence for the CMA relationship. A one-standard-deviation increase in the exposure-weighted CMA change is associated with a 1.071-percentage-point increase in the three-month materials-minus-downstream spread from $t + 2$ through $t + 4$ ($n = 247$, $p = 0.0045$; $q_{\text{all}} = 0.0358$). The corresponding structural coefficient is 0.834 percentage points but does not remain significant after the global correction ($p = 0.0411$; $q_{\text{all}} = 0.1314$).

The adjusted R^2 of the Chinese exposure-weighted specification is approximately 0.004. The Chinese result therefore represents a precisely estimated incremental association rather than a model that explains a large fraction of overall return variation. Its sign is nevertheless consistent with upstream–downstream return differentials toward firms positioned closer to the supply of critical materials.

These coefficients describe relative equity-market responses rather than aggregate market returns. A positive spread coefficient may arise because upstream firms appreciate, downstream firms underperform, or both occur simultaneously. The portfolio-leg decomposition clarifies these components.

4.2. Upstream and Downstream Decomposition

Table 5 estimates the upstream and downstream legs separately using common samples and identical regressors. Under this common specification, the coefficient identity

$$\hat{\beta}^{Spread} = \hat{\beta}^{Upstream} - \hat{\beta}^{Downstream} \quad (12)$$

holds up to numerical precision.

For U.S. CMA, the two-month-ahead one-month response reflects positive returns in both portfolio legs, with a substantially larger upstream response. The exposure-weighted coefficient is 1.898 percentage points for upstream firms and 0.798 percentage points for downstream firms. Both coefficients are statistically significant after correction within the component family, and their difference produces the 1.100-percentage-point spread response. The result is therefore better characterized as an upstream-led rally than as a decline in downstream equities during the subsequent return month.

Over the three-month window from $t + 2$ through $t + 4$, the U.S. upstream coefficient is 0.833 percentage points, whereas the downstream coefficient is -0.162 percentage points. Neither portfolio-leg coefficient is individually precise, but their difference is 0.995 percentage points and remains significant after the global correction applied to the baseline spread regressions. This pattern indicates that the relative performance of the two supply-chain positions is estimated more precisely than either leg in isolation.

The Chinese CMA decomposition displays a clearer directional divergence at three months. The upstream materials coefficient is 0.647 percentage points, whereas the downstream coefficient is -0.424 percentage points. Their difference produces the 1.071-percentage-point spread coefficient. The upstream estimate is marginally significant at the 10% level after correction within the component family, while the downstream estimate is imprecise. Taken together, the U.S. and Chinese decompositions indicate that increases in measured CMA vulnerability are associated with a relative revaluation toward upstream materials exposure.

The decomposition does not separately identify scarcity rents, changes in expected input costs, anticipated policy support, or changes in risk premia. It identifies which portfolio leg contributes to the estimated relative-return relationship.

4.3. Structural versus Exposure-Weighted SSIRI

The horse-race specifications in Appendix Table 7 include structural and exposure-weighted changes jointly. At the two-month-ahead one-month horizon, the exposure-weighted CMA coefficient remains positive and statistically significant at 0.908 percentage points ($p = 0.0005$; $q = 0.0054$ within the corresponding horse-race family). The structural coefficient falls to 0.291 percentage points and is not statistically distinguishable from zero ($p = 0.1584$; $q = 0.3802$).

The difference between the two coefficients is not statistically significant after multiple-testing correction ($p = 0.1111$; $q = 0.3333$). The evidence therefore does not establish that exposure-weighted and structural SSIRI have different population coefficients. It instead shows that, conditional on both measures, the two-month-ahead one-month exposure-weighted coefficient retains statistically significant incremental information, whereas the structural coefficient does not.

This result is consistent with the different purposes of the two measures. Equal product-family weights preserve the influence of narrow structural bottlenecks irrespective of trade value. Fixed import-exposure weights place greater emphasis on vulnerabilities associated with larger trade flows and may therefore align more closely with the aggregate economic exposure represented by broad equity portfolios. Structural SSIRI remains useful for identifying low-value strategic bottlenecks that may not map closely into publicly traded portfolio weights.

4.4. Sectoral Heterogeneity and the BAT Results

The U.S. BAT results provide an informative null. Neither structural nor exposure-weighted BAT changes generate a stable upstream-minus-downstream response after the global FDR correction. The smallest unadjusted p -values are approximately 0.10, and the corresponding q_{all} -values exceed 0.18. The U.S. CMA relationship is therefore not a generic consequence of an increase in any measure of strategic import vulnerability.

Chinese BAT estimates depend materially on the public equity proxy. The original new-energy thematic spread yields positive coefficients, whereas a battery-industry-minus-automobile spread yields coefficients of the opposite

Table 5

Exposure-Weighted SSIRI: Upstream, Downstream, and Spread Decomposition

System	Leg	Horizon	N	Coefficient (pp)	HAC SE (pp)	Test statistic	p-value	Within-system q-value	Global FDR q-value	Adjusted R2
China BAT	Downstream	1	122	-2.0768	1.5638	-1.3280	0.1842	0.2210	0.2210	0.0065
China BAT	Spread	1	122	1.4107	0.7850	1.7971	0.0723	0.1085	0.1183	0.0140
China BAT	Upstream	1	122	-0.6661	0.9812	-0.6789	0.4972	0.4972	0.5593	0.0014
China BAT	Downstream	3	120	-6.9383	2.4037	-2.8865	0.0039	0.0117	0.0134	0.0315
China BAT	Spread	3	120	3.5521	1.0198	3.4832	0.0005	0.0030	0.0030	0.0479
China BAT	Upstream	3	120	-3.3862	1.7660	-1.9174	0.0552	0.1085	0.1104	-0.0118
China CMA	Downstream	1	249	-0.0799	0.2513	-0.3180	0.7505	0.7505	0.7505	-0.0193
China CMA	Spread	1	249	0.5490	0.3231	1.6994	0.0893	0.1339	0.1339	-0.0081
China CMA	Upstream	1	249	0.4691	0.2520	1.8618	0.0626	0.1253	0.1127	0.0005
China CMA	Downstream	3	247	-0.4241	0.2689	-1.5772	0.1147	0.1377	0.1589	-0.0147
China CMA	Spread	3	247	1.0713	0.3769	2.8425	0.0045	0.0269	0.0134	0.0038
China CMA	Upstream	3	247	0.6472	0.2939	2.2022	0.0277	0.0830	0.0622	0.0118
U.S. CMA	Downstream	1	261	0.7983	0.3044	2.6223	0.0087	0.0131	0.0225	0.0675
U.S. CMA	Spread	1	261	1.0993	0.2497	4.4030	0.0000	0.0000	0.0001	0.0498
U.S. CMA	Upstream	1	261	1.8977	0.4370	4.3429	0.0000	0.0000	0.0001	0.0857
U.S. CMA	Downstream	3	259	-0.1618	0.3932	-0.4116	0.6806	0.6806	0.7207	0.0504
U.S. CMA	Spread	3	259	0.9947	0.3454	2.8794	0.0040	0.0080	0.0134	-0.0227
U.S. CMA	Upstream	3	259	0.8328	0.5504	1.5132	0.1302	0.1563	0.1674	0.0128

Notes: The table decomposes the exposure-weighted upstream-minus-downstream return response into its portfolio components. A standardized exposure-weighted SSIRI first difference dated t is related to one- or three-month cumulative arithmetic returns beginning in $t+2$. Thus, an h -month outcome covers $t+2$ through $t+h+1$. The upstream, downstream, and spread regressions within each system use a common sample and identical regressors; consequently, the spread coefficient equals the upstream coefficient minus the downstream coefficient up to numerical precision. All specifications include the current return spread and the full set of financial, commodity, volatility, economic-policy-uncertainty, and geopolitical-risk controls described in the text. Newey–West HAC standard errors use 6 + h monthly lags. Coefficients and standard errors are reported in percentage points for a one-standard-deviation increase in the monthly SSIRI change. Reported p -values are two-sided. Within-system q -values apply the Benjamini–Hochberg correction across the six estimates for each system (three portfolio legs and two horizons); global q -values apply the correction across all 18 estimates in the table. Values reported as 0.0000 are positive values rounded to four decimal places. The China BAT results are exploratory because they are sensitive to the choice of public equity proxy. CMA denotes Critical Minerals and Advanced Materials; BAT denotes Battery and Energy Storage Materials; and pp denotes percentage points.

sign over the common sample. The two spreads are negatively correlated, and their long-horizon coefficient difference is statistically significant. China BAT is therefore classified as exploratory, and the paper's principal conclusions do not depend on either proxy. Appendix D reports the complete proxy comparison and discusses the measurement implications.

4.5. Robustness and Timing Placebos

The principal CMA findings remain under the expanded commodity, volatility, policy-uncertainty, and geopolitical-risk controls. Moreover, no future-signal coefficient survives the timing-placebo correction. Appendix E reports the full-horizon joint regressions, timing placebos, and additional specification checks. These exercises strengthen the timing and specification audit but do not establish causality.

4.6. Supplementary Pseudo-Out-of-Sample Evidence

With a 120-month initial training sample, the exposure-change model achieves an out-of-sample R^2 of 5.27% relative to the AR(1) benchmark, and the Clark–West comparison remains significant after the forecast-family correction ($q = 0.0337$). Appendix F reports the nested-model comparisons, training-window sensitivity, bootstrap uncertainty, and cumulative forecast gains. Because the exercise uses finalized trade data, it measures statistical forecast content rather than implementable trading performance.

5. Discussion

The empirical results document a measurable association between changes in strategic import vulnerability and the relative performance of upstream and downstream equities. The relationship is strongest for critical minerals and advanced materials. In the United States, increases in exposure-weighted CMA vulnerability are followed two months later by higher upstream returns relative to downstream returns. A narrower three-month relationship also appears in China over the window beginning two months after the signal. These findings provide evidence that changes in SSIRI have observable financial-market correlates beyond mechanical variation in bilateral import shares.

5.1. Economic Interpretation

A rise in SSIRI indicates that U.S. dependence on China has increased, the non-China supplier base has become more concentrated, or both have occurred. Such a change may alter investors' assessments of how risks and potential rents are distributed across a supply chain. Firms positioned upstream may benefit from expectations of greater scarcity rents, stronger pricing power, investment in alternative capacity, or policy support. Downstream firms may instead face higher expected input costs, narrower margins, additional inventory requirements, or greater uncertainty about access to specialized materials.

The portfolio-leg decomposition is consistent with these potential mechanisms but does not identify them separately. In the United States, the two-month-ahead one-month CMA relationship is primarily an upstream-led revaluation rather than an outright decline in downstream equity values. Both portfolio legs have positive coefficients, but the upstream response is substantially larger. Over the three-month window beginning at $t + 2$, neither leg is individually precise, while their difference remains statistically significant. The relevant financial object is therefore the relative performance of the two supply-chain positions rather than the return of either portfolio in isolation.

In China, the three-month point estimates combine stronger upstream performance with weaker downstream performance, although the individual leg estimates are less precise than the spread. This pattern is consistent with a reallocation of relative value toward Chinese materials producers when U.S. dependence on China becomes more difficult to replace. Other interpretations remain possible. Greater dependence may increase the perceived market power of Chinese suppliers, but expectations of export controls, foreign diversification, or lower future export volumes could produce offsetting valuation effects.

These patterns should not be interpreted as evidence that greater vulnerability benefits the economy or improves aggregate welfare. Higher upstream equity values may reflect anticipated scarcity rents or policy transfers rather than an expansion of productive capacity. Likewise, an upstream-minus-downstream return spread does not measure the economy-wide costs of supply disruption, substitution, stockpiling, or investment in resilience. The results describe the distribution of relative market valuations across supply-chain positions, not the net social consequences of strategic dependence.

5.2. Structural versus Exposure-Weighted Vulnerability

The more precisely estimated short-run relationship for exposure-weighted SSIRI helps clarify the different roles of the two index constructions. Structural SSIRI assigns equal importance to each strategic product family and is therefore suited to identifying narrow bottlenecks that might otherwise receive little attention. Exposure-weighted SSIRI gives greater weight to product families representing larger import flows and is more closely aligned with the aggregate economic scale represented in broad equity portfolios.

The joint regressions do not establish that the structural measure is irrelevant. In the two-month-ahead one-month U.S. CMA specification, the exposure-weighted coefficient remains statistically significant after both measures are included, whereas the structural coefficient does not. The formal coefficient-difference test is insignificant, however, and therefore does not establish that the population effects of the two measures differ. The evidence shows that the exposure-weighted measure contains more precisely estimated incremental information in the short-run specification, not that structural vulnerability has no financial relevance.

The distinction reflects the different purposes of the measures. Structural SSIRI is designed to identify product-level vulnerabilities regardless of their current import value. A low-value input may remain strategically important if it has few substitutes or is essential to a larger production process. Exposure-weighted SSIRI, by contrast, emphasizes vulnerabilities associated with larger observed trade flows and may therefore map more closely into sector-level revenues, costs, and valuations. The appropriate measure depends on the application rather than on statistical significance alone.

5.3. Sectoral and Proxy Heterogeneity

The absence of robust U.S. BAT results is substantively informative. It indicates that the equity relationship is not an automatic consequence of an increase in strategic import vulnerability. Critical materials map comparatively naturally into upstream extraction and processing activities and downstream industrial uses. Battery-related equity portfolios are more heterogeneous and may include mining, processed materials, battery components, cell manufacturing, charging equipment, and vehicle production.

Battery-sector valuations may also be dominated by technological expectations, subsidies, electric-vehicle demand, capacity expansion, and firm-specific innovation. These influences can weaken the correspondence between a trade-based BAT measure and the returns of broad public equity portfolios. The U.S. BAT null results therefore constrain the scope of the paper's conclusions and show that the CMA findings should not be generalized automatically to other strategic sectors.

The sensitivity of the Chinese BAT results to alternative equity proxies reinforces this interpretation. The CSI new-energy and new-energy-vehicle indices and the alternative CNI battery-industry and Shenwan automobile indices cover overlapping but economically distinct activities. The resulting return spreads are negatively correlated and produce exposure-weighted coefficients with opposite signs. They should therefore not be treated as interchangeable measures of the same upstream–downstream relationship.

Reporting both constructions is preferable to selecting the proxy that produces the strongest result. It also illustrates a broader measurement challenge for research connecting product-level trade vulnerability to financial markets: a coherent strategic-trade classification does not necessarily have a direct counterpart among publicly available equity indices. Firm-level measures of input dependence, sourcing relationships, and revenue exposure would provide a cleaner mapping but are not consistently available over the full sample.

5.4. Implications for Multinational Financial Management

The findings have implications for multinational firms, corporate risk managers, and international investors. For firms that depend on strategic inputs, changes in SSIRI can complement conventional indicators of commodity prices and policy uncertainty by identifying situations in which dependence on China coincides with limited diversification among alternative suppliers. Such information may be relevant to supplier diversification, inventory planning, long-term procurement contracts, capital expenditure, and assessments of production-continuity risk.

The distinction between structural and exposure-weighted vulnerability is also relevant to corporate risk management. Exposure-weighted SSIRI may be more useful for evaluating vulnerabilities associated with large current trade flows and near-term financial exposure. Structural SSIRI may instead draw attention to smaller inputs whose limited substitutability could create disproportionate operational consequences. Firms may therefore benefit from considering both the economic scale of an exposure and the strategic importance of individual inputs.

For international investors, upstream-minus-downstream return spreads provide a way to evaluate how changes in strategic dependence are distributed across supply-chain positions. The results suggest that broad market or sector indices may conceal offsetting valuation effects between input suppliers and input users. SSIRI may consequently be useful as an input to scenario analysis, portfolio-exposure assessment, and the evaluation of supply-chain-sensitive industries. The modest explanatory power of the regressions nevertheless means that it should be considered alongside, rather than in place of, established financial and macroeconomic risk indicators.

The results also have limited policy relevance. Equity-market responses can provide supplementary information about whether changes in measured trade vulnerability are economically salient to investors. They are not sufficient to justify government intervention. Policy assessment requires additional information on inventories, domestic production, supplier capacity, input quality, technological substitutability, allied-country relationships, adjustment costs, and the probability and expected duration of a disruption. Moreover, equity prices may reflect expectations of government support as well as underlying supply-chain conditions. The broader policy-screening framework and limitations of SSIRI are developed in the companion study by Oka (2026a).

5.5. Limitations

Several limitations qualify the interpretation of the findings. First, the regressions identify conditional associations rather than causal effects. Trade exposure, supplier concentration, commodity prices, policy expectations, and equity valuations may respond to common economic information. The main specifications partially alleviate contemporaneous endogeneity concerns by imposing a two-month signal-to-return gap: SSIRI changes formed in month t are related to returns beginning in month $t + 2$. This temporal separation ensures that the monthly customs data underlying the signal are complete before the return window begins, eliminates mechanical overlap between signal formation and return measurement, and reduces concerns about same-month simultaneity and reverse causality. It does not, however, eliminate endogeneity arising from persistent omitted variables, anticipated policy developments, or common news that affects both trade patterns and subsequent equity returns. The factor, commodity, volatility, policy-uncertainty, and geopolitical-risk controls further reduce these concerns but do not provide exogenous variation in SSIRI. Timing placebos support the temporal coherence of the main relationship without establishing causality.

Second, the empirical scope is confined to two strategic sectors and two equity markets. The marked heterogeneity between CMA and BAT underscores that our findings may not extend automatically to semiconductors, medical supplies, defense inputs, or other critical industries. Accordingly, broadening the sectoral coverage and incorporating firm-level measures of cross-border input exposure represent natural priorities for future research.

6. Conclusion

This paper examines whether a trade-based measure of strategic import vulnerability is reflected in sectoral equity returns. Using monthly data from 2004 through 2025, it relates changes in structural and exposure-weighted SSIRI to upstream-minus-downstream return spreads in the United States and China.

The central result is a positive short-run relationship between changes in exposure-weighted CMA vulnerability and the relative performance of U.S. upstream firms. A one-standard-deviation monthly increase in the exposure-weighted index is associated with an approximately 1.1 percentage-point increase in the one-month upstream-minus-downstream spread observed two months later. A related response appears in China over the three-month window beginning at $t+2$. The U.S. decomposition indicates an upstream-led rally, while the Chinese result reflects a combination of stronger upstream and weaker downstream performance.

The exposure-weighted index is more precisely related to equity returns than the equal-family structural measure, particularly when both enter the same specification. In contrast, U.S. BAT results are generally null, and Chinese BAT estimates vary substantially across equity proxies. These findings demonstrate that the financial-market relationship is sector-specific and depends on how trade vulnerability maps into publicly traded industries.

The results contribute to the measurement of strategic dependence by showing that changes in SSIRI have observable financial-market correlates. They do not establish that SSIRI changes cause equity returns or that equity-market responses measure the welfare costs of dependence. Instead, they show that investors appear to revalue upstream and downstream exposure when economically important forms of critical materials vulnerability change.

Future research could strengthen identification by combining firm-level input sourcing with product-level trade exposure, exploiting discrete export-control or supply-disruption events, and constructing real-time vintages of strategic-risk measures. Extending the analysis to additional countries and strategic sectors would also clarify when structural import vulnerability is incorporated into financial-market valuations.

7. Funding sources

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Data availability

Data are drawn from monthly U.S. Census Bureau merchandise-trade records (U.S. Census Bureau, 2026), the U.S. factor files available from the Kenneth R. French Data Library (French, 2026), and daily Chinese index prices compiled by China Securities Index Company and publicly disseminated through Tencent Finance (China Securities Index Company Limited, 2026; Tencent Finance, 2026). A complete replication package containing the aligned public analysis dataset, one self-contained Jupyter notebook, expected-results files, and generated tables and figures will be deposited in a public repository upon acceptance.

Declaration of competing interest

The author declares no known competing financial interests or personal relationships that could have influenced the work reported in this paper.




Figure 1: Monthly Changes in Structural and Exposure-Weighted SSIRI

Notes: The figure plots standardized monthly first differences in the structural and fixed exposure-weighted SSIRI measures for Critical Minerals and Advanced Materials (CMA) and Battery and Energy Storage Materials (BAT) from February 2004 through December 2025.




Figure 2: Upstream, Downstream, and Spread Returns for CMA

Notes: The figure plots the monthly upstream return, downstream return, and upstream-minus-downstream spread used to construct the dependent variables for Critical Minerals and Advanced Materials (CMA). U.S. portfolio-leg returns are measured relative to the U.S. market; Chinese portfolio-leg returns are measured relative to the CSI 300. Returns are in percentage points.




Figure 3: Upstream, Downstream, and Spread Returns for BAT

Notes: The figure plots the monthly upstream return, downstream return, and upstream-minus-downstream spread used to construct the dependent variables for Battery and Energy Storage Materials (BAT). U.S. portfolio-leg returns are measured relative to the U.S. market; Chinese portfolio-leg returns are measured relative to the CSI 300. Returns are in percentage points.

A. Data Sources and Portfolio Construction

Tables 2 and 3 summarize the data sources, coverage, and economic roles assigned to the portfolios. This appendix provides the constituent details needed to reproduce those portfolio definitions.

The U.S. portfolios use the value-weighted 49-industry returns from the Kenneth R. French Data Library. Portfolio membership and weights are fixed before estimation. The BAT upstream portfolio equally weights Mines and Chemicals; its downstream portfolio equally weights Automobiles and Electrical Equipment. The CMA upstream portfolio equally weights Mines, Steel, and Chemicals. The CMA downstream portfolio equally weights Machinery, Electrical Equipment, Automobiles, Aerospace, Defense, Computer Hardware, Semiconductors, and Laboratory Equipment.

For China, the CMA upstream leg is the CSI 300 Materials Index. The downstream leg is an equal-weighted composite of CSI 300 Industrials and CSI 300 Information Technology. Each leg is expressed relative to the CSI 300 market return before decomposition; subtracting the two relative returns yields the same upstream-minus-downstream spread as subtracting the raw leg returns. The original exploratory BAT spread is CSI New Energy minus CSI New Energy Vehicles. The alternative proxy is the CNI Battery Industry Index minus the Shenwan Automobile Industry Index.

The frozen monthly analysis file is supplied because third-party endpoints, ticker histories, and revised provider files can change. The replication notebook reconstructs all first differences, predictor lags, cumulative return outcomes, and regression samples from the raw columns in that file. All factor and equity returns are stored in decimal units; tables multiply estimated return coefficients by 100 and report percentage points.

B. Stationarity and Variable Transformations

Appendix Table 6 reports the stationarity tests that motivate the transformation used in the main analysis. The ADF tests do not reject a unit root at the 5% level for either CMA level, and the KPSS tests reject level stationarity. By contrast, the standardized first differences strongly reject a unit root and do not reject KPSS stationarity. The same qualitative pattern holds for the BAT measures: level persistence is substantial, whereas monthly first differences are suitable for the return regressions.

Alternative autoregressive innovations were also considered. Their ADF tests reject a unit root, but their KPSS statistics continue to reject stationarity for the two CMA measures. First differences therefore provide the cleaner common transformation and avoid conditioning the signal on a particular autoregressive model. In-sample regressions standardize the monthly differences over the corresponding estimation sample. Recursive forecasts use lagged raw differences; with an intercept, rescaling a predictor does not change the OLS forecast.

C. Additional Regression Results and Control Specifications

Table 4 reports the short-horizon estimates and their control sensitivity. The complete output for all four country-sector spreads, both SSIRI measures, four horizons, and three control sets is supplied as `all_main_regressions_32_tests_by_specification.csv` in the replication package. The file reports every coefficient, HAC standard error, test statistic, unadjusted p -value, FDR-adjusted q -value, sample size, and adjusted R^2 . Timing, controls, and multiple-testing families follow the definitions in Sections 3.3–3.7.

D. BAT Results and Equity-Proxy Sensitivity

Complete U.S. BAT estimates are included in the machine-readable regression output described in Appendix C. This section focuses on the proxy sensitivity that limits interpretation of the Chinese BAT estimates.

The Chinese BAT evidence is more directly constrained by proxy validity. Over the common proxy sample, the original CSI new-energy-minus-new-energy-vehicles spread produces exposure-weighted coefficients of 3.034, 3.430, and 2.960 percentage points at 3, 6, and 12 months. The 6- and 12-month estimates survive correction within the original-proxy family. In contrast, the CNI battery-industry-minus-Shenwan automobile spread produces coefficients of -0.529 , -3.179 , and -6.171 percentage points at the same horizons.

The proxy spreads are negatively correlated, with horizon-specific correlations between approximately -0.18 and -0.33 . At 12 months, the alternative-minus-original coefficient difference is -9.131 percentage points and survives correction across the proxy-comparison tests ($q = 0.0365$). These indices therefore represent different economic

Table 6
Stationarity Tests for CMA and BAT SSIRI Levels and First Differences

Sector	SSIRI measure	Transformation	Observations	ADF statistic	ADF p-value	KPSS statistic	KPSS p-value
CMA	Structural	Level	264	-2.8109	0.0567	1.0886	0.0100
CMA	Structural	First difference	263	-16.3881	0.0000	0.1556	0.1000
CMA	Exposure weighted	Level	264	-2.4670	0.1237	0.9372	0.0100
CMA	Exposure weighted	First difference	263	-13.3642	0.0000	0.1111	0.1000
BAT	Structural	Level	264	-2.7229	0.0702	0.6402	0.0190
BAT	Structural	First difference	263	-11.3683	0.0000	0.3920	0.0806
BAT	Exposure weighted	Level	264	-1.4300	0.5678	1.2774	0.0100
BAT	Exposure weighted	First difference	263	-10.7015	0.0000	0.0719	0.1000

Notes: The table reports augmented Dickey–Fuller (ADF) and Kwiatkowski–Phillips–Schmidt–Shin (KPSS) tests for the structural and exposure-weighted SSIRI measures in the CMA and BAT sectors. The level series contain 264 monthly observations from January 2004 through December 2025; first differencing removes the initial observation. Both tests include an intercept but no deterministic time trend. The ADF lag length is selected using the Akaike information criterion, and the KPSS bandwidth is selected automatically.

The ADF null hypothesis is that the series contains a unit root, whereas the KPSS null hypothesis is that the series is stationary around a constant. For all four SSIRI measures, the level tests fail to reject the ADF unit-root null at the 5% level and reject the KPSS stationarity null. In contrast, the first-difference tests strongly reject the ADF unit-root null and do not reject KPSS stationarity at the 5% level. The evidence therefore supports the use of monthly first differences rather than index levels in the equity-return regressions. The first differences are subsequently standardized for estimation; standardization does not alter their stationarity classification.

ADF *p*-values reported as 0.0000 are positive values rounded to four decimal places. KPSS *p*-values of 0.0100 and 0.1000 represent, respectively, the lower and upper bounds of the test's tabulated *p*-value approximation and are not exact probabilities. SSIRI denotes the Structural Strategic Import Risk Index; CMA denotes Critical Minerals and Advanced Materials; and BAT denotes Battery and Energy Storage Materials. Structural SSIRI uses equal product-family weights, whereas exposure-weighted SSIRI uses fixed 2003 import-exposure weights.

themes rather than interchangeable upstream and downstream portfolios. Both constructions are disclosed, but neither is treated as confirmatory evidence. The principal conclusions rest on CMA and the U.S. BAT null.

E. Joint-Measure Regressions, Timing Placebos, and Additional Robustness

Appendix Table 7 extends the headline joint U.S. CMA comparison to horizons from 1 through 12 months. It reports the two coefficients, their individual and joint tests, the direct coefficient-difference test, signal correlations, and adjusted R^2 . Separate Benjamini–Hochberg adjustments are applied to the structural, exposure, coefficient-difference, and joint-test families.

Timing placebos retain the correctly dated signal and add an otherwise identical SSIRI change observed 6 or 12 months in the future. None of the 16 future-signal coefficients survives correction across the placebo family; the smallest global placebo *q*-value is 0.93. The correctly dated coefficients remain economically larger and more precisely estimated. The placebos therefore provide no evidence that the result is generated by a generic correlation with distant future values of the index. They improve the timing audit but do not establish causality.

Additional checks preserve the signal date, the two-month gap, the current spread control, and the $6+h$ Newey–West bandwidth while varying the financial and global-risk controls. The short-horizon U.S. and Chinese CMA exposure-weighted results remain statistically and economically similar across these specifications.

F. Supplementary Pseudo-Out-of-Sample Analysis

Appendix Table 8 reports the complete expanding-window forecast comparison for the U.S. CMA spread under the forecasting design defined in Section 3.8. It presents benchmark and model loss measures, out-of-sample R^2 , and the Clark–West tests for the 120-month baseline training window.

Direct nested comparisons clarify the ranking. The combined model does not improve on the exposure-only model: structural changes do not add beyond exposure changes (one-sided $p = 0.3324$). Exposure changes do add beyond the structural model (one-sided $p = 0.0080$; $q = 0.0161$). Relative to AR(1), the exposure model produces out-of-sample R^2 values of 4.63%, 5.27%, and 5.38% with initial windows of 84, 120, and 156 months.

Table 7
Structural and Exposure-Weighted SSIRI in Joint U.S. CMA Regressions

Horizon	N	Structural beta (pp)	Structural p	Structural q	Exposure beta (pp)	Exposure p	Exposure q	Beta difference (pp)	Difference p	Difference q	Joint p	Joint q	Signal correlation	Adjusted R ²
1	261	0.2909	0.1584	0.3802	0.9081	0.0005	0.0054	-0.6172	0.1111	0.3333	0.0000	0.0003	0.3452	0.0365
2	260	0.3387	0.1554	0.3802	0.5323	0.0502	0.1003	-0.1936	0.6381	0.8401	0.0161	0.0484	0.3436	-0.0036
3	259	0.5305	0.0740	0.3802	0.7548	0.0351	0.0843	-0.2243	0.6754	0.8401	0.0034	0.0204	0.3437	-0.0080
4	258	0.0750	0.7993	0.8720	1.0991	0.0067	0.0400	-1.0241	0.0888	0.3333	0.0059	0.0236	0.3487	-0.0058
5	257	-0.1394	0.6913	0.8296	0.9246	0.0205	0.0615	-1.0641	0.0916	0.3333	0.0566	0.0849	0.3527	-0.0169
6	256	0.5205	0.2157	0.4315	0.7766	0.0711	0.1160	-0.2561	0.7186	0.8401	0.0204	0.0491	0.3570	-0.0078
7	255	0.2334	0.5220	0.7573	0.3653	0.4137	0.4137	-0.1319	0.8367	0.8401	0.4961	0.4961	0.3573	-0.0218
8	254	0.0257	0.9421	0.9421	0.9793	0.0113	0.0451	-0.9536	0.0975	0.3333	0.0334	0.0572	0.3590	-0.0167
9	253	0.5980	0.0771	0.3802	0.7479	0.0870	0.1160	-0.1499	0.8025	0.8401	0.0246	0.0492	0.3593	-0.0155
10	252	0.2204	0.5680	0.7573	0.8207	0.0786	0.1160	-0.6003	0.3480	0.8353	0.1565	0.1878	0.3463	-0.0116
11	251	0.4136	0.2807	0.4811	0.5586	0.3050	0.3327	-0.1449	0.8401	0.8401	0.2616	0.2853	0.3507	-0.0088
12	250	0.6606	0.1167	0.3802	0.8412	0.1068	0.1281	-0.1806	0.7819	0.8401	0.0906	0.1208	0.3505	0.0010

Notes: The table reports horizon-specific joint regressions for the U.S. CMA upstream-minus-downstream return spread. Standardized monthly first differences in structural and exposure-weighted SSIRI, both dated t , enter each regression simultaneously. For horizon h , the dependent variable is the cumulative arithmetic spread return from $t+2$ through $t+h+1$. The specifications therefore maintain a two-month separation between the signal month and the beginning of the return window.

All regressions include the current U.S. CMA return spread, the market excess return, and the size, value, profitability, investment, and momentum factors, all dated t . Newey–West HAC standard errors use $6+h$ monthly lags. Coefficients are reported in percentage points for a one-standard-deviation monthly increase in the corresponding SSIRI change. Structural SSIRI equally weights strategic product families, whereas exposure-weighted SSIRI uses fixed 2003 import-exposure weights.

The beta difference is defined as $\hat{\beta}^{\text{Structural}} - \hat{\beta}^{\text{Exposure}}$. Its p -value tests the null that the two population coefficients are equal. The joint p -value is from a Wald test of the null that both SSIRI coefficients equal zero. Individual-coefficient and coefficient-difference p -values are two-sided. Separate Benjamini–Hochberg corrections are applied across the 12 horizons for the structural-coefficient, exposure-coefficient, coefficient-difference, and joint-test families. Accordingly, the four q -value columns should not be interpreted as one combined correction across all 48 reported tests.

Signal correlation is the sample correlation between the standardized structural and exposure-weighted changes for the corresponding regression sample. Statistical significance of one coefficient but not the other does not, by itself, establish that the two coefficients differ; that inference is based on the reported difference test. Values reported as 0.0000 are positive values rounded to four decimal places. Negative adjusted R^2 values can result from the degrees-of-freedom adjustment. CMA denotes Critical Minerals and Advanced Materials; pp denotes percentage points; and N is the regression sample size.

For the 120-month design, the moving-block bootstrap places 95.6% of the AR(1)-relative out-of-sample R^2 distribution above zero, but the two-sided interval $[-0.008, 0.110]$ includes zero. This distinction is important: the one-sided nested-model test supports forecast improvement, whereas the two-sided bootstrap indicates substantial sampling uncertainty about its magnitude.

The cumulative gains rise over the evaluation period but include episodes of deterioration, showing that the forecast advantage is not uniform over time. Moreover, finalized customs data and subsequent revisions are used. The exercise therefore measures incremental statistical information and is not presented as a real-time trading strategy.

Table 8
Out-of-Sample Forecast Performance for the U.S. CMA Spread

Benchmark	Model	OOS months	Forecast start	Forecast end	Benchmark RMSE	Model RMSE	Model MAE	OOS R2	Clark-West stat.	Clark-West p	Global FDR q
AR(1)	Both changes	141	2014-04	2025-12	0.0434	0.0424	0.0333	0.0469	2.9896	0.0014	0.0217
AR(1)	Exposure change	141	2014-04	2025-12	0.0434	0.0423	0.0334	0.0527	2.6273	0.0043	0.0337
AR(1)	Structural change	141	2014-04	2025-12	0.0434	0.0434	0.0334	0.0029	1.5057	0.0661	0.3171
Historical mean	Both changes	141	2014-04	2025-12	0.0433	0.0424	0.0333	0.0420	2.9102	0.0018	0.0217
Historical mean	Exposure change	141	2014-04	2025-12	0.0433	0.0423	0.0334	0.0479	2.5359	0.0056	0.0337
Historical mean	Structural change	141	2014-04	2025-12	0.0433	0.0434	0.0334	-0.0023	1.3639	0.0863	0.3452

Notes: The table reports expanding-window, one-month-ahead forecasts of the U.S. CMA upstream-minus-downstream return spread. The initial estimation window contains 120 months, and the out-of-sample evaluation period extends from April 2014 through December 2025. For target return r_t , the historical-mean benchmark uses the recursively estimated training-sample mean, while the AR(1) benchmark uses r_{t-1} . The competing models add the structural SSIRI change dated $t-2$, the exposure-weighted change dated $t-2$, or both changes. Consequently, the forecasting exercise preserves the two-month signal-to-return separation used in the in-sample analysis. All model parameters and transformations are estimated recursively using information in the corresponding training sample.

RMSE and MAE are calculated from monthly returns expressed in decimal units. Out-of-sample R^2 is defined relative to the benchmark shown in each row as

$$R_{\text{OOS}}^2 = 1 - \frac{\sum_t (r_t - \hat{r}_t^{\text{Model}})^2}{\sum_t (r_t - \hat{r}_t^{\text{Benchmark}})^2}.$$

A positive value therefore indicates that the SSIRI model has a lower mean squared forecast error than the corresponding benchmark. Clark-West statistics and p -values are based on the one-sided test that the larger, nested model does not improve forecast accuracy; a positive statistic favors the SSIRI model. Global FDR q -values apply the Benjamini-Hochberg correction across the complete 24-comparison forecast family: four country-sector systems, two benchmarks, and three SSIRI models. The table displays the six U.S. CMA comparisons.

The exercise uses finalized monthly trade data rather than real-time data vintages and does not explicitly model the historical publication calendar. It should therefore be interpreted as evidence of incremental statistical forecasting information, not as the performance of an implementable trading strategy. CMA denotes Critical Minerals and Advanced Materials; OOS denotes out of sample; RMSE denotes root-mean-square error; MAE denotes mean absolute error; and FDR denotes false discovery rate.




Figure 4: Cumulative Out-of-Sample Forecast Gains

Notes: The figure plots cumulative differences between the AR(1) benchmark squared forecast error and each SSIRI-model squared forecast error. A rising line indicates lower cumulative squared loss for the SSIRI model. For target return r_t , the benchmark uses r_{t-1} and the predictors are SSIRI changes dated $t-2$. The evaluation uses finalized monthly data.

G. Replication and Computational Details

The replication archive is designed to run without any project-specific Python module. After installing the pinned requirements, the user opens `SSIRI_Equity_Replication.ipynb`, restarts the kernel, and runs all cells. The notebook reads `data/analysis_monthly_public_complete.csv`, which contains 264 monthly observations from January 2004 through December 2025, and writes machine-readable and Overleaf-ready outputs to `outputs/tables/` and `outputs/figures/`.

The timing parameter is explicit and must remain `SIGNAL_TO_RETURN_GAP = 2` to reproduce the manuscript. A signal dated t is paired with the one-month return in $t + 2$, and an h -month outcome sums returns from $t + 2$ through $t + h + 1$. The spread at t and all other controls are dated t ; no variable from $t + 1$ or the dependent return window is used as a regressor.

The notebook reproduces Tables 6, 4, 5, 7, and 8, together with Figure 4. Frozen expected-results files are verification targets rather than substitutes for estimation. Automated checks verify the upstream-minus-downstream portfolio identities, the component-coefficient identity in Equation (10), selected manuscript estimates, and the recursive forecast path. Exact software versions and the random seed are recorded in the package README, requirements file, and reproduction manifest.

The executable core reproduces the original thematic China BAT spread from the frozen public data. Because the alternative AKShare source observations were not retained in the frozen notebook export, their public retrieval commands are supplied in an optional disabled cell. The appendix proxy comparison is archived as a reported sensitivity exercise, but it is not included in automatic endpoint-dependent verification.

References

- Acemoglu, D., Carvalho, V.M., Ozdaglar, A., Tahbaz-Salehi, A., 2012. The network origins of aggregate fluctuations. *Econometrica* 80, 1977–2016. doi:10.3982/ECTA9623.
- Amiti, M., Gomez, M., Kong, S.H., Weinstein, D.E., 2021. Trade Protection, Stock-Market Returns, and Welfare. NBER Working Paper 28758. National Bureau of Economic Research. URL: <https://www.nber.org/papers/w28758>, doi:10.3386/w28758.
- Barrot, J.N., Sauvagnat, J., 2016. Input specificity and the propagation of idiosyncratic shocks in production networks. *The Quarterly Journal of Economics* 131, 1543–1592. doi:10.1093/qje/qjw018.
- Benjamini, Y., Hochberg, Y., 1995. Controlling the false discovery rate: A practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society: Series B (Methodological)* 57, 289–300. doi:10.1111/j.2517-6161.1995.tb02031.x.
- Caldara, D., Iacoviello, M., Molloy, P., Prestipino, A., Raffo, A., 2020. The economic effects of trade policy uncertainty. *Journal of Monetary Economics* 109, 38–59.
- Carvalho, V.M., Nirei, M., Saito, Y.U., Tahbaz-Salehi, A., 2021. Supply chain disruptions: Evidence from the great east japan earthquake. *Quarterly Journal of Economics* 136, 1255–1321. doi:10.1093/qje/qjaa044.
- China Securities Index Company Limited, 2026. Index data and methodologies. <https://www.csindex.com.cn/>. Accessed August 8, 2026.
- Choi, S.Y., 2024. Sectoral responses to economic policy uncertainty and geopolitical risk in the US stock market. *Journal of Multinational Financial Management* 76, 100874. doi:10.1016/j.mulfin.2024.100874.
- Clark, T.E., West, K.D., 2007. Approximately normal tests for equal predictive accuracy in nested models. *Journal of Econometrics* 138, 291–311. doi:10.1016/j.jeconom.2006.05.023.
- Cohen, L., Frazzini, A., 2008. Economic links and predictable returns. *The Journal of Finance* 63, 1977–2011. URL: <https://doi.org/10.1111/j.1540-6261.2008.01379.x>, doi:10.1111/j.1540-6261.2008.01379.x.
- Dickey, D.A., Fuller, W.A., 1979. Distribution of the estimators for autoregressive time series with a unit root. *Journal of the American Statistical Association* 74, 427–431. doi:10.1080/01621459.1979.10482531.
- Ersahin, N., Giannetti, M., Huang, R., 2024. Supply chain risk: Changes in supplier composition and vertical integration. *Journal of International Economics* 147, 103854. doi:10.1016/j.jinteco.2023.103854.
- French, K.R., 2026. Data library. Tuck School of Business, Dartmouth College. URL: https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html. accessed August 8, 2026.
- Gofman, M., Segal, G., Wu, Y., 2020. Production networks and stock returns: The role of vertical creative destruction. *The Review of Financial Studies* 33, 5856–5905. URL: <https://doi.org/10.1093/rfs/hhaa034>, doi:10.1093/rfs/hhaa034.
- Hassan, T.A., Hollander, S., van Lent, L., Tahoun, A., 2019. Firm-level political risk: Measurement and effects. *The Quarterly Journal of Economics* 134, 2135–2202. URL: <https://doi.org/10.1093/qje/qjz021>, doi:10.1093/qje/qjz021.
- Kwiatkowski, D., Phillips, P.C.B., Schmidt, P., Shin, Y., 1992. Testing the null hypothesis of stationarity against the alternative of a unit root: How sure are we that economic time series have a unit root? *Journal of Econometrics* 54, 159–178. doi:10.1016/0304-4076(92)90104-Y.
- Newey, W.K., West, K.D., 1987. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica* 55, 703–708. doi:10.2307/1913610.
- Oka, A., 2026a. Measuring u.s. strategic import vulnerability to china: Theory and evidence from batteries and critical materials. Manuscript under review.
- Oka, A., 2026b. Policy uncertainty and U.S. equity returns: A sector-level analysis of disaggregated international EPU. *Journal of Multinational Financial Management* 82, 100959. URL: <https://linkinghub.elsevier.com/retrieve/pii/S1042444X26000150>, doi:10.1016/j.mulfin.2026.100959.
- Pankratz, N.M.C., Schiller, C.M., 2024. Climate change and adaptation in global supply-chain networks. *The Review of Financial Studies* 37, 1729–1777. doi:10.1093/rfs/hhad093.
- Said, S.E., Dickey, D.A., 1984. Testing for unit roots in autoregressive-moving average models of unknown order. *Biometrika* 71, 599–607. doi:10.1093/biomet/71.3.599.
- Tencent Finance, 2026. Chinese equity index price data. <https://gu.qq.com/>. Accessed August 8, 2026.
- U.S. Census Bureau, 2026. USA Trade Online: U.s. merchandise trade data. <https://www.census.gov/foreign-trade/reference/products/catalog/usatradeonline.html>. Accessed August 8, 2026.