

# When Tariff Averages Misrank Markets: A Business Intelligence Audit of WTO 2025 Profiles

Seyfettin Artan<sup>1,\*</sup>

<sup>1</sup> Department of Economics, Karadeniz Technical University, Trabzon, Turkiye

Email: [artan@ktu.edu.tr](mailto:artan@ktu.edu.tr)

Received: August 02, 2026 Revised: September 14, 2026 Accepted: October 07, 2026 ★ Corresponding author

## ABSTRACT

Procurement teams often rank destination markets using a single published tariff average. Whether that ordering survives a change of measurement is a practical question for business intelligence. We audit the 2025 *World Tariff Profiles*, retaining 101 economies with matching 2024 most-favoured-nation (MFN) applied tariffs, reported trade-weighted tariffs, and 2023 import totals. A reproducible screen compares simple and trade-weighted averages, ranking changes, duty-free subheading shares, and the share of subheadings facing duties above 15%. The median simple average is 7.5% and the median trade-weighted average is 6.9%; the within-economy weighted-minus-simple median is –0.6 percentage points. Thirty-five economies differ by at least two points. With ties retained at each top-25 cutoff, 28 economies qualify under the simple rate and 26 under the weighted rate, but only 20 qualify under both. Türkiye falls from 17.3% to 7.0%, whereas Dominica rises from 10.7% to 15.9%. Disagreement persists among larger import markets: 13 of 31 economies with at least US\$50 billion in 2023 imports have a gap of two points or more. A layered dashboard should therefore trigger product-level verification where aggregate measures disagree. Neither average identifies the duty paid by a particular firm or the effect of changing trade policy.

**Keywords:** Business intelligence ▪ Tariff analytics ▪ Market screening ▪ Import sourcing ▪ Decision support

## 1. INTRODUCTION

A firm selecting a sourcing location or export market may begin with a table of national tariffs. The table compresses thousands of product lines into one number per destination, saving the analyst's time at the first screening stage. But a ranking built on a simple average of tariff lines can differ from a ranking that gives more weight to products a country actually imports. Neither statistic is a transaction quote. When procurement analysts mistake either for a landed-cost forecast, a seemingly inexpensive market may be assigned too little attention, or a viable market may be excluded too early.

The decision question is specific: *how often does a firm's initial market-priority list change when a simple tariff average*

*is replaced by the published trade-weighted average?* We audit that choice in a defined set of economies. The resulting diagnostic identifies where analysts should request a product-level rate; it cannot estimate the rate payable on a particular shipment.

## 2. RELATED WORK AND RESEARCH GAP

Business analytics research emphasizes the path from integrated data to a defensible decision. In a survey of 305 firms, Szukits and Móricz [1] associate perceived data quality and managerial support for analytics with an analytical decision culture; Liu and Yoon [2] design a goal-driven method for integrating data around a user's decision need. These contributions motivate a transparent sourcing dashboard, but neither

assesses whether two legitimate summaries of the same tariff schedule put different economies in a manager's review queue. Our analysis treats that disagreement as an auditable property of the information presented to decision makers.

Trade studies show why the omitted product and firm detail matters. Handley et al. [3] link tariff uncertainty to firms' adoption of imported inputs, while Alessandria et al. [4] document inventory and entry responses to uncertainty about tariff treatment. Poilly and Tripier [5] construct a regional uncertainty measure using sectoral tariff volatility and the composition of imports. At the firm level, Greenland et al. [6] infer heterogeneous exposure to trade liberalization from equity-market responses. These studies estimate behavioral or economic effects with information beyond the country summaries available here. They support treating the import basket and sector mix as relevant to the decision, without licensing a causal interpretation of the difference between our two reported averages.

The applicable border charge is also more specific than a national MFN mean. Firms may weigh the cost of complying with rules of origin against a tariff penalty [7]; the effects of preferences vary across products and countries [8]. Import processing adds costs that tariff rates alone omit [9]. We therefore address a distinct gap between trade-cost evidence and business intelligence practice: a reproducible test of how published aggregate indicators change a market screen, coupled with a rule for handling flagged cases to product-level verification. The 2025 WTO, International Trade Centre, and UN Trade and Development profiles supply the common tariff vintage and definitions needed for that test [10].

### 3. EVIDENCE AND DIAGNOSTIC DESIGN

#### 3.1 A defined comparison set

We read the all-products summary tables and Part A.1 of the individual profiles in *World Tariff Profiles 2025* [10]. The unit is an economy, not a firm or an import shipment. We retain a profile only if the report gives a 2024 MFN applied simple average, a 2024 MFN applied trade-weighted average, a 2023 import total, and corresponding 2024 summary-table shares for duty-free subheadings and duties exceeding 15%. This complete-case rule yields 101 economies. The report covers more economies overall; the results below describe the 101 retained economies, not a census of global trade. In particular, they are unweighted by economic size unless an import threshold is stated.

The WTO defines the simple MFN mean from HS six-digit subheading averages. Its trade-weighted mean weights HS six-digit MFN averages using imports from the most recently available year. The separate import-total field in all retained profiles is dated 2023. We use it only for sensitivity groups, not to recompute tariff weights. Duty-free and above-15% indicators are percentages of HS six-digit subheadings, with partial subheadings allocated pro rata under the report's method; they are *not* percentages of import value. The all-products simple rate in each retained profile reconciles to the corresponding summary entry to the published tenth of a point. Thus every comparison uses figures reported for the same economy and tariff vintage, while the tariff and import reference years remain visible.

#### 3.2 What counts as a changed decision

Let  $S_i$  and  $W_i$  denote economy  $i$ 's reported simple and trade-weighted MFN applied averages, in percent. The descriptive difference  $D_i = W_i - S_i$  is measured in percentage points. Negative  $D_i$  means imports give greater weight to lower-tariff products than the equal-subheading average; positive  $D_i$  means the reverse. This is a statement about the published composition, not evidence that changing a duty would cause firms to change their purchases by  $D_i$ .

$$D_i = W_i - S_i, \quad R_i = |r_i(S) - r_i(W)|, \quad (1)$$

where  $r_i$  is a descending rank among the 101 economies, with average ranks for ties. We report the count with  $|D_i| \geq 2$  points, membership above each top-25 boundary, including every economy tied at that boundary, and crossings of a 5% screening cutoff. The two-point and 5% boundaries are transparent illustrations of a dashboard rule, not estimated optimal thresholds. Spearman's rank correlation describes overall similarity; the actual membership changes describe the queue a manager would see. We also recompute the disagreement count among markets with at least US\$10 billion, US\$50 billion, and US\$100 billion in reported 2023 imports, thereby testing whether the result depends entirely on small markets.

### 4. WHERE THE SCREENS DISAGREE

#### 4.1 Rank stability has a material tail

The two measures agree broadly, with Spearman's  $\rho = 0.901$ , yet the level and decision comparisons in Figure 1 show why a high correlation does not settle a procurement decision. The median of the 101 simple averages is 7.5%; the median of their weighted counterparts is 6.9%. The median of the *paired* differences is  $-0.6$  points, distinct from subtracting the two sample medians. Sixty-six differences are negative, 28 positive, and seven exactly zero at reported precision. The median absolute difference is 1.1 points, while 35 economies differ by at least two points.

The median absolute midrank movement is seven places. Türkiye moves 45 places, the largest change in the retained set. At the boundary of the 25 highest rates, rounding creates ties: the simple-rate cutoff of 12.0% admits 28 economies, and the weighted-rate cutoff of 9.9% admits 26. Twenty belong to both sets; eight qualify only by the simple measure and six only by the weighted measure. Retaining all boundary ties prevents an alphabetical or file-order rule from deciding membership. Six economies cross the illustrative 5% screen, including Chinese Taipei (6.3% simple; 2.0% weighted) and Montenegro (3.7%; 5.8%). Rankings are relative to this selected set; adding countries with missing fields could change every position. The underlying percentage-point differences, which do not rely on sample ranks, remain the stronger basis for individual-market follow-up.

Table 1 demonstrates why the direction of the gap matters. The Turkish simple average of 17.3% exceeds the reported trade-weighted average of 7.0% by 10.3 points, yet 35.2% of its HS6 subheadings face applied rates above 15%. A buyer whose product sits in that tail cannot infer a 7.0% duty from the weighted figure. Dominica's weighted measure instead exceeds its simple rate by 5.2 points; its 2023 imports are




**Figure 1.** A four-part audit of 101 economies. (A) Each dot compares published 2024 MFN applied averages; the dashed line denotes equality. (B) Economy counts by weighted-minus-simple tariff difference; boundary values are assigned by the code in the reproducibility package. (C) Duty-free and above-15% indicators describe HS six-digit subheadings rather than trade value. (D) Absolute descending midrank changes within this sample. Source: authors' calculations from WTO et al. (2025).

only US\$0.3 billion. These examples illustrate different types of triage: Turkey requires product-level verification despite a low weighted rate relative to its simple rate, while Dominica may matter greatly to a specialized supplier even though its aggregate market is small.

**Table 1.** Selected discrepancies in published MFN averages

| Economy            | <i>S</i><br>rate (%); <i>D</i> in pp | <i>W</i> | <i>D</i> | > 15%<br>HS6 |
|--------------------|--------------------------------------|----------|----------|--------------|
| Türkiye            | 17.3                                 | 7.0      | −10.3    | 35.2         |
| Korea, Republic of | 13.4                                 | 8.5      | −4.9     | 10.5         |
| Viet Nam           | 9.5                                  | 5.1      | −4.4     | 23.8         |
| Dominica           | 10.7                                 | 15.9     | +5.2     | 25.4         |
| Belize             | 11.9                                 | 16.1     | +4.2     | 26.8         |

*S*: simple mean; *W*: trade-weighted mean; *D* = *W* − *S* in percentage points. Last column: share of HS6 subheadings subject to applied duties above 15%, not share of import value. Source: WTO et al. (2025); calculations in accompanying files.

#### 4.2 Product breadth and market size

The median duty-free share is 31.4% of HS6 subheadings and the median share facing applied rates above 15% is 12.9%. The distribution in Figure 1C confirms that a national mean omits the breadth of zero-rate and high-rate lines. The two shares describe different regions of the tariff schedule, so neither can be converted to a firm's cost without its HS classification and relevant sourcing arrangement. Even a low average can coexist with a commercially important tail. Similarly, the published agricultural and non-agricultural weighted




**Figure 2.** Selected economies with large changes or a 5% screen crossing. Lines connect two reported MFN applied averages; they are not time trends or estimates of tariffs paid. Source: WTO et al. (2025).

means do not always move together: 56 economies have a higher weighted than simple mean for agricultural products, while 75 have a lower weighted than simple mean for non-agricultural products. A sourcing mix therefore matters more than a generic “country tariff” label.

Disagreement remains visible when the analysis excludes smaller import markets. Among 67 economies reporting at least US\$10 billion in 2023 imports, 25 have  $|D| \geq 2$  points and the median difference is  $-0.8$  points. Among 31 with at least US\$50 billion, the corresponding count is 13 and the median is  $-1.7$  points. Eleven of 20 economies above US\$100 billion pass the same disagreement threshold; their median is  $-2.0$  points. These are nested sensitivity groups, not representative samples or a weighted estimate of the world's trade. They show that the gap cannot be dismissed as an artefact

confined to tiny importers. For example, Türkiye, Korea, and Viet Nam each report more than US\$300 billion in 2023 imports and a weighted rate at least 4.4 points below the simple rate.

#### 4.3 Two independent reasons to investigate

A wide gap between  $S$  and  $W$  flags instability in the national ordering. A large share of HS6 subheadings above 15% flags breadth of comparatively high statutory rates. These are complementary questions: the former compares how tariff averages are constructed, and the latter summarizes the underlying schedule. For an illustrative review grid we set  $|D| \geq 2$  percentage points and an above-15% share of at least 25%. Table 2 summarizes the resulting counts for the 101 economies. The numerical cutoffs are management triggers, not welfare criteria. They are chosen to make the screening logic observable and are kept fixed across the sample.

**Table 2.** A two-signal tariff review grid (101 economies)

| $ W - S $         | HS6 subheadings above 15% duty |              |
|-------------------|--------------------------------|--------------|
|                   | Below 25%                      | At least 25% |
| Below 2 points    | 49                             | 17           |
| At least 2 points | 14                             | 21           |

All four cells count economies, not shipments. The HS6 measure uses the WTO's prorated definition; the horizontal tariff gap is in percentage points. Source: WTO et al. (2025); authors' calculations.

The grid separates the two reasons for investigation. Twenty-one economies trigger both signals; 14 trigger only the mean-disagreement signal, and 17 trigger only the high-duty-breadth signal. The latter 17 would not be selected by a policy that looks only for large differences between averages. Conversely, a 2-point difference does not imply that a quarter of all product subheadings have duties above 15%. Both indicators should therefore remain visible, not be collapsed into a score with an arbitrary exchange rate between them. For a business with known product codes, the grid is only a queue for checking those codes; it does not classify the firm's exposure as high or low.

The grid also guards against an intuitive but incorrect reading of the negative gap. A weighted rate lower than the simple rate need not mean the tariff schedule is benign. For India, the simple and weighted all-products rates are 16.2% and 12.0%, respectively, but 32.5% of HS6 subheadings exceed 15%. Türkiye likewise triggers both signals, whereas Korea's 4.9-point negative gap is accompanied by a smaller 10.5% above-15% share. These market cases point to distinct due-diligence tasks: identify the imported product line and its legal duty, then ask whether the national line distribution or observed import basket has any bearing on the proposed order. The report does not contain an origin-specific preferential rate for that order.

#### 4.4 Sensitivity to the screen's threshold

The two-point flag is deliberately a managerial convention. If the allowable difference were one point, 52 of 101 economies would be flagged; at three points, 19 would be flagged. The interpretation does not change: the thresholds regulate the analyst's workload rather than establish the true cost of importing. Reporting all three counts lets a firm decide how

many cases its trade-compliance team can review. The same caution applies to the 5% market-entry cutoff. Twenty-four economies fall at or below 5% on the simple average and 28 on the weighted average, but six change sides; a count of qualifying markets alone would conceal the opposing switches.

Two very large gaps may dominate a casual reading of the scatterplot. Excluding the Bahamas (−17.5 points) and Türkiye (−10.3 points) leaves 99 economies: their median paired difference remains −0.6 points and 33 still have an absolute gap of at least two points. The gap is therefore more than a pair of extreme observations, though it is not uniform. Sector means strengthen this reading. The agricultural weighted-minus-simple median is +0.7 points; the non-agricultural median is −0.8 points. These within-sector summaries should not be combined by taking an unweighted mean of two sector rates: the product counts and import composition differ, and the report already publishes the appropriate all-products measures.

### 5. A DEFENSIBLE BUSINESS INTELLIGENCE WORKFLOW

The audit suggests a three-stage workflow for an import or export planning team. First, maintain the simple average, the trade-weighted average, the gap, and the report year as separate database fields. A wide  $|D|$  raises a review flag; it does not automatically remove a destination from the sourcing universe. Second, display the duty-free and above-15% HS6 shares alongside the market's import scale. This helps an analyst distinguish a broad tariff schedule from one whose mean conceals strong product variation. Third, before a quote or contract decision, replace the country-wide screen with the firm's product code, current tariff schedule, origin, applicable preference, and other border and processing charges. The published MFN means cannot supply any of those shipment-specific details.

The difference between  $S$  and  $W$  is particularly easy to misuse. Trade-weighted rates describe products that entered the market; the composition of those imports may itself be affected by tariffs, demand, production patterns, and trade agreements. A low  $W$  can reflect heavy imports in zero-duty lines without proving that a prospective firm's good will face a low duty. Treating an MFN screen as evidence of preferential access would erase a relevant difference. Conversely, a high simple average does not mean every prospective transaction is expensive. Market intelligence is improved when the dashboard records why a case was flagged and which exact line-level information can clear it.

The rule is operationally testable. A procurement manager can log which products prompted a 2-point difference flag, then compare their verified applied rates with the initial market ordering. The feedback belongs at the product and origin level, with dated tariff versions, purchase quantities, and preference eligibility. If the pipeline yields repeated false alarms for one product family, the organization can refine its screening threshold using its own realized decisions; the present cross-section cannot estimate such a threshold. For longer-term commitments, the analyst should separately monitor tariff-policy uncertainty. This diagnostic does not model future changes.

## 6. FROM SCREEN TO PROCUREMENT EVIDENCE

### 6.1 A concrete review queue

The two-signal grid can be converted to a manageable sequence without pretending to optimize a firm's purchases. The 21 economies in its upper-right cell enter first if the trade team needs to address both ranking instability and a broad high-duty tail. The 17 economies with high-duty breadth but a small mean gap require a different check: the averages are similar, yet the prospective product may still belong to a costly subheading. The 14 with a large gap but less than one-quarter of subheadings above 15% require attention to the difference between the import basket and the full tariff schedule. The 49 remaining economies receive standard product-level confirmation before contracting, rather than an assurance that no risk exists. This ordering is a work-allocation example; a firm's own purchase volume and product concentration can reverse its priority.

That last point is especially relevant to the large-market cases in Figure 1. A commercial team serving one product category should not average its exposure across all product lines or all 101 economies. For the Republic of Korea, the overall simple average is 13.4% and the weighted figure is 8.5%; the published agricultural rates are 57.0% simple and 89.6% weighted, while the corresponding non-agricultural rates are 6.5% and 3.5%. The all-products mean offers very different guidance to a food importer and an industrial input buyer. The apparent precision of an all-products figure is thus a consequence of aggregation, not evidence that it is precise for either procurement category. A manager should record the category and the rate actually relevant to that category before any cash-flow calculation.

A useful exception ticket would contain five fields: the market considered, the firm's actual HS code, the proposed country of origin, the dated MFN tariff line and any supported preference claim, and the intended shipment value. The two national averages and their source date remain attached as the reason the ticket entered the queue. The customs team verifies the line and legal treatment; finance then estimates the transaction's cost using the verified rate and the relevant value basis. Separating these responsibilities prevents an analyst from applying a nationally trade-weighted percentage to the firm's purchase order. Processing expenses, documentation, freight, and timing have their own entries; treating them as zero simply because the tariff profile does not report them would defeat the landed-cost calculation [9].

### 6.2 How the signal should be evaluated

A dashboard is successful if it changes the quality of a decision, not if it produces a large number of red cells. The firm could preserve each initial market shortlist, record whether the two-signal grid sent a case to review, and archive the line-level verification returned by the customs specialist. The most direct process measure is the proportion of flagged cases that yield a materially different line-level assessment or alter the shortlist after verification. That outcome would be observed in the firm's own future cases. We do not calculate it from the country profiles, because the WTO publication has no record of the firm's HS codes, counterfactual shortlist, verification cost, or final orders.

The appropriate comparison is an internal historical or prospective workflow in which the same candidate markets are first ranked under a simple-average-only procedure and then reviewed with the extra indicators. Any time or accuracy gain depends on the case mix and the capacity of the trade team. A prospective evaluation would set its materiality criterion before looking at outcomes, use the same purchase-value basis in both procedures, and distinguish corrected false exclusions from corrected false approvals. Otherwise a dashboard could appear valuable simply because analysts chose difficult markets after seeing the signal. The current paper supplies a replicable screening benchmark and a specified verification handoff, not an estimate of organizational return on analytics investment.

### 6.3 Why the denominator matters

The above-15% measure counts HS six-digit subheadings, with the report's pro-rata treatment of partially covered lines; it does not measure the share of import value paying high duties. The trade-weighted average instead uses imports allocated across product categories. Their different denominators make the indicators complementary. If a country's imported value is concentrated in duty-free products, its weighted rate can remain low while a sizable fraction of the schedule has high duties. If a firm plans to import an item from that high-duty fraction, the weighted rate ceases to be an informative approximation for that firm.

This distinction also blocks a causal interpretation of the negative median gap. Observed imports may be concentrated in low-tariff lines because buyers favor such lines, because a country lacks domestic supply in those categories, because demand is concentrated there, or because of many other factors. The cross section does not identify the separate contribution of any mechanism. It cannot support a claim that firms saved 0.6 percentage points by choosing a particular origin, or that a country would gain import volume if its simple average converged to its weighted average. The stated results are a diagnostic of reported measures and of the priorities they produce when used in a business screen.

## 7. INTERPRETATION AND LIMITS

The main result is a measurement disagreement, not a treatment effect. WTO profiles report national MFN tariff schedules and import-weighted counterparts at HS six-digit resolution. Trade-weighting uses imports from the latest available year, while our import-scale filter uses the profiles' 2023 total. Published rates are rounded to one decimal place, so exact zero gaps and near-threshold cases are sensitive to reporting precision. We have no transaction microdata, product-level values within the report, preference utilization, or importer-level landed costs. We cannot tell whether a firm's specific sourcing choice would change if it were shown both measures, or quantify the money saved by adopting this dashboard.

The complete-case sample is another constraint. Some economies in the report lack a matching 2024 weighted rate or the other fields required for our screen; their omission can affect cross-economy ranks. The top-25 overlap is therefore deliberately described as a decision comparison *within 101 economies*. The point differences for named economies are directly traceable to their published profiles, and the import-

size checks show the broad observation persists in sizable markets. Still, a firm must update the underlying tariff before any live decision. The July 2025 report and every cited study were available by December 2025. Later policy revisions are outside this historical comparison.

### 7.1 Scope of the market comparison

Each economy contributes one observation regardless of import size. The import-size checks ask whether disagreement survives the exclusion of small markets; they do not yield a global shipment-weighted tariff. Nor should the reported 2023 total be assumed to equal the underlying product-flow weights: WTO uses the most recently available HS six-digit flows.

Dominica, with US\$0.3 billion in 2023 imports, has the largest positive gap at 5.2 points. This changes its position on a country screen but cannot be scaled to a global cost effect. Market size and product relevance should remain separate dashboard filters.

Country profiles also report agriculture and non-agriculture separately. Their aggregation matters when an analyst compares firms in different sectors. In India's profile, the agricultural simple mean is 36.7% and the trade-weighted agricultural mean is 64.3%, even though the all-products weighted rate is below the all-products simple rate. The signs differ because the sector-specific and combined distributions answer different questions. The headline negative all-products gap therefore does not justify a claim that an agricultural buyer would encounter lower charges. A firm that already knows its sector can use the published sector panel as a more relevant intermediate check, then proceed to the actual tariff line and origin. In the absence of a known sector, the only defensible statement is about the report's aggregate statistics.

Finally, the 2024 tariff vintage is a historical snapshot of statutory MFN measures, not a standing quotation for a purchase order written later. Changes in national tariff schedules, trade remedies, origin qualification, and import processing arrangements can each invalidate a historical screen for a present transaction. The proposed system would store the archived WTO profile for reproducible research and link future decisions to a freshly checked tariff schedule. This separation allows a manager to audit how a shortlist was made at the time without relabeling an old profile as live legal advice. It also keeps the conclusion appropriately bounded: the study identifies fragile rankings in a published data product and specifies what additional evidence a firm needs before acting.

## 8. CONCLUSION

A business intelligence system can make the discrepancy visible and route consequential cases to product-level customs review. This is a transparent, testable screening procedure; neither aggregate average is a firm-specific tariff or a forecast of commercial performance.

A simple tariff average is a useful first sorting device, but it is not a stable complete ranking of market conditions. In the 101 comparable WTO 2025 profiles, 35 economies have at least a two-point gap between simple and trade-weighted 2024 MFN rates, the two tie-inclusive high-rate groups have only 20 economies in common, and the mismatch is present among large import markets.

## DATA AVAILABILITY

The source is *World Tariff Profiles 2025*, available at [https://www.wto.org/english/res\\_e/booksp\\_e/world\\_tariff\\_profiles25\\_e.pdf](https://www.wto.org/english/res_e/booksp_e/world_tariff_profiles25_e.pdf).

## REFERENCES

- [1] A. Szukits and P. Móricz, "Towards data-driven decision making: The role of analytical culture and centralization efforts," *Review of Managerial Science*, vol. 18, pp. 2849–2887, 2024.
- [2] D. Liu and V. Y. Yoon, "Developing a goal-driven data integration framework for effective data analytics," *Decision Support Systems*, vol. 180, p. 114197, 2024.
- [3] K. Handley, N. Limão, R. D. Ludema, and Z. Yu, "Firm input choice under trade policy uncertainty," *Journal of International Economics*, vol. 150, p. 103909, 2024.
- [4] G. Alessandria, S. Y. Khan, and A. Khederlarian, "Taking stock of trade policy uncertainty: Evidence from china's pre-wto accession," *Journal of International Economics*, vol. 150, p. 103938, 2024.
- [5] C. Poilly and F. Tripier, "Regional trade policy uncertainty," *Journal of International Economics*, vol. 155, p. 104078, 2025.
- [6] A. Greenland, M. Ion, J. Lopresti, and P. K. Schott, "Using equity market reactions to infer exposure to trade liberalization," *Journal of International Economics*, vol. 152, p. 104000, 2024.
- [7] K. Head, T. Mayer, and M. J. Melitz, "The laffer curve for rules of origin," *Journal of International Economics*, vol. 150, p. 103911, 2024.
- [8] F. Forge, J. Garred, and K. L. Kwon, "When are tariff cuts not enough? heterogeneous effects of trade preferences for the least developed countries," *Journal of International Economics*, vol. 152, p. 103971, 2024.
- [9] J. Carballo, A. G. Graziano, G. Schaur, and C. Volpe Martincus, "Import processing and trade costs," *Journal of International Economics*, vol. 154, p. 104060, 2025.
- [10] World Trade Organization and International Trade Centre and UN Trade and Development, "World tariff profiles 2025," World Trade Organization, 2025.