

# A product-level Trade-LCA framework for embodied carbon accounting in global trade

Florian Pothin

`florian.pothin@univ-rennes.fr`

Université de Rennes

---

## Method Article

**Keywords:** embodied carbon, international trade, life-cycle assessment, carbon border adjustment, input–output analysis, product-level accounting

**Posted Date:** September 10th, 2026

**DOI:** <https://doi.org/10.21203/rs.3.rs-10627706/v1>

**License:**   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

**Additional Declarations:** Competing interest reported. The author is employed by Toovalu (Nantes, France) under a CIFRE doctoral agreement. Toovalu develops carbon accounting software and provided access to the ecoinvent database used in this study. Toovalu had no role in the design of the study, the analysis, or the decision to publish. The author has no other financial or non-financial competing interests.

---

# A product-level Trade-LCA framework for embodied carbon accounting in global trade

Florian Pothin

Université de Rennes, CREM UMR 6211, Rennes, France

Toovalu, Nantes, France

`florian.pothin@univ-rennes.fr`

## Abstract

Carbon accounts are organized by territory or final demand at sectoral resolution; border measures and procurement standards need product-level information. We build a product-level Trade-LCA framework linking bilateral HS6 trade flows to life-cycle emission factors (ecoinvent cut-off and Agribalyse, IPCC 2013 GWP 100a) across 5,000 products, 120 countries and 1995–2023. A large language model assists the matching from customs nomenclature to life-cycle processes, restricted to classification among documented candidates; to our knowledge it is the first such pipeline with measured rather than asserted reliability. A blind expert audit of 100 stratified products found 17 wrong assignments; repairing the filter with deterministic rules, including a processing-stage ladder read from the nomenclature, brings a second blind audit to 8 wrong in 100 (1.2% of emissions). The audited account puts world embodied carbon in traded goods at 14,368 MtCO<sub>2e</sub> in 2023, with 71.7% of emissions on direct product-specific evidence and the remainder on declared fallbacks, of which 0.6% of emissions is monetary. It reconstructs 96% of the benchmark consumption-minus-production wedge, captures 90% of the EXIOBASE footprint at the emissions-weighted median, and matches an independent material database at a median ratio of 1.04. Carbon intensity varies thirteen-fold *within* the median input–output sector, against thirty-five-fold before the repair: part of the dispersion unaudited accounts report is matching error. A sector average still mis-estimates, by more than a factor of two, products carrying 45.5% of traded emissions and 59.9% of trade value.

**Keywords:** embodied carbon; international trade; life-cycle assessment; carbon border adjustment; input–output analysis; product-level accounting. **JEL classification:** F18; Q56; Q54; C67; H23.

# 1 Introduction

Production-based accounts assign emissions to the territories where they occur, while consumption-based accounts reallocate emissions to the final demand that drives production [Davis and Caldeira, 2010, Peters et al., 2011]. The gap between these two principles is the carbon embodied in international trade, and it is large: for major economies it can exceed a quarter of the production-based inventory [Hertwich and Peters, 2009, Peters and Hertwich, 2008]. Emissions embodied in global trade have also grown substantially since 1990, as production has fragmented across global value chains [Peters et al., 2011]. Quantifying this gap at fine resolution has become a policy priority. Instruments such as the EU Carbon Border Adjustment Mechanism increasingly target individual traded products [European Parliament and Council, 2023, Böhringer et al., 2022, OECD, 2023].

Environmentally extended multi-regional input-output models provide macroeconomic consistency. They are the dominant approach for estimating consumption-based carbon footprints [Tukker and Dietzenbacher, 2013, Lenzen et al., 2013, Stadler et al., 2018]. Their sectoral resolution, however, averages over products with very different physical characteristics and carbon intensities. Within a single input-output sector, the carbon intensity of individual products can differ by more than an order of magnitude. Sectoral averages are therefore poorly suited to product screening, border measures and low-carbon procurement [Minx et al., 2009].

Carbon-account resolution is also inadequate geographically [Davis et al., 2025] and in the choice of life-cycle database [Gudipudi et al., 2026]; product resolution within the sector is the third axis, and it is the one this paper measures. Life-cycle assessment supplies product-specific evidence at this resolution. But LCA databases have not been systematically connected to global bilateral trade panels at customs-product detail [Lenzen, 2000, Suh et al., 2004, Crawford, 2008].

Product-level embodied carbon accounting has a direct precedent in Sato [2014a] and Kanemoto et al. [2016] (see also Sato, 2014b), which mapped monetary intensities onto trade categories for a single year; the framework here attaches process-LCA cradle-to-gate factors to physical quantities at HS6, bilaterally, over twenty-nine years, organized around an auditable evidence hierarchy [Gaulier and Zignago, 2010].

Services are not observed at product resolution, so goods are the core of the paper; a monetary-intensity services extension is in the Supporting Information.

Aggregation bias in embodied-carbon accounts is established at input-output resolution [Lenzen, 2011, Su et al., 2010, Steen-Olsen et al., 2014, de Koning et al., 2015]; what has not been available is a measurement of the dispersion below the finest input-output partition, in physical units, across the full traded-product space (Supporting Information Section S10 positions this claim in that literature).

The paper addresses four questions. First, can HS6 trade flows be linked to life-cycle

emission factors at global scale with transparent evidence layers, and which products and countries dominate embodied carbon exposure in trade? Second, and centrally, how much does carbon intensity vary *within* input-output sectors, and how much embodied trade carbon does a sector average, the resolution of MRIO models, therefore mis-estimate? Third, how do the resulting product-level Trade-LCA adjustments compare with benchmark consumption-based accounts, and where do they diverge? Fourth, how robust are these estimates to the evidence layer, emission-factor uncertainty and outliers?

One element of the implementation carries a claim of its own. The link from HS6 customs lines to life-cycle processes is built with the assistance of a large language model, used for classification and disambiguation among documented candidate processes and for nothing else. A literature now applies such models to life-cycle data. It matches free-text product descriptions to emission factors [Balaji et al., 2023a,b, 2025], and it names background-data matching as a grand challenge these models can address [Tu et al., 2024]. The same literature documents hallucinated inventory values [Preuss et al., 2024] and unstable performance in inventory gap-filling [Zhao et al., 2025]. To our knowledge, no work in this line addresses customs nomenclature, and none publishes a measured error rate. This paper does both, in deliberately narrow terms. Blind expert coding of stratified samples measures the matching at 17 wrong assignments in 100 before repair, and 8 in 100 after. The repair itself is a set of deterministic rules, not model judgement (Section 3.4). The cautionary half of that measurement is part of the result: a model-based pre-adjudication of match quality agreed with the blind expert at chance level. The reliability of a model-assisted pipeline therefore cannot be certified by the model itself. Throughout, we publish the audited account and state its coverage honestly, because a wrong factor is worse than a declared fallback.

The factors are representative-technology intensities, so the account measures the *composition* of embodied trade and its dispersion is a lower bound on true heterogeneity; where its sign differs from a consumption account, the difference diagnoses where production technology matters most (Section 5.5).

## 2 Product-level Trade-LCA framework

### 2.1 Accounting logic

The unit of observation for the goods component is a bilateral product flow,

$$(i, j, p, t), \tag{1}$$




Figure 1: **The product-level Trade-LCA framework.** Bilateral HS6 trade flows (BACI) are linked to life-cycle emission factors (ecoinvent, Agribalyse) through product harmonization, unit checks and an auditable four-layer evidence hierarchy; product-level embodied imports and exports are then aggregated to country-level trade-adjusted production emissions, validated against the EXIOBASE footprint, the ICE database and CBAM default values.

where  $i$  is the exporter,  $j$  is the importer,  $p$  is the traded product and  $t$  is the year. For each bilateral product flow, embodied emissions are estimated as

$$E_{i,j,p,t} = Q_{i,j,p,t} \times EF_{p,r,t}, \quad (2)$$

where  $Q_{i,j,p,t}$  is the traded quantity and  $EF_{p,r,t}$  is the assigned life-cycle emission factor for product  $p$ , exporter or region  $r$ , and year  $t$ .

Country-product imports and exports are

$$\text{Imported}E_{c,p,t} = \sum_i E_{i,c,p,t}, \quad (3)$$

$$\text{Exported}E_{c,p,t} = \sum_j E_{c,j,p,t}. \quad (4)$$

The country-level goods trade adjustment is

$$\text{Adj}_{c,t}^{\text{goods}} = \text{Imported}E_{c,t}^{\text{goods}} - \text{Exported}E_{c,t}^{\text{goods}}. \quad (5)$$

We construct country-level trade-adjusted production emissions as

$$\text{PBE}_{c,t}^{\text{Trade-LCA}} = \text{PBE}_{c,t} + \text{Adj}_{c,t}^{\text{goods}}. \quad (6)$$

An optional services term, documented in the Supporting Information, adds to Equation (6); the core analysis is goods-only.

This accounting logic preserves product-level structure in the trade component. National production-based emissions remain the territorial baseline. The result measures product-level embodied trade exposure and country-level trade-adjusted production emissions. It




Figure 2: **What the account counts, next to the frames it complements.** Production-based accounting (PBE) records emissions where they occur, at territorial resolution, and includes the combustion of fuels at the burner. Consumption-based accounting (CBE) reallocates PBE to final demand through the Leontief inverse of an input–output table, at country by sector resolution. The account built here weights bilateral product shipments by cradle-to-gate life-cycle factors (IPCC 2013 GWP 100a), at country by country by HS6 resolution; mineral fuels carry upstream factors only, because their combustion is already in the importer’s territorial account. The net country adjustment closes to PBE and the world total is conserved by construction; gross flows measure exposure. The bottom strip states what each frame can and cannot see.

does not allocate domestic production emissions to products. The emission factors  $EF_{p,r,t}$  are largely product-generic (Section 3.3). The trade adjustment therefore varies with *what* and *how much* a country trades, holding production technology at representative intensities. It is a composition measure. The exporter-region index  $r$  enters only through the limited subset of factors for which geographically specific evidence is available.

Figure 2 places the account among the two frames it complements. Production-based inventories count emissions where they occur, and consumption-based accounts reallocate them to final demand at sector resolution; the account built here counts the carbon that products carry across borders, and its net country adjustment closes to the territorial ledger.

## 2.2 Evidence-layer hierarchy

The framework assigns each traded flow to one mutually exclusive evidence layer (Figure 1). The first layer uses direct HS6 physical matches, where the HS6 product description, quantity unit and emission-factor unit are compatible and the matched donors survive the admissibility filter of Section 3.4.

Below it sit a stage-matched physical fallback (a code whose donors the filter rejected inherits the median of surviving family codes at the same processing stage), plain HS4/HS2/global physical fallbacks, and a monetary fallback where physical evidence is unavailable or fails the validity screen. All outputs retain the assignment layer so that

results can be filtered by evidence quality.

## 2.3 Gross and net embodied trade exposure

Gross embodied trade exposure measures the scale of embodied carbon moving through a country-product position:

$$\text{GrossTradeE}_{c,p,t} = \text{ImportedE}_{c,p,t} + \text{ExportedE}_{c,p,t}. \quad (7)$$

Net embodied trade exposure measures direction:

$$\text{NetTradeE}_{c,p,t} = \text{ImportedE}_{c,p,t} - \text{ExportedE}_{c,p,t}. \quad (8)$$

Gross and net exposure answer different questions. Net exposure identifies whether a country or product position is a net importer or exporter of embodied carbon. Gross exposure identifies positions with large two-way flows. Those positions can stay transition-relevant even when imports and exports offset in net terms.

Because each product carries its full cradle-to-gate factor, summing gross flows double-counts the shared upstream of goods traded at several processing stages, so gross exposure is a risk measure and the net adjustment is the accounting-consistent one (Supporting Information Section S10).

## 3 Data and implementation

### 3.1 BACI trade data

We use BACI as the empirical backbone of the international goods trade network. BACI reports harmonized bilateral trade flows by exporter, importer, product and year using Harmonized System product codes and reconciled mirror trade statistics [Gaulier and Zignago, 2010]. Each observation includes a trade value and, where available, a physical quantity and unit. The final account covers approximately 5,000 HS6 product categories, 120 countries and the period 1995–2023. The source dataset is BACI HS92 V202501, comprising 258,605,562 bilateral trade flows.

### 3.2 Product harmonization

BACI supplies all flows in a single fixed nomenclature (HS92), which removes the need for a concordance [Pierce and Schott, 2012, Van Beveren et al., 2012]; products affected by definitional drift within codes are aggregated to the lowest level that preserves temporal comparability.

### 3.3 Life-cycle emission factors

Life-cycle emission factors convert traded quantities into greenhouse gas emissions. The factors are process-level data. Industrial and material products use *ecoinvent* cradle-to-gate processes; agri-food products (HS chapters 1–24) use the ADEME *Agribalyse* database; both come through the Base Empreinte distribution [Wernet et al., 2016, Koch and Salou, 2020]. The ecoinvent extract uses the **cut-off** system model, with cumulative life-cycle impact assessment at the process level. It is not the APOS or consequential variant.

All results are characterized with IPCC 2013, climate change, GWP 100a, the horizon on which every benchmark below is built; an earlier version of this account had inherited the extract’s superseded CML 2001 GWP 500a column by first match, and the re-characterization, applied at record level and re-accumulated from flow level, raises the pre-repair 2023 total by 11.8% (Supporting Information Section S11 documents the provenance and the product-level heterogeneity of the switch). The upstream boundary for mineral fuels excludes the combustion of traded fuels, which territorial accounts already carry. Factors are expressed per kilogram, per physical unit, or, only where no physical factor exists, per monetary unit. They cover the cradle-to-gate stages from raw-material extraction to the factory gate. We assign exporter- or region-specific factors where they exist, and we fall back to global averages otherwise. Specific factors are the exception: 67% of the underlying records carry a generic geography. The factors are therefore representative-technology intensities, as discussed above. Arbitration and validation initially accepted direct physical factors for 4,667 HS6 codes; the admissibility filter of Section 3.4, which is a result of this paper, retains 3,022 of them, and the remainder fall to the stage-matched fallback layer.

Each code’s factor is the median across the admissible records mapped to it, and the characterization is applied at record level (98.3% of accepted records resolve to a named process; Supporting Information Section S11).

### 3.4 HS-LCA matching

The mapping procedure links customs product descriptions to life-cycle emission-factor categories. HS product descriptions and emission-factor descriptions are cleaned by normalizing case, removing non-informative terms, standardizing units and extracting material or product keywords. Candidate factors are retrieved using text similarity and sectoral context. Candidates are then filtered by seven tests, each of which a record must pass before it can serve as a donor for an HS6 code.

The seven tests, stated in full in Supporting Information Section S11: unit compatibility; the waste boundary; reference-flow substance; material composition (a fifty-class taxonomy over the HS description and the donor’s reference product); product hierarchy; processing stage (an ordinal ladder read from the nomenclature itself, so that refined silver at

459 kgCO<sub>2</sub>e/kg cannot price a precious-metal ore); and traceability (a code with no admissible donor falls to the stage-matched fallback rather than take a wrong factor).

The division of labour is exact: the model assists candidate classification and disambiguation, and nothing else; every admissibility test, like the repair, is a deterministic rule over recorded metadata, and all carbon intensities come from documented life-cycle sources.

Restricting the model to arbitration among documented candidates, never to producing a number, keeps the provenance of every factor auditable [Preuss et al., 2024, Zhao et al., 2025].

## Blind expert audit of the matching

Those seven tests are the outcome of a two-iteration repair that is itself a result of the paper: the pipeline that first produced this account applied five of the seven not at all, and the before-and-after pair of blind audits is a statistic the literature has not previously reported.

Accuracy was measured on a stratified sample of 100 HS6 codes from the direct-physical layer (the 30 highest-emission codes plus 70 drawn proportionally to emissions), coded blind on a workbook carrying no model verdict and no pre-filled value, each match graded correct, acceptable or wrong (protocol in Supporting Information Section S12).

The expert found 17 of 100 assignments wrong, 14 acceptable and 69 correct. A model pre-adjudication of the same sample agreed with the expert column at chance level (exact agreement 39/100, Cohen’s  $\kappa = -0.003$ ) and had validated fifteen of the seventeen wrong assignments, among them the diborane factor on apples; we therefore withdrew all model-assisted agreement statistics, and the blind human audit is the only reliability evidence this paper treats as such (Supporting Information Section S12).

Inspection of the seventeen identified four failure modes, each attributable to a pipeline stage: ingestion (a column-header block parsed as data, setting the factor of 719 codes); semantically unrelated donors (fresh apples priced by *market for diborane*, uncooked pasta by *printed wiring board production* at 278 kgCO<sub>2</sub>e/kg); components standing for assemblies; and waste-treatment datasets standing for virgin production (Supporting Information Section S12).

The cause is mechanical: the text the similarity score read was not the process name but a record string of numeric codes and hashes, so the activity name was never compared, and the arbitration that followed tested neither hierarchy nor material composition (Supporting Information Section S12).

## Scope of the defect and its repair

Applied as detection rules over the whole factor set, the tests price the defect: 42,864 of 68,102 donor records (62.9%) fail at least one, the codes holding an inadmissible donor carry 73.2% of 2023 embodied emissions, and 1,645 of 4,667 codes retain no admissible donor and fall to the stage-matched fallback rather than carry a wrong factor. The stage test, added in a second iteration, was validated on the three cases whose direction is known a priori: precious-metal ores lose their refined-silver donors (median 171.6 to 1.02 kgCO<sub>2</sub>e/kg), crude petroleum its refinery-gate donors, unwrought aluminium its rolled-product donors (Supporting Information Section S12).

Re-accumulating all 258.6 million flows on admissible donors raises the 2023 world total from 10,282 to 14,368 MtCO<sub>2</sub>e (+39.7%): the excluded waste-treatment datasets had a median intensity of 0.023 kgCO<sub>2</sub>e/kg against 1.53 for the rest, and were pulling every median they touched toward zero. Coverage falls as wrong matches are dropped rather than replaced, and the treatment of the 1,645 orphaned codes brackets the total: plain family inheritance 14,291, stage-matched inheritance (central) 14,368, monetary 13,596, uncovered 11,463 with 33.0% of trade value then carrying no admissible physical evidence, the honest coverage statement that accompanies the central figure. On the central treatment the direct layer carries 71.7% of emissions and 57.7% of trade value, the stage-matched fallback 20.1% and 32.6% (Table 2; full decomposition in Supporting Information Section S12).

## The post-repair audit

The repaired assignment was then audited the same way the defective one was. A new stratified sample of 100 codes was drawn from the repaired direct layer (top 30 plus 70 emissions-weighted draws, 50.0% of the direct-layer total). The author coded it blind again, under the protocol stated above, on a workbook that lists *every* retained donor with its reference product, unit and processing stage, with no model verdict anywhere. The verdict: **8 wrong in 100**, against 17 in 100 before the repair (2 of 30 in the certainty stratum, 6 of 70 in the emissions-weighted stratum, none acceptable). The eight residual errors are all of one kind: semantic mismatches *within* the correct material class and processing stage. That is precisely the failure mode a lexical taxonomy and a structural stage ladder cannot reach. The eight carry 1.22% of the repaired account's emissions and 1.45% of its trade value. We do not hand-correct them: correcting the codes the sample happened to expose, and not the codes it did not, would bias the very error rate the audit measures. The pair 17/100 → 8/100, each measured blind on a stratified emissions-weighted sample, is the paper's measurement-quality result. Every number in this paper is computed on the audited, repaired assignment.

A factor-level plausibility diagnostic accompanies the filter as a flag, never a truncation;

its main catch is 38 codes whose assigned factor is exactly zero, an open class error inside the waste exception that neither audit sample happened to draw. Four limits bound what the audit pair establishes, chiefly that the samples cover only the direct-physical layer, so the fallback layers (27.7% of emissions) are unaudited inheritance (Supporting Information Section S12).

### 3.5 Assignment layers

Each flow is assigned in the strict hierarchy of Section 2.2: HS6 direct physical match on admissible donors, then the stage-matched fallback for codes whose donors the filter rejected, then HS4 and HS2 physical fallbacks and a global physical median, then a monetary fallback (kgCO<sub>2</sub>e per USD applied to trade value) for the remainder. No flow with positive trade value is dropped, and every flow keeps its layer label, so results can be stratified by evidence quality (Table 2, Supporting Information Figure S3).

### 3.6 Quantities and units

The central estimates are products of physical quantities and per-kilogram factors, independent of currency conversion; monetary conversion enters only the services module and the EXIOBASE comparison, which do not share a rate (Supporting Information Section S11).

### 3.7 Services (supplementary extension)

Services are not observed at product resolution; the Supporting Information develops a monetary-intensity services module and shows it does not improve the aggregate reconstruction, so all central results are goods-only.

### 3.8 Benchmark data and residuals

We compare reconstructed trade adjustments with benchmark production- and consumption-based emissions from the Our World in Data CO<sub>2</sub> and greenhouse gas emissions dataset [Ritchie et al., 2024]. For country  $c$  and year  $t$ , the benchmark wedge is

$$\text{Wedge}_{c,t}^{\text{bench}} = \text{CBE}_{c,t}^{\text{bench}} - \text{PBE}_{c,t}^{\text{bench}}. \quad (9)$$

The residual is

$$\text{Residual}_{c,t} = \text{Wedge}_{c,t}^{\text{bench}} - \text{Adj}_{c,t}^{\text{recon}}. \quad (10)$$

We interpret this residual as diagnostic. It reflects the difference between observed product-level trade exposure and benchmark consumption-accounting estimates, which can arise from services, re-exports, indirect upstream requirements, domestic allocation choices and system-boundary differences.

## 4 Validation and uncertainty

### 4.1 Coverage and layer robustness

We evaluate coverage separately for trade value and embodied emissions, because products with small value shares can carry large emissions shares (Supporting Information Figure S3). In the latest goods account, direct HS6 physical matches on audited, admissible donors cover 57.7% of trade value and 71.7% of embodied emissions; the stage-matched fallback carries 32.6% and 20.1%, the plain HS4/HS2/global physical fallbacks 8.5% and 7.7%, and the monetary fallback only 1.2% and 0.6%; this structure is stable across the panel (Supporting Information Figure S4). Before the repair the direct layer read 90.2% of value and 94.9% of emissions; that figure measured the reach of an unaudited filter, and the physical layers together still carry 98.8% of the account.

### 4.2 Emission-factor uncertainty

For products with uncertainty ranges or multiple sources, we compare lower, central and upper factors and test whether the product rankings, country rankings and aggregate emissions hold. A Monte Carlo analysis (4,000 draws) propagates emission-factor uncertainty through the full pipeline. It perturbs each product factor with a mean-preserving log-normal shock calibrated to its spread across the admissible donor records; codes on the stage-matched fallback carry the spread of the donor pool they inherit from. The spreads are therefore empirical, read from the donor pool each code actually draws on, rather than taken from the pedigree matrix that supplies default uncertainty factors in life-cycle practice [Ciroth et al., 2013]. The empirical route ties the uncertainty to the evidence this account holds for each code; the pedigree route would carry database-average judgements that our admissibility filter has already overridden in places. For world goods in 2023 the deterministic central estimate is 14,367.8 MtCO<sub>2e</sub> (after the factor corrections below; 14,648.3 as assigned). On the audited upstream-only factor set, the Monte Carlo distribution has a median of 14,108.5 MtCO<sub>2e</sub> and a 90% interval of [13,257.0, 15,127.9] MtCO<sub>2e</sub>, a coefficient of variation of 4.0%. The deterministic central estimate sits inside this interval. The correction shifts the central by 1.9%, well within the band.

Assignment uncertainty dominates factor uncertainty: the four orphan treatments span 2,905 MtCO<sub>2e</sub> against 1,871 for the Monte Carlo interval (Section S13).

The interval is narrower than the pre-repair one (coefficient of variation 4.0% against 7.8%): inadmissible donors had inflated the spreads that calibrate each factor's shock (Supporting Information Section S13).

Its width rests on the independence of the draws, not on the factor spreads [Groen and Heijungs, 2017, Schulte et al., 2024]. With the marginals held fixed, dependence widens the interval 1.20 times within donor processes, 2.34 times within HS2 chapters and 4.4 times

under a background component carrying a fifth of each factor’s log variance. Correlation cannot be calibrated from a characterised-score extract, so these are scenarios and the 4.0% figure is conditional on independence; rankings are stable across draws (Supporting Information Section S13).

### 4.3 Top-product audit

We audit the largest contributors to embodied emissions, because a few products dominate the total. The twenty largest HS6 products account for 27.9% of all embodied emissions; fourteen of the twenty are direct HS6 physical matches (crude petroleum, natural gas liquefied and gaseous, urea, soya beans, unwrought aluminium, maize, coal, palm oil, frozen beef) and four ride on the stage-matched fallback (refrigerators, lamps and light fittings among them), which is where the audited filter’s coverage cost is most visible. Twenty codes out of about 4,700 carrying a quarter of the total is concentration enough to make a top-product audit worthwhile, but not enough for a small set to stand in for the whole. The concentration is lower than on a combustion-inclusive basis, where twenty products carry half the account, because charging traded fuels with their eventual combustion loads the total onto a handful of petroleum lines.

As an external check, the factors for the carbon-dominant traded materials sit inside published cradle-to-gate ranges (the matching repair moved unwrought aluminium and semi-finished steel into range), and the audited corrections anchor four factors the extract misplaced; the correction cuts the total by 1.9% and the fertilizer share from 6.5% to 4.5%. Against the EU CBAM total default values the median ratio across the ten regulated groups is 1.00 (Supporting Information Section S3 and Section S13).

### 4.4 Benchmark comparison

We validate the account against benchmark consumption-minus-production wedges on aggregate scale, sign agreement, country-year correlations and residual diagnostics; exact equality is not expected across different system boundaries. Over the full country-year panel the goods-only account captures 96% of the benchmark wedge in aggregate magnitude (ratio 0.956). The scale-robust (log-signed) cross-country correlation is 0.53, the Spearman rank correlation 0.50, the panel-level Pearson correlation 0.69, and sign agreement holds in 69.1% of observations. Among the thirty largest non-entrepôt positions in the latest year, the absolute magnitude ratio is 0.888, the Spearman correlation 0.66 and sign agreement 87% (Table 3).

Two repairs separate these numbers from an earlier version of this account, and each is identified against the benchmark in its own way: the re-characterization moved levels but not ranks (a near-monotone rescaling cannot reorder countries), while the matching repair moved both, taking the magnitude ratio from 0.767 to 0.956 and the panel Pearson

correlation from 0.385 to 0.69. A repair motivated entirely by internal evidence moves the account most of the remaining way toward an external benchmark it was never fitted to, which we read as confirmation that the audit measured real error (Supporting Information Section S13).

Adding the supplementary services extension barely moves these figures (Supporting Information Section S6). Across the 28 available years the log-linear correlation between the sector-level estimates and EXIOBASE is stable at  $r = 0.76$ – $0.79$ ; Supporting Information Section S13 traces the ratio’s history through the combustion-inclusive basis, where two errors of opposite sign once cancelled to a spurious near-unity ratio.

The appropriate level benchmark is the EXIOBASE footprint intensity (the total-requirements multiplier), which like our cradle-to-gate factors embodies the whole upstream chain. Against it, 82.8% of estimated embodied emissions fall within a factor of three and 65.1% within a factor of two; the emissions-weighted median capture ratio is 0.90 and the unweighted median 0.98 (Figure 3). Three mechanisms in this gap have been separated on the same cells: the characterization switch (largest where methane matters, the signature a horizon correction must produce), matching error (the larger of the two, raising the emissions-weighted median from 0.67 to 0.90), and the boundary residue that remains, undershoot in machinery and fuels against overshoot in agriculture and wood. What an earlier version reported as process-LCA truncation was in large part mis-assigned donors, and the account is approximately unbiased at the median rather than a conservative lower bound (Supporting Information Section S13).

As an independent check, the eleven material factors matched to the ICE database (compiled outside the ecoinvent–Agribalyse backbone, on the same cradle-to-gate boundary) agree at a median ratio of 1.04 with a log–log correlation of 0.91; nine of eleven fall within a factor of two, and one of the two outside sits on an inherited fallback factor (Supporting Information Section S4 and Section S13).

One structural feature deserves its own record: before the repair the largest position, China, took the opposite sign to the benchmark and held the latest-year Pearson correlation near zero; on the audited account China’s adjustment is  $-812$  MtCO<sub>2</sub>e against a wedge of  $-1,315$ , the sign agrees, and the cross-section correlation rises to 0.69. India’s divergence survives, and among the thirty largest non-entrepôt positions sign agreement is 87% (Supporting Information Section S13).

A truncation test separates a real divergence from a capture artefact (Table 1). India’s positive adjustment survives all three treatments, so it is a genuine composition-versus-technology divergence; China’s changes sign, so China’s sign is not identified.




Figure 3: **Validation against the EXIOBASE footprint intensity.** (a) Median ratio of the Trade-LCA intensity to the EXIOBASE footprint (total-requirements, **fL**) intensity, by broad sector. Material-intensive sectors that dominate embodied trade carbon (blue) agree closely; complex manufactures (red) are under-captured. On the repaired account the residual gap is no longer read as process-LCA boundary truncation. Most of what an earlier version of this paper attributed to truncation was mis-assigned donors: repairing the matching raised the emissions-weighted median capture from 0.67 to 0.90, so matching error was the larger of the two effects and the account is approximately unbiased at the median rather than a conservative lower bound (Section 4.4). (b) Cumulative share of estimated embodied carbon by the Trade-LCA/footprint ratio (emissions-weighted, log scale): 65.1% of embodied carbon falls within a factor of two of the footprint intensity (shaded band) and 82.8% within a factor of three.

Table 1: Truncation-asymmetry test on the two largest residual divergences. Trade-LCA adjustment under three treatments of differential sector capture, MtCO<sub>2</sub>e, latest year.

| Treatment of differential capture            | India  | China  |
|----------------------------------------------|--------|--------|
| No correction                                | +346   | −763   |
| Sectors both accounts capture at $\geq 0.60$ | +137   | −1,400 |
| First-order de-truncation, both legs         | +998   | +751   |
| Benchmark wedge (CBE − PBE)                  | −519.5 | −1,315 |

A country that imports from well-captured sectors while exporting from poorly captured ones would show a positive adjustment for measurement reasons alone. India’s adjustment keeps its positive sign under all three treatments, against a negative benchmark wedge, so the divergence is not that artefact. China’s adjustment changes sign between treatments, so its sign is not identified. The de-truncation column has known limits: dividing the fuel leg by its capture ratio of 0.24 partly re-imports the combustion emissions the upstream basis excludes, and 44.7% of the de-truncated total still sits on sectors captured below 0.30.

## 4.5 Residual diagnostics

We flag countries where residuals may be affected by re-exports, entrepôt trade or services. These include Singapore, Hong Kong, the Netherlands, Belgium and the United Arab Emirates. For these economies, the residual can reflect the difference between observed customs flows and the allocation of indirect requirements in input-output accounts.

## 5 Results

### 5.1 Coverage by evidence layer

The central account is not driven by residual assumptions: direct HS6 physical matches on admissible donors carry 71.7% of embodied emissions and 57.7% of trade value, physical evidence including the stage-matched and family fallbacks carries 98.8% of both, and the monetary fallback only 0.6% of emissions (Table 2; layer detail in Supporting Information Section S14).

### 5.2 Concentration of embodied carbon

Embodied carbon in traded goods is highly concentrated (Figure 6). Across the reconstructed goods account, gross embodied trade (the sum of all bilateral flows) amounts to 14,367.8 MtCO<sub>2</sub>e in the latest year (Section 4.2).

The gross figure double-counts the shared upstream of multi-stage trade and is an exposure measure (Section 2.3); concentration is reported on it because border measures act on gross flows. Mineral fuels account for 2,184.1 MtCO<sub>2</sub>e, or 15.2% of embodied emissions, while representing 14.0% of trade value. Iron and steel contribute 847.2 MtCO<sub>2</sub>e, or 5.9% of emissions, on 2.1% of trade value. Fertilizers account for 644.5 MtCO<sub>2</sub>e, or 4.5% of emissions, but only 0.5% of trade value, followed by plastics, organic chemicals, fats and oils, aluminium and meat.

Machinery and electrical machinery owe their positions mainly to the stage-matched fallback, and the fertilizer level reflects the audited nitrogen cap; both mechanics are documented in Supporting Information Section S14. Trade value is therefore a poor proxy for embodied carbon exposure. The heavy chapters are the material-production sectors that have become a large and growing share of global emissions [Hertwich, 2021].

### 5.3 Within-sector dispersion of carbon intensity

The central advantage of product resolution is that carbon intensity varies by an order of magnitude *within* the sectors at which input-output models operate. The audit changes what can honestly be claimed about how much. We measure the dispersion on two bases

that answer different questions. The *physical* basis (kgCO<sub>2</sub>e/kg) isolates carbon from price. On the audited factor set, the trade-weighted ratio between the 90th and 10th percentile product intensities in the median HS2 sector is about thirteen-fold (precisely 12.7, across the 71 chapters with at least five directly matched codes). The *monetary* basis (kgCO<sub>2</sub>e/USD) adds price variation and reaches 12.2-fold. Before the matching repair the same statistics read 35.0 and 82.7. Roughly two thirds of the dispersion an unaudited product-level account measures, and an earlier version of this paper reported, is therefore matching error rather than product heterogeneity: wrong donors are drawn from the far tails of the factor distribution. The audited thirteen-fold figure is the one we defend, and it is still far beyond anything a sector average can carry.

Monetary intensity compounds carbon with price: unit values span 4.4-fold in the median chapter, but carbon intensity and unit value are positively correlated within sectors, so the covariance cancels most of the price channel and the two bases nearly coincide at the median (the full decomposition, and the pre-repair artefact it corrects, are in Supporting Information Section S14).

We report the physical figure as the headline measure of carbon heterogeneity, and the monetary figure as the composite statistic that a value-based instrument would in fact inherit, because a value-based MRIO applies one sector-average kgCO<sub>2</sub>e/USD to every dollar of trade in the sector.

A sector-average intensity mis-estimates the embodied carbon of a large share of trade even on the audited factors (Figure 4). On the monetary basis, the basis on which an MRIO average operates, products that differ from their sector average by more than a factor of two account for 45.5% of traded embodied emissions and 59.9% of trade value (beyond a factor of five: 14.7% and 29.6%). Price does not drive the mis-estimated share: on the pure physical basis, these products still carry 34.2% of embodied emissions. The result also survives outlier cleaning. Capping the 0.7% of products still sitting above thirty times their sector median changes total emissions by 1.9% and leaves these shares essentially unchanged.

The mis-estimates concentrate exactly where border measures and procurement standards must discriminate: ores, base metals, chemicals and machinery, where products like primary aluminium and petroleum coke sit far above their sector mean (Supporting Information Section S14).

Finer sectors do not dissolve the heterogeneity: at HS4 resolution (341 groups, finer than any operational MRIO) a group average still mis-estimates 8.7% of emissions and 14.6% of value beyond a factor of two, so the mis-estimated share is bounded between 9–46% of emissions and 15–60% of value at any feasible MRIO granularity, and the shares are stable across 1995–2023 (Supporting Information Section S14).




**Figure 4: Carbon intensity varies by orders of magnitude within input-output sectors.** (a) Trade-weighted p10–p90 spread of product carbon intensities, in kgCO<sub>2</sub>e per kg, within the top-emitting HS2 sectors (log scale); the dot is the sector median. The median chapter spans a factor of 12.7 ( $n = 71$  chapters). (b) The same spread computed per dollar of trade, in kgCO<sub>2</sub>e per USD, which is the unit a value-based MRIO applies. The median chapter spans a factor of 12.2 ( $n = 96$ ), slightly *less* than on the physical basis: carbon intensity and unit value are positively correlated within sectors, so the price channel cancels most of its own contribution instead of adding to it (Section 5.3). (c) Share of a total carried by products that deviate from their own HS2 sector average by at least the factor on the horizontal axis. Deviation is symmetric, the larger of the ratio and its inverse, so overstatement and understatement count alike. At a factor of two (left dotted line) these products carry 45.5% of traded embodied emissions and 59.9% of trade value on the monetary basis, and 34.2% of emissions on the physical basis (dashed). A sector-average intensity, the resolution of MRIO models, cannot capture this heterogeneity.

## 5.4 Country-level Trade-LCA adjustments and the aggregate reconstruction ratio

Across all country-years the goods-only reconstruction captures 119,642.6 MtCO<sub>2</sub>e of the 125,179.1 benchmark wedge magnitude (ratio 0.956), with sign agreement in 69.1% of country-years, Spearman rank correlation 0.50 and Pearson 0.69. Supporting Information Figure S7 shows the annual trend. Supporting Information Section S13 reports the detail.

## 5.5 Residuals identify where product trade and MRIO accounting diverge

For several economies the goods reconstruction captures most of the benchmark wedge (the United States 448.7 against 513.3 MtCO<sub>2</sub>e, Russia −344.4 against −327.2); the mid-sized residuals are read in Supporting Information Section S14.

China is where the repair changed the account’s reading most: from a net importer at +266 MtCO<sub>2</sub>e (sign contradicting the benchmark wedge of −1,315) to a net exporter at −812, the defective donors having priced its dominant export chapters near zero (Supporting Information Section S14).

The most informative remaining residuals are India (+349 against a wedge of −520) and Vietnam: consumption accounts make them net exporters of embodied carbon through dirtier-than-average production, while the representative-technology account values their carbon-intensive imports at generic intensity, and for India the divergence survives every truncation treatment (Section 4.4); the gap pinpoints the economies for which production technology, the dimension a border measure acts on, matters most. These residuals are diagnostic outputs (Figure 5). They identify where observed product-level trade exposure differs from benchmark consumption-accounting conventions through composition, services, re-exports, indirect supply chains and system-boundary choices.

# 6 Discussion

## 6.1 Contribution to the literature

Product-level Trade-LCA provides a reproducible bridge between life-cycle evidence and global trade statistics. It adds product specificity to embodied carbon trade accounting, and it keeps a transparent diagnostic of the evidence layer used in each assignment. The contribution complements MRIO-based carbon footprints. Those provide economy-wide consistency and final-demand allocation, but they typically operate at more aggregated sectoral resolution [Tukker and Dietzenbacher, 2013, Stadler et al., 2018].

The paper advances three steps beyond the product-level accounts of Sato [2014a] and Kanemoto et al. [2016] and the aggregation-bias quantification of Su et al. [2010], which

worked within input–output resolution rather than below it: a documented process-LCA backbone mapped to HS6 over a 29-year panel with a measured error rate (17 then 8 wrong in 100, the first published for an account of this kind); a quantitative policy statistic, audited against the account’s own measurement error, for how much embodied carbon a sector average mis-estimates; and a composition-versus-technology reading of the sign disagreements with consumption accounts (Supporting Information Section S15).

The audit pair is also the paper’s contribution to the literature on language models in life-cycle assessment [Balaji et al., 2023a,b, 2025, Tu et al., 2024]: the first model-assisted customs-to-LCA pipeline with measured reliability, whose cautionary half is that the model-based pre-adjudication agreed with blind expert coding at chance level ( $\kappa = -0.003$ ) while everything that fixed the pipeline was a deterministic rule. The exportable lesson is one sentence: model assistance in classification is useful, model self-validation is worthless, and a blind human audit is not optional.

A monetary-intensity services extension is reported in the Supplementary Material; because services lack product resolution and the extension does not improve the reconstruction, the product-level goods account remains the core of the paper.

## 6.2 Implications for product-level climate policy

One scope statement: these are default values in the border-adjustment sense, a contested basis for a charge, and we take no position in that debate (Section S15). The factors are an accounting and screening instrument, not a liability basis. Measurement error in a policy’s pollution basis redistributes burden across heterogeneous firms [Lyubich et al., 2018], which is why we publish an error rate.

With that stated, the framework supports policy applications that require product-level carbon information, and the cost of *not* having it is concrete. Consider a screen that prices or flags traded goods by their carbon intensity, the logic of a border measure or a procurement standard. A screen built on sector-average intensities, the resolution an MRIO provides, would *under-price by more than half* a set of 640 products. Together they carry 5,412 MtCO<sub>2e</sub>, or 37.7% of all embodied trade carbon, because each is far more carbon-intensive than its sector. The list is specific and policy-relevant: unwrought aluminium alloys (3.1 times their sector average), photosensitive and solar apparatus (5.3 times, inside the otherwise moderate electrical-machinery sector), nickel ores (about 14 times), lignite (13 times), cement clinker, methanol and semi-finished steel. A product-level account flags exactly these goods; a sectoral account cannot. The framework thus identifies where monetary trade value, or a sectoral average, understates physical carbon exposure [European Parliament and Council, 2023, OECD, 2023].

Because the factors are technology-normalized, a country’s embodied-trade position also decomposes into a composition and a production-technology component, with compo-

sition the larger on average; border instruments act on the technology component only (Supporting Information Section S15).

### 6.3 Relationship to MRIO approaches

MRIO models allocate emissions to final demand; Trade-LCA measures the carbon embodied in observed product flows. They answer different questions and should be used together, which is where the field’s own consensus assessment points, ranking environmental data quality ahead of economic resolution [Tukker et al., 2018]. Our sector-level estimates correlate with EXIOBASE at  $r = 0.76$ – $0.79$  in all 28 years (Section S15).

MRIO consumption-based estimates carry documented uncertainty of their own [Rodrigues et al., 2018, Moran and Wood, 2014, Wiedmann and Lenzen, 2018], and the geographic correction of Davis et al. [2025] is orthogonal to the within-sector dispersion measured here, so the two corrections compound rather than compete (Supporting Information Section S15).

### 6.4 Limitations

The framework has several limitations, and the first is the largest: the matching is error-prone at a measured rate (17 then 8 wrong in 100), the residual errors are the kind neither a lexical taxonomy nor a stage ladder can catch, the audit covers only the direct-physical layer, and 33.0% of trade value carries no admissible direct physical evidence. The 38 zero-factor codes exposed by the plausibility diagnostic remain open. Nothing else in this list is of that order. Three further limits: factors are representative-technology intensities held fixed over 1995–2023, so the annual series is a lower bound on decarbonization while the within-year dispersion results are unaffected; HS categories can hold heterogeneous products; and entrepôt trade distorts country-level readings. The rest are in the Supporting Information.

The remaining limitations are listed in Supporting Information Section S15: factors are representative-technology averages held fixed over time, so the annual series is a lower bound on decarbonization; HS categories contain heterogeneous products; BACI quantities carry reporting error; re-exports can distort country-level readings; and the account does not propagate emissions through complete value chains or allocate domestic production.

### 6.5 Future work

Future work should develop country- and technology-specific factors, integrate the account with MRIO models, and bring services onto the same product-resolution footing as goods.

## 7 Conclusion

Product-level Trade-LCA provides a reproducible bridge between life-cycle evidence and global trade statistics. It complements production- and consumption-based accounts. It identifies the traded products and country positions through which embodied carbon exposure is created.

The results show that trade value is a poor proxy for embodied-carbon exposure, that carbon intensity varies by an order of magnitude within the sectors input-output models resolve, and, by auditing this account's own matching before and after repair, that unaudited product-level accounts overstate that heterogeneity: the audit is as much a part of the method as the mapping. By making the evidence layer of each assignment explicit, the framework provides a transparent basis for product-level carbon accounting in global trade.

Country decomposition: OWID wedge vs goods net imports (2023, upstream-only basis)




Figure 5: **Country-level decomposition.** Top countries by absolute OWID wedge in the latest year. Bars show the OWID CBE-PBE wedge and goods net imports of embodied carbon.




Figure 6: **Product-group concentration.** Left bars show trade-value shares; right bars show emissions shares for the top HS2 sectors. Fertilizers are the clearest case of emissions concentration without value concentration: 0.5% of trade value against 4.5% of emissions. Mineral fuels are not, on this boundary. HS27 carries 14.0% of value against 15.2% of emissions, so it is close to proportional, because the upstream-only treatment excludes the combustion of imported fuels and combustion is what made fuels look disproportionate in territorial and consumption accounts. The concentration that survives is therefore a statement about production intensity, not about the energy content of the goods.

Table 2: Goods coverage by evidence layer, latest year (2023), on the audited basis: iteration-2 admissible donors, stage-matched fallback for the codes whose donors the filter rejected, upstream-only mineral-fuel factors and the factor corrections described in Section 4.3. The emissions column sums to 14,367.8 MtCO<sub>2</sub>e, the central estimate. Before the matching repair the direct layer carried 90.2% of value and 94.9% of emissions; the difference is the coverage cost of dropping inadmissible donors, and a factor change moves flows across the physical-validity cap, which is why the account is re-accumulated at flow level rather than rescaled.

| Layer                                          | Trade value (trn USD) | Share of trade value | Emissions (MtCO <sub>2</sub> e) | Share of emissions | Products | Flows   |
|------------------------------------------------|-----------------------|----------------------|---------------------------------|--------------------|----------|---------|
| HS6 direct physical (audited donors)           | 13.04                 | 57.7%                | 10298.0                         | 71.7%              | 2797     | 5538317 |
| Stage-matched physical fallback                | 7.38                  | 32.6%                | 2884.9                          | 20.1%              | 1505     | 3810670 |
| HS4 physical fallback                          | 1.25                  | 5.5%                 | 896.5                           | 6.2%               | 179      | 550919  |
| HS2 physical fallback                          | 0.67                  | 3.0%                 | 202.4                           | 1.4%               | 155      | 452262  |
| Global physical fallback                       | 0.01                  | 0.04%                | 1.1                             | 0.01%              | 5        | 16072   |
| Monetary fallback (capped or missing quantity) | 0.27                  | 1.2%                 | 84.9                            | 0.6%               | 4561     | 468035  |

Table 3: Validation against OWID CBE–PBE benchmark, on the audited upstream-only, factor-corrected, IPCC 2013 GWP 100a basis. Before the matching repair the full-panel row read Pearson 0.39, Spearman 0.50, log-signed 0.50, same-sign 70%, absolute ratio 0.767, and the latest-year Pearson was 0.05; on a combustion-inclusive basis the full-panel ratio of 1.007 reflected two offsetting errors rather than agreement (Section 4.4).

| Sample                      | $n$  | Pearson $r$ | Spearman $\rho$ | Log-signed $r$ | Same-sign | Abs. ratio |
|-----------------------------|------|-------------|-----------------|----------------|-----------|------------|
| Full panel (1995–2023)      | 3421 | 0.69        | 0.50            | 0.53           | 69%       | 0.956      |
| Latest year (2023)          | 119  | 0.69        | 0.52            | 0.54           | 77%       | 1.049      |
| Top 30, latest, ex-entrepôt | 30   | 0.73        | 0.66            | 0.70           | 87%       | 0.888      |

## Declarations

## Funding

This work was supported by a CIFRE agreement (Association Nationale de la Recherche et de la Technologie, ANRT, grant no. 2022/1627) between Toovalu and the Université de Rennes (CREM UMR 6211).

## Competing interests

The author declares no competing interests. Funding, including the CIFRE agreement under which the author was employed by Toovalu, is stated in the funding declaration.

## Author contributions

F.P. is the sole author and conceived the study, developed the product-level Trade-LCA framework, designed the empirical strategy, performed the analysis and wrote the manuscript.

## Data and code availability

The code, derived data, blind-audit workbooks and audit trail that reproduce every number, figure and table in this article are deposited on Zenodo under the concept DOI [10.5281/zenodo.20719018](https://doi.org/10.5281/zenodo.20719018) (data and documents CC BY 4.0, code MIT). The licensed raw inputs (ecoinvent, BACI) are not redistributed and are available from their providers; Agribalyse, EXIOBASE, Our World in Data, World Bank, ICE and the CBAM default values are public.

## Acknowledgements

This work was supported by a CIFRE doctoral fellowship (Association Nationale de la Recherche et de la Technologie, grant n° 2022/1627) in partnership with Toovalu, which provided access to the ecoinvent database. The study also draws on publicly available data from CEPII (BACI), the ADEME Base Empreinte (Agribalyse), EXIOBASE, Our World in Data and the World Bank, whose providers the author gratefully acknowledges.

## Ethics approval

Not applicable. The study analyses public and licensed statistical and life-cycle databases and involves no human participants, human data, human material or animals.

## Consent to participate

Not applicable.

## Consent to publish

Not applicable.

## Supporting information

Supporting Information is available for this article. Sections S1 to S9 carry the data provenance, the residual interpretation, the CBAM and ICE comparisons, the dispersion robustness, the services extension and the country-level decomposition. Sections S10 to S15 carry the supplementary notes, which give the full detail of arguments the main text states in compressed form: S10 the framework, the account boundary and the prior work; S11 the data pipeline, the emission factors and the admissibility rules; S12 the blind audit, the defect census and the repair; S13 uncertainty, benchmark comparison and independent validation; S14 concentration, within-sector dispersion and country patterns; S15 the discussion, the limitations and the future work. Section S16 collects the supplementary figures.

## References

- Bharathan Balaji, Venkata Sai Gargeya Vunnava, Nina Domingo, Shikhar Gupta, Harsh Gupta, Geoffrey Guest, and Aravind Srinivasan. Flamingo: Environmental impact factor matching for life cycle assessment with zero-shot machine learning. *ACM Journal on Computing and Sustainable Societies*, 1(2):1–23, December 2023a. ISSN 2834-5533. doi: 10.1145/3616385. URL <http://dx.doi.org/10.1145/3616385>.
- Bharathan Balaji, Venkata Sai Gargeya Vunnava, Geoffrey Guest, and Jared Kramer. Caml: Carbon footprinting of household products with zero-shot semantic text similarity. In *Proceedings of the ACM Web Conference 2023*, WWW '23, page 4004–4014. ACM, April 2023b. doi: 10.1145/3543507.3583882. URL <http://dx.doi.org/10.1145/3543507.3583882>.
- Bharathan Balaji, Fahimeh Ebrahimi, Nina Gabrielle G Domingo, Venkata Sai Gargeya Vunnava, Abu-Zaher Faridee, Soma Ramalingam, Shikha Gupta, Anran Wang, Harsh Gupta, Domenic Belcastro, Kellen Axten, Jeremie Hakian, Jared Kramer, Aravind Srinivasan, and Qingshi Tu. Emission factor recommendation for life cycle assessments with generative ai. *Environmental Science & Technology*, 59(18):9113–9122, March 2025. ISSN 1520-5851. doi: 10.1021/acs.est.4c12667. URL <http://dx.doi.org/10.1021/acs.est.4c12667>.
- Christoph Böhringer, Carolyn Fischer, Knut Einar Rosendahl, and Thomas F. Rutherford. Potential impacts and challenges of border carbon adjustments. *Nature Climate Change*, 12(1):22–29, 2022. doi: 10.1038/s41558-021-01250-z. URL <https://doi.org/10.1038/s41558-021-01250-z>.
- Andreas Ciroth, Stéphanie Muller, Bo Weidema, and Pascal Lesage. Empirically based uncertainty factors for the pedigree matrix in ecoinvent. *The International Journal of Life Cycle Assessment*, 21(9):1338–1348, December 2013. ISSN 1614-7502. doi: 10.1007/s11367-013-0670-5. URL <http://dx.doi.org/10.1007/s11367-013-0670-5>.
- Robert H. Crawford. Validation of a hybrid life-cycle inventory analysis method. *Journal of Environmental Management*, 88(3):496–506, 2008. doi: 10.1016/j.jenvman.2007.03.024. URL <https://doi.org/10.1016/j.jenvman.2007.03.024>.
- S. J. Davis and K. Caldeira. Consumption-based accounting of CO<sub>2</sub> emissions. *Proceedings of the National Academy of Sciences*, 107(12):5687–5692, 2010. doi: 10.1073/pnas.0906974107. URL <https://doi.org/10.1073/pnas.0906974107>.
- Steven J. Davis, Andrew Dumit, Mo Li, Yohanna Maldonado, Michael Steffen, Martha Stevenson, Tatiana Boldyreva, and Sangwon Suh. The importance of multiregional

accounting for corporate carbon emissions. *Nature Communications*, 17, 2025. doi: 10.1038/s41467-025-67759-5. URL <https://doi.org/10.1038/s41467-025-67759-5>.

Arjan de Koning, Martin Bruckner, Stephan Lutter, Richard Wood, Konstantin Stadler, and Arnold Tukker. Effect of aggregation and disaggregation on embodied material use of products in input–output analysis. *Ecological Economics*, 116:289–299, August 2015. ISSN 0921-8009. doi: 10.1016/j.ecolecon.2015.05.008. URL <http://dx.doi.org/10.1016/j.ecolecon.2015.05.008>.

European Parliament and Council. Regulation (EU) 2023/956 of the European Parliament and of the Council of 10 May 2023 establishing a carbon border adjustment mechanism. Official Journal of the European Union, L 130/52, 2023. OJ L 130, 16.5.2023, p. 52–104; entry into force 17 August 2023.

Guillaume Gaulier and Soledad Zignago. BACI: International Trade Database at the Product-Level. The 1994–2007 Version. *CEPII Working Paper*, 2010-23, 2010. doi: 10.2139/ssrn.1994500. URL <http://www.cepii.fr/CEPII/en/publications/wp/abstract.asp?NoDoc=2726>.

E.A. Groen and R. Heijungs. Ignoring correlation in uncertainty and sensitivity analysis in life cycle assessment: what is the risk? *Environmental Impact Assessment Review*, 62:98–109, January 2017. ISSN 0195-9255. doi: 10.1016/j.eiar.2016.10.006. URL <http://dx.doi.org/10.1016/j.eiar.2016.10.006>.

Ramana Gudipudi, Luis Costa, Ponraj Arumugam, Matthew Agarwala, Felix Creutzig, and Jürgen P. Kropp. From corporate net-zero pledges to credible climate action. *Nature Sustainability*, 2026. doi: 10.1038/s41893-026-01908-6. URL <https://doi.org/10.1038/s41893-026-01908-6>.

Edgar G. Hertwich. Increased carbon footprint of materials production driven by rise in investments. *Nature Geoscience*, 14:151–155, 2021. doi: 10.1038/s41561-021-00690-8. URL <https://doi.org/10.1038/s41561-021-00690-8>.

Edgar G. Hertwich and Glen P. Peters. Carbon footprint of nations: a global, trade-linked analysis. *Environmental Science & Technology*, 43(16):6414–6420, 2009. doi: 10.1021/es803496a. URL <https://doi.org/10.1021/es803496a>.

Keiichiro Kanemoto, Daniel Moran, and Edgar G. Hertwich. Mapping the carbon footprint of nations. *Environmental Science & Technology*, 50(19):10512–10517, 2016. doi: 10.1021/acs.est.6b03227. URL <https://doi.org/10.1021/acs.est.6b03227>.

Peter Koch and Thibault Salou. AGRIBALYSE: Methodology for the food products, version 3.1. ADEME (French Environment and Energy Management Agency), Angers, France, 2020. Distributed via the ADEME Base Empreinte database.

- Manfred Lenzen. Errors in conventional and input-output-based life-cycle inventories. *Journal of Industrial Ecology*, 4(4):127–148, 2000. doi: 10.1162/10881980052541981. URL <https://doi.org/10.1162/10881980052541981>.
- Manfred Lenzen. Aggregation versus disaggregation in input–output analysis of the environment. *Economic Systems Research*, 23(1):73–89, March 2011. ISSN 1469-5758. doi: 10.1080/09535314.2010.548793. URL <http://dx.doi.org/10.1080/09535314.2010.548793>.
- Manfred Lenzen, Daniel Moran, Keiichiro Kanemoto, and Arne Geschke. Building EORA: a global multi-region input–output database at high country and sector resolution. *Economic Systems Research*, 25(1):20–49, 2013. doi: 10.1080/09535314.2013.769938. URL <https://doi.org/10.1080/09535314.2013.769938>.
- Eva Lyubich, Joseph S. Shapiro, and Reed Walker. Regulating mismeasured pollution: Implications of firm heterogeneity for environmental policy. *AEA Papers and Proceedings*, 108:136–142, May 2018. ISSN 2574-0776. doi: 10.1257/pandp.20181089. URL <http://dx.doi.org/10.1257/pandp.20181089>.
- Jan C. Minx, Thomas Wiedmann, Richard Wood, Glen P. Peters, Manfred Lenzen, Anne Owen, Kate Scott, John Barrett, Klaus Hubacek, Giovanni Baiocchi, Arne Paul, Edmond Dawkins, Jon Briggs, Dabo Guan, Sangwon Suh, and Frank Ackerman. Input-output analysis and carbon footprinting: an overview of applications. *Economic Systems Research*, 21(3):187–216, 2009. doi: 10.1080/09535310903541298. URL <https://doi.org/10.1080/09535310903541298>.
- Daniel Moran and Richard Wood. Convergence between the Eora, WIOD, EXIOBASE, and OpenEU’s consumption-based carbon accounts. *Economic Systems Research*, 26(3): 245–261, 2014. doi: 10.1080/09535314.2014.935298. URL <https://doi.org/10.1080/09535314.2014.935298>.
- OECD. Effective carbon rates 2023: Pricing greenhouse gas emissions through taxes and emissions trading systems. OECD Publishing, Paris, 2023. URL <https://doi.org/10.1787/b84d5b36-en>.
- G. P. Peters, J. C. Minx, C. L. Weber, and O. Edenhofer. Growth in emission transfers via international trade from 1990 to 2008. *Proceedings of the National Academy of Sciences*, 108(21):8903–8908, 2011. doi: 10.1073/pnas.1006388108. URL <https://doi.org/10.1073/pnas.1006388108>.
- Glen P. Peters and Edgar G. Hertwich. CO<sub>2</sub> Embodied in International Trade with Implications for Global Climate Policy. *Environmental Science & Technology*, 42(5):1401–1407, 2008. doi: 10.1021/es072023k. URL <https://doi.org/10.1021/es072023k>.

- Justin R. Pierce and Peter K. Schott. Concoring u.s. harmonized system categories over time. *Finance and Economics Discussion Series*, 2012.0(16):1–13, January 2012. ISSN 1936-2854. doi: 10.17016/feds.2012.16. URL <http://dx.doi.org/10.17016/feds.2012.16>.
- Nathan Preuss, Abdulelah S. Alshehri, and Fengqi You. Large language models for life cycle assessments: Opportunities, challenges, and risks. *Journal of Cleaner Production*, 466:142824, August 2024. ISSN 0959-6526. doi: 10.1016/j.jclepro.2024.142824. URL <http://dx.doi.org/10.1016/j.jclepro.2024.142824>.
- Hannah Ritchie, Max Roser, and Pablo Rosado. Our World in Data CO<sub>2</sub> and greenhouse gas emissions dataset. Our World in Data, 2024. URL <https://ourworldindata.org/co2-and-greenhouse-gas-emissions>. Based on data from Global Carbon Project and other sources; accessed 2024.
- João F. D. Rodrigues, Daniel Moran, Richard Wood, and Paul Behrens. Uncertainty of Consumption-Based Carbon Accounts. *Environmental Science & Technology*, 52: 7577–7586, 2018. doi: 10.1021/acs.est.8b00632. URL <https://doi.org/10.1021/acs.est.8b00632>.
- Misato Sato. Product level embodied carbon flows in bilateral trade. *Ecological Economics*, 105:106–117, September 2014a. ISSN 0921-8009. doi: 10.1016/j.ecolecon.2014.05.006. URL <http://dx.doi.org/10.1016/j.ecolecon.2014.05.006>.
- Misato Sato. Embodied carbon in trade: a survey of the empirical literature. *Journal of Economic Surveys*, 28(5):831–861, 2014b. doi: 10.1111/joes.12027. URL <https://doi.org/10.1111/joes.12027>.
- Simon Schulte, Arthur Jakobs, and Stefan Pauliuk. Estimating the uncertainty of the greenhouse gas emission accounts in global multi-regional input–output analysis. *Earth System Science Data*, 16(6):2669–2700, June 2024. ISSN 1866-3516. doi: 10.5194/essd-16-2669-2024. URL <http://dx.doi.org/10.5194/essd-16-2669-2024>.
- Konstantin Stadler, Richard Wood, Tatyana Bulavskaya, Carl-Johan Södersten, Moana Simas, Sarah Schmidt, Arkaitz Usubiaga, José Acosta-Fernández, Jeroen Kuenen, Martin Bruckner, Stefan Giljum, Stephan Lutter, Stefano Merciai, Jannick H. Schmidt, Michaela C. Theurl, Christoph Plutzar, Thomas Kastner, Nina Eisenmenger, Karl-Heinz Erb, Hein de Vries, and Arnold Tukker. EXIOBASE 3: developing a time series of detailed environmentally extended multi-regional input–output tables. *Journal of Industrial Ecology*, 22(3):502–515, 2018. doi: 10.1111/jiec.12715. URL <https://doi.org/10.1111/jiec.12715>.

- Kjartan Steen-Olsen, Anne Owen, Edgar G. Hertwich, and Manfred Lenzen. Effects of sector aggregation on CO<sub>2</sub> multipliers in multiregional input–output analyses. *Economic Systems Research*, 26(3):284–302, July 2014. ISSN 1469-5758. doi: 10.1080/09535314.2014.934325. URL <http://dx.doi.org/10.1080/09535314.2014.934325>.
- Bin Su, H. C. Huang, B. W. Ang, and P. Zhou. Input–output analysis of CO<sub>2</sub> emissions embodied in trade: the effects of sector aggregation. *Energy Economics*, 32(1):166–175, 2010. doi: 10.1016/j.eneco.2009.07.010. URL <https://doi.org/10.1016/j.eneco.2009.07.010>.
- Sangwon Suh, Manfred Lenzen, Graham J. Treloar, Hiroki Hondo, Arpad Horvath, Gjalt Huppes, Olivier Joliet, Uwe Klann, Wolfram Krewitt, Yuichi Moriguchi, Jesper Munksgaard, and Greg Norris. System boundary selection in life-cycle inventories using hybrid approaches. *Environmental Science & Technology*, 38(3):657–664, 2004. doi: 10.1021/es0263745. URL <https://doi.org/10.1021/es0263745>.
- Qingshi Tu, Jing Guo, Nan Li, Jianchuan Qi, and Ming Xu. Mitigating grand challenges in life cycle inventory modeling through the applications of large language models. *Environmental Science & Technology*, 58(44):19595–19603, October 2024. ISSN 1520-5851. doi: 10.1021/acs.est.4c07634. URL <http://dx.doi.org/10.1021/acs.est.4c07634>.
- Arnold Tukker and Erik Dietzenbacher. Global multiregional input–output frameworks: An introduction and outlook. *Economic Systems Research*, 25(1):1–19, 2013. doi: 10.1080/09535314.2012.761179. URL <https://doi.org/10.1080/09535314.2012.761179>.
- Arnold Tukker, Arjan de Koning, Anne Owen, Stephan Lutter, Martin Bruckner, Stefan Giljum, Konstantin Stadler, Richard Wood, and Rutger Hoekstra. Towards robust, authoritative assessments of environmental impacts embodied in trade: Current state and recommendations. *Journal of Industrial Ecology*, 22(3):585–598, February 2018. ISSN 1530-9290. doi: 10.1111/jiec.12716. URL <http://dx.doi.org/10.1111/jiec.12716>.
- Ilke Van Beveren, Andrew Bernard, and Hylke Vandenbussche. Concording eu trade and production data over time. NBER Working Paper 18604, National Bureau of Economic Research, December 2012. URL <http://dx.doi.org/10.3386/w18604>.
- Gregor Wernet, Christian Bauer, Bernhard Steubing, Jürgen Reinhard, Emilia Moreno-Ruiz, and Bo Weidema. The ecoinvent database version 3 (part I): overview and methodology. *The International Journal of Life Cycle Assessment*, 21(9):1218–1230, 2016. doi: 10.1007/s11367-016-1087-8. URL <https://doi.org/10.1007/s11367-016-1087-8>.
- Thomas Wiedmann and Manfred Lenzen. Environmental and social footprints of international trade. *Nature Geoscience*, 11(5):314–321, 2018. doi: 10.1038/s41561-018-0113-9. URL <https://doi.org/10.1038/s41561-018-0113-9>.

800 Bu Zhao, Jitong Jiang, Ming Xu, and Qingshi Tu. A data-centric investigation on  
801 the challenges of machine learning methods for bridging life cycle inventory data  
802 gaps. *Journal of Industrial Ecology*, 29(3):955–966, April 2025. ISSN 1530-9290. doi:  
803 10.1111/jiec.70022. URL <http://dx.doi.org/10.1111/jiec.70022>.

## Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [supplementary.pdf](#)
- [supplementary.tex](#)