Commentary and Critique **Research Commentary**
Vol. 9 (2026) · ISSN 3006-6220 journal.pe.ac/Commentary-and-Critique

Imputed Weights and Unit Values in Public UN Comtrade Evidence on the Cost of Friend-Shoring

Nargiza Qodirova[a], Matthew Ellerton[a]

[a]Academy of Public Policy and Economics, London E1 0SG, United Kingdom

Corresponding author: Nargiza Qodirova, nargiza.qodirova@pe.ac.

Published 20 March, 2026.

DOI: https://doi.org/10.62432/APPE.CC.2026.9.36

Article webpage: https://journal.pe.ac/Commentary-and-Critique/paper/view/645842

Abstract

The proposition that friend-shoring is costly has been established on the model side: a quantitative multi-country, multi-industry trade model calibrated to 2018 input–output data and shocked with an exogenous 20 per cent rise in trade costs returns real GDP losses of up to 4.7 per cent in some economies, with a world GDP loss of 1.3 per cent in its headline scenario (Javorcik et al., 2024). This article compares observed sourcing shifts with the supplier unit-value evidence available from the public preview of the United Nations Comtrade database, testing what the recorded values and weights establish about the cost of those shifts.

The model literature estimates the cost of imposed friend-shoring scenarios. This article instead measures observed sourcing changes and tests whether public customs unit values can price them. The data yield a clear sourcing result and a binding measurement limit. At chapter level the shift away from China is real and large: across all 97 HS two-digit chapters, China's share of United States import value fell between 2017 and 2024 by 34.5 percentage points in leather goods and 23.2 points in furniture, and Viet Nam's share rose by 12.7 points in furniture and 10.6 points in footwear in the chapters where a comparable 2017 baseline exists. A goods-level panel of 28 six-digit codes and seven years, requested against eight partners of which seven return rows, then returns 1,305 partner-detail rows — of which **1,114 rows, or 85.4 per cent, carry isNetWgtEstimated = true**, so their net weight is an imputed value rather than a reported one. The imputation often behaves like a near-deterministic replay: within commodity-year cells in which every partner row is imputed, the coefficient of variation of the implied unit value averages 0.0380 with a median rounding to 0.0000, and 67 of those 148 cells (45 per cent) have a coefficient of variation below 1e-9, whereas the 15 cells in which no partner row is imputed average 0.8693 and contain no such case. Only seven of the 28 codes contain any reported-weight row at all, and only four of them ever yield a cell in which three or more partners report weight.

In this panel of 28 explicitly selected six-digit codes, the unit-value premium of a sourcing switch is not identifiable, and any computation of one from these rows that does not filter the flag is at substantial risk of reproducing the imputation pattern rather than supplier prices. Whether the same holds across the six-digit universe is unresolved: the panel is what was measured, and Section 5.4 sets out why the universe cannot be enumerated from this endpoint. On the 20 cells that survive the filter the dispersions are large — the widest is +1,406.9 per cent against the largest supplier — but they cannot be read as price differences, because a six-digit customs unit value mixes within-code product composition with price. The chapter-level data establish sourcing shifts. In the goods panel, mean within-cell dispersion differs roughly twentyfold between fully imputed and fully reported cells (coefficient of variation 0.0380 against

© 2026 Academy of Public Policy and Economics · Open Access (CC BY 4.0).

0.8693). Fully imputed cells return premiums with absolute values below 1e-9 (expressed as proportions) for every partner in 56 of the 148 fully imputed cells (38 per cent), while the other 92 (62 per cent) contain at least one premium at or above that absolute-value threshold, reaching +709.6 per cent. These different imputation patterns require separate treatment when interpreting supplier unit values.

Keywords: friend-shoring; sourcing reallocation; unit values; UN Comtrade; net weight imputation; trade statistics; import measurement

1. Introduction

Friend-shoring — the reorganisation of supply towards politically aligned partners — moved from slogan to policy instrument within a few years, and economics responded first where response is easiest: in models. The most careful of those models takes inter-country, inter-industry linkages seriously, calibrates to 2018 OECD input–output tables aggregated to 39 countries or country groups and 16 industries, and implements friend-shoring as an exogenous increase of 20 per cent in tariffs or in iceberg trade costs (Javorcik et al., 2024). Its headline is a real GDP loss of up to 4.7 per cent in some economies, with world GDP falling 1.3 per cent in the scenario the authors foreground. The same literature is explicit about its own boundary conditions: the model has no time dimension and permits no formation of new trade partnerships, so the authors describe the result as medium-run (Javorcik et al., 2024).

That modelling literature answers the question ”what would friend-shoring cost if it were imposed at this magnitude?” It does not, and does not claim to, answer the question ”what has the sourcing that has already switched actually cost?” The latter is an observational question, and answering it requires trade records at a resolution that names the supplier, the product and the price. The obvious object of measurement is the customs unit value — declared value divided by declared quantity — and that object is not exotic: a recent study in this literature uses import unit values as the dependent variable and regresses them on the similarity of non-tariff measures across partners, which establishes more narrowly that customs unit values are used as an outcome in this field (Nakhoda, 2025). The question this article asks is narrower and, it turns out, harder: if the United States bought less from China and more from Viet Nam in a given six-digit product, did it pay a different unit value for what it bought from Viet Nam?

The question is not merely rhetorical. A systematic review of the near-shoring and friend-shoring literature shows a field mature enough to be reviewed (da Rocha et al., 2025); firm-level work examines whether geopolitical tension actually changes foreign sourcing, supply-base complexity and sub-tier supplier sharing, but its title names no price or unit-value variable (Charpin & Cousineau, 2024); third-country effects of the United States–China tariff escalation are studied as reallocation away from China (Hayakawa & Sudsawasd, 2025); and critical-mineral friend-shoring has its own institutional and governance literature (Vivoda, 2023) (Vivoda & Matthews, 2023). The closest empirical neighbour found in this project’s search examines the impact of source diversification on China’s soybean import prices — the same question form, ”does changing where you buy change what you pay”, at single-commodity, single-importer resolution (Zheng & He, 2024).

This article’s contribution is therefore neither a new country nor a new year and, after the second narrowing below, not even a number for the premium. It is a measurement result with a falsifiable test: on the public Comtrade preview, the unit value used to price a sourcing switch is, in the large majority of partner-detail rows, an imputed quantity, and the imputed rows frequently reproduce a near-common per-commodity unit value; therefore an unfiltered premium cannot be interpreted as a supplier price difference. The test is run below and it holds. The article also states what that implies operationally, because the implication is not pessimistic for the field: it is the filter that has to be applied before any such computation means anything.

Two boundaries are recorded at the outset so that nothing later is read more widely than it was measured. First, the chapter-level result in Section 3 is a result about *shares of import value*; it establishes that sourcing moved, not what the move cost. Second, the goods-level panel in Sections 4 to 6 covers 28 explicitly selected six-digit codes, not the universe of traded goods, and the reasons the universe cannot be enumerated are documented in Section 2.

2. Data channel, and what it will not do

2.1. The source

All trade figures in this article come from the public preview endpoint of the United Nations Comtrade database, https://comtradeapi.un.org/public/v1/preview/C/A/HS, queried without credentials and stored as raw JSON responses before any reduction (United Nations Statistics Division, 2026). The reference table that the service itself publishes was used for country codes rather than recalled from memory. Annual imports were extracted for calendar years 2017, 2018, 2019, 2021, 2022, 2023 and 2024; 2020 is omitted throughout.

That year list crosses a change in the classification version the rows carry. Every raw row has a classificationCode field, and it is not constant across the list: it is H5 for 2017, 2018, 2019 and 2021 and H6 for 2022, 2023 and 2024 — the codes the service associates with the HS2017 and HS2022 versions of the nomenclature, respectively, though no first-hand UN/WCO HS correspondence table was obtained to confirm this mapping independently. The 2022 version is also the one under which the two goods added in that revision exist at all. Every comparison in this article that sets a 2017 to 2021 figure against a 2022 to 2024 figure therefore compares rows recorded under two different versions of the nomenclature, and not merely two different years. Whether a six-digit code that appears under both versions denotes the same commodity basket in both is not established here: no HS correspondence table was obtained, so the size of any reclassification effect on the figures below cannot be confirmed. The two cross-version comparisons are flagged where they are made, in Sections 3 and 7.

The extraction was performed in 2026-09 and the stored responses, not a derived table, are the input to every number printed here.

The literature cut-off for this article is 2026-02-20. No cited item was created or published after that date, and for each cited item the registry's own creation timestamp is recorded in the accompanying verification record.

2.2. The reporter code: 842, not 840

For the United States as reporting country the code is **842**, and the trap is worth stating because it is silent. The service's published reporter reference table contains 255 entries and includes 842 with the description "USA"; it contains no entry for **840** at all. A query with reporterCode=840 does not fail: it returns HTTP 200 with count = 0 and an empty data list, which is indistinguishable in form from a genuine statement that there was no trade. The same is true when 840 is used as a partner: the 842-by-840 cell returns zero rows for every year examined. A reader who assumed the ISO 3166 numeric code would report, without any error message, that the United States neither imported nor was imported from. Every figure in this article uses 842.

2.3. What the endpoint returns, and what it does not

Three properties of the preview determine the design of the panel in Section 4, and each was established by direct request rather than assumed.

First, the response is **capped at 500 rows and offers no pagination**. A request for the aggregate six-digit universe returns count = 500 and exactly 500 rows; page, offset, limit and top are silently ignored

and repeated requests return the same 500 rows. A request for two-digit chapters as a group (AG2) does return all 97 chapters, and that is why chapter-level coverage in Section 3 is complete — but every one of those 97 rows carries netWgt of zero or null, so chapter-level data can carry shares and cannot carry unit values. A request for four-digit groups (AG4) is truncated at 500 rows and is not the full set of four-digit codes.

Second, an explicitly named commodity code returns only the aggregate row itself, not its children. The panel below therefore names its six-digit codes explicitly. Comma-separated code lists work — a two-code list returns two rows — and comma-separated partner lists work — a three-partner list returns three rows. Two neighbouring syntaxes do not: partnerCode=all is rejected with HTTP 400 as an invalid parameter value, and a multi-period request is rejected with the message that the maximum number of periods for the preview is 1.

Third, the preview is **rate-limited more tightly than its own documentation suggests**. Interval tests between consecutive calls produced HTTP 429 at spacings of 2.05 and 2.29 seconds; the collection underlying this article therefore used a 5-second interval, and even then recorded throttling responses that had to be retried. Two chapter-level cells in the sweep were left unfilled by this and are documented in Section 3 rather than silently imputed.

Finally, a limitation that matters for units. The response carries a monetary value field (primaryValue) and, on import rows, an identical cifvalue, so the value is a landed (cost, insurance and freight) value paid by the importer; this was verified row by row, not assumed. But the response does not include a field naming the monetary unit. UN Statistics Division documentation for the same processing system states that values reported in national currencies are converted to United States dollars before dissemination (United Nations Statistics Division, 2020). The USD labels therefore follow official processing documentation; scale and ordering provide only a consistency check. The service also publishes reporterDesc and partnerDesc fields, and on preview responses both are present with the value null; country names therefore come from the reference table.

2.4. The import of all this

The channel can support chapter-level share arithmetic with complete coverage of the 97 chapters. It can support a goods-level panel only if the goods are named explicitly and the row count stays under the cap. It cannot support enumeration of the six-digit universe, and it does not generate a chapter-level unit value. Those are the constraints within which the rest of the article works.

3. The shift is real at chapter level

The first question is whether there is anything to price. Table 1 answers it with complete coverage: one request per partner and year returns all 97 two-digit chapters, and the shares below are computed from the World row (partner code 0) as denominator and the China row as numerator, both read directly from the stored responses.

Table 1. China's share of United States import value by two-digit chapter, and Viet Nam's change, 2017 to 2024.

HS chapter	US imports 2024 (USD bn)	China share 2017	China share 2024	Change	Viet Nam change
42 Leather articles	14.3	57.2%	22.7%	−34.5 pp	+1.8 pp
94 Furniture, bedding	72.6	52.0%	28.8%	−23.2 pp	+12.7 pp
64 Footwear	28.4	55.8%	36.2%	−19.6 pp	+10.6 pp
84 Machinery	531.2	32.2%	16.0%	−16.2 pp	+4.6 pp

(continued)

Table 1 (continued)

HS chapter	US imports 2024 (USD bn)	China share 2017	China share 2024	Change	Viet Nam change
85 Electrical machinery	485.9	42.1%	26.1%	−16.0 pp	+5.7 pp
62 Apparel, not knitted	36.3	35.7%	21.4%	−14.4 pp	+5.9 pp
61 Apparel, knitted	47.4	32.5%	22.4%	−10.1 pp	+2.7 pp
83 Miscellaneous base metal	16.1	43.6%	34.9%	−8.6 pp	+2.5 pp
15 Animal and vegetable fats	17.2	0.8%	7.4%	**+6.7 pp**	+0.6 pp
96 Miscellaneous manufactures	9.1	50.5%	55.7%	**+5.1 pp**	+0.7 pp

Note. shares are of the World record for the same chapter and year. Value is in USD billions under the official processing documentation cited in Section 2.3. The Viet Nam column is a change in Viet Nam's share of the same chapter between the same two years.

Three things follow, and they are the only three things this table establishes. First, the reallocation is large and is not confined to marginal chapters. Among the 55 of the 97 chapters whose 2024 United States import value is USD 5 billion or more, the five steepest falls include the two largest chapters in the table, machinery at −16.2 points and electrical machinery at −16.0 points; across all 97 chapters those two rank 11th and 12th by the size of the fall. Second, it is not universal. Chapter 15 rose 6.7 points and chapter 96 rose 5.1 points, so "the United States is de-China-fying its imports" is not a statement this data supports; the correct statement is that reallocation is chapter-specific. Third, Viet Nam is the only alternative supplier whose change can be measured across the whole sweep: for every other partner the 2017 baseline is missing from the collection, in most cases because those cells were never requested, and in the case of Mexico because the collected 2017 response is a throttling body rather than data. It follows that the standard summary sentence "Viet Nam is the main taker" is **not** established at chapter level by this sweep; what is established is that Viet Nam's share rose in the chapters where it can be measured, most sharply in furniture and footwear. Ranking takers requires a partner set with comparable baselines, and Section 5 supplies one at goods level within a restricted universe.

The boundary of this section is therefore exactly as stated at the outset: Table 1 shows that sourcing moved. It says nothing about what the move cost, and chapter rows cannot say anything about cost, because their net weights are zero.

A second boundary of Table 1 belongs beside that one. Its two endpoint years are read from rows recorded under different versions of the nomenclature (Section 2): the 2017 column is read from HS2017 rows and the 2024 column from HS2022 rows. The same 97 two-digit headings return rows under both versions, so the chapter list itself is unchanged. Whether the two versions place the same goods under the same chapter is a different question, and it is not established here, because no correspondence table was obtained. The chapter-level changes in Table 1 are therefore changes between two versions of the nomenclature as well as between two years.

4. A goods-level panel designed to price the shift

To reach unit values, the panel must name six-digit codes. The codes were selected from chapters where Table 1 shows the largest reallocation, supplemented by two goods that friend-shoring policy names explicitly, lithium-ion accumulators and photovoltaic modules. The result is 28 codes and seven years, requested against eight partners — China, Mexico, Viet Nam, Japan, Korea, Germany, Canada and Taiwan

— as 49 successful grouped calls with a 5-second interval and retries after throttling. Taiwan returned no row in any cell, so the effective panel is seven partners; that is the second silent-empty case in this article and it is reported rather than dropped. A supplementary request shows that the flow is not absent from the channel: for reporter 842, year 2023 and code 847130, partner 158 returns no row while partner 490, "Other Asia, nes", returns one row with a landing value of USD 2.09 billion. Why the flow appears under 490 and not under 158 is not established by anything this study read, and the panel does not reconstruct it.

The design produced **1,305 partner-detail rows**. The 28 codes over seven years give 196 possible commodity–year cells, of which 188 return data — the two 2022-nomenclature codes return nothing before 2022 — and the returning cells carry seven, six or five partners rather than the eight that were requested. That is the whole of the gap between 1,305 rows returned and the 1,568 partner-cells requested. Each row was checked for duplicates: the 1,305 rows are 1,305 distinct year–commodity–partner combinations. This matters because this endpoint family is known to split one flow across carriage, customs-procedure and second-partner dimensions; in this panel every row carries motCode = 0, customsCode = C00 and partner2Code = 0, so no such split is present and summing rows across partners would not double-count. That check is reported rather than left implicit, since a naive row sum on this endpoint family can overstate a flow several-fold.

5. The measurement limit: imputed weights

5.1. The flag, and its coverage

Every row carries a quality flag, isNetWgtEstimated. In this panel:

Table 2. Rows returned, and the weight-quality flag.

Item	Value
Goods (six-digit codes)	28
Years	7 (2017–2019, 2021–2024)
Partners per request	8
Partners that return rows	7
Commodity–year cells with data	188 of 196
Potential year–commodity–partner rows	1,568
Rows actually returned	1,305
Rows with isNetWgtEstimated = true	**1,114 (85.4%)**
Rows with isNetWgtEstimated = false	191 (14.6%)
Rows with net weight missing or zero	139 (10.7%)
Goods with at least one reported-weight row	7 of 28
Goods that ever yield a cell with three or more reporting partners	4 of 28

The flag is not a corner case. More than four rows in five do not yield a reported net weight, and the 21 goods that never produce one are not peripheral: they include laptops, toys, wooden furniture, electric cables and photovoltaic modules.

5.2. The imputation pattern is often a near-replay

The decisive question is what an imputed weight does to the unit value. The answer came from the raw rows and is stronger than a suspicion. Consider a commodity–year cell — one six-digit code in one year — and classify it by the flags of its partner rows: fully imputed if every partner row with a weight is imputed, fully reported if none is, mixed otherwise. Compute the coefficient of variation of the implied unit value within the cell.

Table 3. Within-cell dispersion of the implied unit value, by weight status.

Cell class	Cells	Mean CV	Median CV	Cells with CV below 1e-9
Fully imputed	148	0.0380	**0.0000**	67 (45%)
Fully reported	15	0.8693	0.9125	0
Mixed	5	0.6473	0.6067	0

The contrast is two orders of magnitude. In fully imputed cells the dispersion collapses toward zero; in fully reported cells it approaches one. And zero here is not ”small”: 67 of the 148 fully imputed cells are identical to machine precision, and 130 commodity–year cells in the panel — including laptops, toys, wooden furniture and electric cables in every year in which they appear — return a single identical unit value for every partner. The 2019 lithium-ion cell is the clearest single illustration: seven partners, one unit value, 60.976 dollars per kilogram. Laptops are the mirror image: all seven of their commodity–year cells return a single identical unit value for every partner, and not one of their rows carries a reported weight.

The pattern supports a bounded inference. If the weight of a partner row were imputed from a single per-commodity unit value, then dividing declared value by that weight would return that same unit value for every partner. That is what is observed, in the class where it must be observed, and not observed in the class where it must not be. A unit value computed from an imputed row may therefore reproduce an **imputation parameter** and cannot, without provider documentation or reported quantity, be treated as information about the supplier’s price. The mechanism by which the service performs this imputation is not documented in anything this study read; the characterisation rests on the observed zero dispersion and is stated as an inference from the data, not as a citation to a method note.

5.3. What this does to the intended measurement

The intended measurement was a supplier-to-supplier unit-value premium: for each commodity and year, take the largest supplier by landing value as the reference, and compare each other partner’s unit value with it. Run on this panel without filtering, that computation would return a premium for every cell. In fully imputed cells the premium is frequently, but not uniformly, degenerate: 67 of the 148 fully imputed cells have a coefficient of variation below 1e-9, Table 3’s threshold for effectively zero dispersion, and within that set 56 cells (38 per cent of the 148) have every partner premium within an absolute tolerance of 1e-9, with premiums expressed as proportions. The remaining 92 cells (62 per cent) contain at least one premium at or above that absolute-value threshold, up to +709.6 per cent (HS420292, 2018). The premium tolerance and the CV threshold are different criteria; neither defines literal equality to zero. That pattern is still the signature of the replay described in Section 5.2 rather than a finding about supplier prices, but it does not make every fully imputed premium zero by construction. The quantity the article set out to measure is therefore not identifiable in this panel of 28 explicitly selected six-digit codes, and the reason is a property of the data, not a shortage of sample. Whether it is identifiable elsewhere on this channel is a separate question and is unresolved here: the evidence covers 28 codes, not the channel’s commodity space, and the six-digit universe cannot be enumerated from this endpoint (Section 5.4).

5.4. What cannot be claimed about generality

The 85.4 per cent figure is a statement about the 28 explicitly selected codes. It is **not** a statement about the six-digit universe, and this article does not make one. The reason is mechanical: the aggregate six-digit request is truncated at 500 rows with no pagination, so the universe cannot be enumerated from this endpoint. Similarly, the accompanying sum of landing value for the 28 codes — USD 228.3 billion in 2023 across the seven partners that report — is a total for those partners, not for United States imports as a whole, and it is printed here only to size the panel. There is also a version limit on the panel itself: its

2017 to 2021 rows are recorded under the HS2017 version of the nomenclature and its 2022 to 2024 rows under the HS2022 version (Section 2). A commodity observed in both spans is therefore observed under two versions, and no correspondence table was obtained with which to check that the two versions mean the same basket.

6. What the reported-weight subset can and cannot show

Filtering to rows whose weight is reported leaves very little: seven of the 28 goods contain any such row, and only four — cotton knitted T-shirts, cotton knitted pullovers, cotton woven men's trousers and lithium-ion accumulators — ever yield a cell in which three or more partners report weight. There are 20 such cells in total, of which 15 contain no imputed row at all and the remaining 5 are mixed.

Table 4. Unit-value spread within cells where at least three partners report weight.

HS6	Year	Supplier	Partners	Unit value (USD/kg)	Lowest premium	Highest premium
610910	2017	China	7	12.453	−19.8%	+1,116.9%
610910	2018	China	7	13.179	−23.7%	+987.4%
610910	2021	China	7	8.915	+26.5%	**+1,406.9%**
610910	2022	Viet Nam	7	19.104	−45.0%	+570.6%
610910	2023	Mexico	7	12.994	−30.3%	+925.2%
611020	2017	China	6	18.415	−48.3%	+758.8%
611020	2018	China	7	17.984	−41.6%	+878.1%
611020	2021	Viet Nam	7	20.197	−39.3%	+844.0%
611020	2022	Viet Nam	7	20.982	−30.2%	+865.9%
611020	2023	Viet Nam	7	19.854	−38.9%	+996.1%
620342	2017	Mexico	7	11.444	−0.3%	+489.1%
620342	2018	Mexico	7	11.716	+1.4%	+456.6%
620342	2021	Mexico	7	12.773	−63.5%	+448.1%
620342	2022	Mexico	6	14.718	+52.3%	+382.4%
620342	2023	Mexico	7	14.590	−35.2%	+579.3%
850760	2017	China	6	41.223	−11.2%	+224.9%
850760	2018	China	6	39.269	+15.6%	+328.5%
850760	2021	China	7	22.360	+39.4%	+131.1%
850760	2022	China	7	20.888	+8.4%	+137.5%
850760	2023	China	6	14.245	+73.9%	+267.2%

Note. the premium is each other reporting partner's unit value divided by the largest supplier's unit value, minus one. The USD/kg reading follows the official processing documentation cited in Section 2.3.

Two readings must be given together, and neither may be dropped. The first is that reported-weight cells are not uninformative: dispersion is real, and the spread between partners reaches 1,406.9 per cent in the widest case. The second is that a spread of that size cannot be a price difference. Cotton knitted T-shirts under a single six-digit code differ in fabric, yarn count, cut, branding and packing, and those differences dominate a customs unit value; a factor of fifteen between suppliers of the same code is a statement about what was in the box, not only about what it cost. Even on the cleanest subset, therefore, what is measured is a **landed unit-value difference**, and the article does not describe it as a same-quality price difference. This is a hard boundary, and it is the reason the article does not report a premium as its headline: Table 4 shows how much of the apparent "premium" a six-digit code can manufacture on its own.

One further observation is possible without overreach. The identity of the largest supplier within these four goods changes across the period — China in the T-shirt and pullover cells early, Viet Nam in the pullover cells from 2021, Mexico in the trouser cells throughout — which is at least consistent with Table 1's chapter-level shift. It is not offered as evidence of it.

7. A comparable-baseline check on takers

Because the chapter sweep has no comparable baselines for partners other than Viet Nam, the question of who took the share China lost is asked again inside the panel, where all seven reporting partners are present in both endpoint years. The comparison is restricted to the 26 codes whose six-digit strings return rows in both 2017 and 2024, so that the same list of codes forms the aggregate in both years. The two 2022-nomenclature goods, which carry USD 50.1 billion of 2024 landing value between them, are excluded for that reason alone. The restriction holds the code list constant; it does not hold the commodity definition constant, and no claim is made here that the denominator keeps the same composition. The 2017 endpoint is recorded under the HS2017 version of the nomenclature and the 2024 endpoint under the HS2022 version (Section 2), and 26 of the panel's 28 codes return rows under both versions. Whether a code that appears under both versions denotes the same commodity basket in both is therefore unresolved — no HS correspondence table was obtained — and the size of any reclassification effect on the shares below cannot be confirmed. Table 5's changes are changes between two versions of the nomenclature as well as between two years.

Table 5. Shares of the seven-partner panel aggregate, 26 goods common to both years.

Partner	Share 2017	Share 2024	Change
China	70.3%	53.7%	−16.6 pp
Viet Nam	7.3%	19.1%	**+11.7 pp**
Mexico	11.3%	13.8%	+2.5 pp
Canada	4.6%	5.1%	+0.5 pp
Korea	2.5%	3.9%	+1.4 pp
Japan	2.4%	2.6%	+0.3 pp
Germany	1.6%	1.9%	+0.3 pp

Note. Taiwan was requested and returned no row in any cell, so the aggregate covers the seven partners that report. It is USD 139.7 billion in 2017 and USD 179.4 billion in 2024. It is not United States imports; it is the sum of the seven partners that report over the 26 goods.

Within this restricted universe, Viet Nam is the largest gainer by a wide margin: no other tracked partner gains more than 2.5 points, and Viet Nam's gain covers about seven-tenths of China's loss. This is a narrower statement than the one the data cannot support, and it is the widest version of "Viet Nam is the main taker" that this evidence will bear.

8. Scope of supported claims

The archived panel supports the reported chapter-level sourcing shifts, the 1,114-of-1,305 estimated-weight count, and the dispersion contrast between fully imputed and fully reported cells. It shows that imputed rows often collapse toward a near-common unit value. It does not identify the provider's algorithm, a supplier price premium, the economic cost of friend-shoring, or the pattern outside the selected 28 codes.

9. Discussion

9.1. What the result is

The result is a measurement limit, cleanly located. In the panel read from the public preview of UN Comtrade, the observed price of a sourcing switch is not recoverable at partner-detail resolution, because the quantity that would convert landed value into a unit value is imputed in 85.4 per cent of the rows in this panel, and because the imputed rows often collapse toward a common unit value within commodity and year, obscuring precisely the between-partner variation that the measurement needs. The test that

establishes this is falsifiable and was run: dispersion must be near zero in fully imputed cells and clearly positive in fully reported ones, and it is 0.0380 against 0.8693.

9.2. What the result is not

It is not a claim that no one can price friend-shoring. It is not a claim that the Comtrade database is unusable; its chapter-level coverage is complete and its goods-level rows are usable once filtered. It is not a claim that the reallocation observed at chapter level is friend-shoring rather than tariffs, logistics, cost or demand — this article measures where purchases came from, not why. And it is not a claim that the neighbour literature is wrong: the model-based study contains no unit-value measurement at all, so nothing here contradicts it. What this article contradicts is a data practice, not a paper.

9.3. The operational implication

The implication for anyone attempting this measurement is concrete and cheap. Read isNetWgtEstimated on every row; discard or flag rows where it is true; report the size of the sample that survives; check that every requested partner actually returns rows, since a partner code can be valid and silent; and if that sample is small — here, four of 28 goods and 20 cells — say so in the abstract rather than in a footnote. Any premium computed without the filter should be labelled as potentially driven by the publisher's imputation rather than supplier prices. Where a genuine measurement is wanted, the data requirement is stated by the failure: reported quantities at a finer product resolution than six digits, with the reporting country's own weight declarations, rather than a reconstructed weight.

9.4. What would change the answer

Two things would. A documented method note from the data provider explaining when and how net weight is imputed would move the characterisation in Section 5.2 from inference to citation. And an official bulk extract with complete six-digit coverage and reported weights would move the generality claim in Section 5.4 from "these 28 goods" to the universe. Neither was available within this project's access to the data channel, and neither is asserted here.

10. Declarations

Evidence cut-off. 2026-02-20. No cited item was created or published after that date; the registry creation timestamp for each cited item is recorded in the accompanying verification file. The trade data are a live service and were extracted in 2026-09; that extraction date is a property of the data collection, not of a cited source.

Reproducibility. Every number printed in this article is re-derived from the raw stored responses by a single script, which recomputes the chapter shares, the panel composition, the flag coverage, the dispersion statistics, the premium table and the panel-level share table, and exits non-zero if any printed value fails to reproduce. The script reads the stored raw responses rather than derived tables.

Figures. None. The article uses tables only.

11. Tables

Tables 1 to 5 appear in Sections 3, 4, 5, 6 and 7 respectively, with their notes.

Statements

Funding. This work received institutional support from the Academy of Public Policy and Economics (APPE).

Competing interests. The authors are affiliated with APPE, which operates this journal. No external competing interests are declared.

References

Charpin, R., & Cousineau, M. (2024). Friendshoring: How geopolitical tensions affect foreign sourcing, supply base complexity, and sub-tier supplier sharing. *International Journal of Operations & Production Management*, *45*(5), 1006–1031. https://doi.org/10.1108/ijopm-01-2024-0067

da Rocha, A., da Fonseca, L. N. M., & Kogut, C. S. (2025). Deciphering relocation paths: A systematic literature review of near-shoring and friend-shoring. *Journal of International Management*, *31*(5), 101282. https://doi.org/10.1016/j.intman.2025.101282

Hayakawa, K., & Sudsawasd, S. (2025). Trade effects of the US–China trade war on a third country: Preventing trade rerouting from China. *The World Economy*, *48*(10), 2333–2343. https://doi.org/10.1111/twec.70014

Javorcik, B., Kitzmüller, L., Schweiger, H., & Yıldırım, M. A. (2024). Economic costs of friendshoring. *The World Economy*, *47*(7), 2871–2908. https://doi.org/10.1111/twec.13555

Nakhoda, A. (2025). Similarity of NTMs and its impact on import unit values. *The World Economy*, *48*(6), 1333–1350. https://doi.org/10.1111/twec.13702

United Nations Statistics Division. (2020). *Request for 2019 annual and 2020 monthly merchandise trade statistics and trade by enterprise characteristics*. United Nations Department of Economic and Social Affairs. https://unstats.un.org/unsd/trade/IMTS/IMTS%20Data%20Request%20EN_2019.pdf

United Nations Statistics Division. (2026). *UN Comtrade database: Public preview API (preview/C/A/HS endpoint)*.

Vivoda, V. (2023). Friend-shoring and critical minerals: Exploring the role of the Minerals Security Partnership. *Energy Research & Social Science*, *100*, 103085. https://doi.org/10.1016/j.erss.2023.103085

Vivoda, V., & Matthews, R. (2023). "Friend-shoring" as a panacea to Western critical mineral supply chain vulnerabilities. *Mineral Economics*, *37*(3), 463–476. https://doi.org/10.1007/s13563-023-00402-1

Zheng, X., & He, M. (2024). The impact of source diversification on soybean's import prices of China. *Resources Science*, *46*(7), 1358–1369. https://doi.org/10.18402/resci.2024.07.09