From Shipments to Supply Chains: Mining Input-Output Links from Firm-level Trade Flows*

Karsten Mau
Maastricht University

Antonio Vicencio
Maastricht University

Mingzhi (Jimmy) Xu
Peking University

Yawen Zheng
Shenzhen University

September 2026

Abstract

We develop a method to recover a granular, product-level input-output structure from firm-level customs transactions. Building on an association-rule mining algorithm, we exploit systematic co-occurrences between what firms import and what they export to infer input use in production. Applying the approach to data from several countries, we show that the inferred mappings closely resemble conventional input-output relationships and perform well in external validation exercises. We illustrate the value of these linkages in two applications. First, following China's WTO entry, export growth is accompanied by a pronounced rise in imports of the inputs our mapping associates with those exports. Second, in the 2018-2019 U.S.-China trade war, tariffed Chinese exports to the United States contract, and the shock propagates upstream: China's imports of exposed inputs fall, and third-country suppliers that specialize in those inputs experience the declines in their exports to China, especially if China is a key destination market. Overall, simple pattern recognition tools and micro-level trade data can deliver high-resolution input–output linkages that improve exposure measures and empirical assessments of supply-chain propagation.

Keywords: Input-output linkages, Machine learning, Firm-level data, International trade
JEL codes: C81, L14, F14, F13, O25

*The authors are grateful to the numerous comments, discussions, and suggestions received during seminars at Oxford University, CEPII, OECD, Centraal Planbureau (CPB), Goethe University Frankfurt, the International Input Output Association (IIOA), Maastricht University, Leuven-Louvain International Trade Online Seminar, Peking University, and Statistics Netherlands (CBS), as well as at the 26th Göttingen Workshop in International Economics, the ETSG conference 2025, the 17th Input-Output Workshop. This study uses, in part, the Panjiva dataset (provided by S&P Global) and China Customs data. The conclusions are the authors' own; S&P Global and China Customs had no role in the analysis or writing and bear no responsibility for the results. Mingzhi Xu acknowledges NSFC support (Grant Nos. 72322007, 72503230). Yawen Zheng acknowledges support from Shenzhen University's Research Center for FinTech and Digital-Intelligent Management and the Shenzhen Humanities & Social Sciences Key Research Bases. Karsten Mau acknowledges NWO support (Grant no. 406.XS.25.01.030). Corresponding author: Karsten Mau, Maastricht University, School of Business and Economics, Maastricht, The Netherlands, e-mail: k.mau@maastrichtuniversity.nl; Antonio Vicencio Betanzo, Maastricht University, School of Business and Economics, Maastricht, The Netherlands, e-mail: antonio.vicenciobetanzo@maastrichtuniversity.nl; Mingzhi Xu, Institute of New Structural Economics (INSE), Peking University, Beijing, China, email: mingzhixu@nsd.pku.edu.cn; Yawen Zheng, Shenzhen Audencia Financial Technology Institute (SAFTI), Shenzhen University, Shenzhen, China, email: yawenzheng@szu.edu.cn.

1 Introduction

The prominent role of intermediate input goods in 21st-century international trade flows entails non-trivial interdependencies between countries.[1] This can be observed whenever local shocks to an economy propagate through the network of international trade relations and impact other economies far away from the original disruption. Numerous studies documented this in a range of different settings, which contributed to growing interest in the supply side of international trade and in firms' sourcing behavior more generally.[2]

While the gains and consequences from trade in intermediate inputs have been studied extensively in recent years (Antràs & Chor, 2021), it is surprising how little is known about disaggregated input-output relationships at the level where trade and industrial policy operate. The most frequently used source of information in quantitative trade models restricts observations of input-output relationships to about 20 manufacturing sectors.[3] Supply-chain disruptions and policy interventions, however, occur at much more granular levels, which implies that aggregated sector-level information can substantially mismeasure the degree of foreign input use, the concentration of critical suppliers, and thereby obscure the propagation of shocks through production networks.[4]

This paper proposes a simple and flexible machine-learning approach to recover product-level input-output linkages from widely available firm-level trade data. By adapting an Association Rule Mining (ARM) approach and disciplining it based on well-known stylized facts about firms in international trade, we generate new insights on critical input-output relationships at the product level. We implement and calibrate our method on a cross-sectional firm-level dataset from China, with detailed information on imports and exports at the product level. Based on this data and auxiliary firm- and product-level information, we infer which goods that are sourced by a firm are (likely) used as inputs for specific goods the firm sells. Our method draws on data-mining tools, which we discipline with two economically transparent assumptions: some inputs are "critical" in the sense that they are difficult to substitute in the short run, and the sampled firms share broadly comparable production technologies. Under these conditions, the large number of firm-product observations becomes

[1] See Johnson (2014) and De Backer & Miroudot (2014) for early accounts and discussions of the evolution of global value chains.

[2] See, for example, the work of Boehm et al. (2019) and Carvalho et al. (2020) on the earthquake and subsequent tsunami in Japan in 2011, Barrot & Sauvagnat (2016) on natural disasters, Bonadio et al. (2021) on the Covid-19 pandemic, or Flaaen & Pierce (2024) and Handley et al. (2020) on the US-China trade war.

[3] This refers to the globally comparable World Input Output Database (WIOD), documented by Timmer et al. (2015). Input-output tables available at the US Bureau of Economic Analysis (BEA) are more disaggregated, but only available for the US.

[4] See Cox (2025) and Sutter (2025) for concrete examples of US import and export restrictions that highlight this discrepancy.

informative about otherwise hidden input-output linkages, allowing us to map supply chains at a level of granularity well-suited for trade and industrial policy analysis.

The aim of this paper is to quantify how economic shocks reshape trade flows by propagating through granular input-output (IO) relationships. Recent trade and industrial policy has shifted toward supply-chain de-risking and resilience, which motivates our analysis, but the issue is more general. Understanding adjustment and incidence in modern trade requires measuring interdependence at the product and firm margins where sourcing decisions are made, substitution is limited, and policy is applied. This matters not only for today's resilience agenda, but also for interpreting the distributional consequences of trade during the still-recent episode of trade liberalization.[5] To bring the data to bear on these mechanisms, we construct a mapping of input-output linkages at the 4-digit level of the Harmonized System (HS), using the universe of Chinese firm-level imports and exports in 2004. In our main application, this firm-level IO structure lets us trace how China's export-led expansion in the early 2000s translated into improved export opportunities for foreign economies supplying the critical intermediate inputs embodied in China's growing export basket, links that would be largely invisible in sector-level IO accounts.

Our setting faces a number of challenges for identification that arise partly from the nature of the data we use, but also from the fact that firms generally tend to engage in multiple activities that are difficult to disentangle. Indeed, multi-product firms dominate manufacturing good production and trade in both advanced and developing economies (e.g., Eckel & Neary, 2010; Bernard et al., 2011; Boehm et al., 2022), which significantly complicates the inference of product-level input-output relationships. Furthermore, regarding the nature of our data (firms' product-level export and import transactions), it is evident that domestically sourced inputs or domestically sold outputs remain unobserved, similar to firms' inventories. Carry-along-trade, exports of goods that the firm has not produced itself (Bernard et al., 2019; Erbahar & Rebeyrol, 2023), can also lead to misleading associations between imported and exported products. We discuss these challenges in detail and implement different calibrations of our algorithm, and data processing steps, to minimize measurement error, which we evaluate in extensive validation exercises. A central contribution of our paper is, therefore, to provide a critical review and empirical assessment of alternative approaches to measuring granular input-output linkages, together with a discussion of when and how they can be transported to other settings.

We document several novel findings. First, product-level input-output networks are sparse: most products have only a few upstream or downstream linkages, while very few

[5]Recent examples are the EU's Open Strategic Autonomy and the Critical Raw Materials Act, China's self-reliance plans, as well as the US Chips Act and Inflation Reduction Act.

products have many linkages. This highlights the susceptibility to (compounding) measurement error in more aggregated input-output data (e.g., Diem et al., 2024), as well as the importance of accurate measures for effective policy making. Second, our inferred mappings are positively and significantly correlated with more aggregated IO coefficients, but less so for inputs where China has a revealed comparative advantage. This indicates that our inferred IO linkages describe international (i.e., cross-border) product-level supply-chain relationships, which is distinct from mappings that are based on technological feasibility (e.g., Fetzer et al., 2024). Third, our mappings perform very well in external validation exercises, but we emphasize that purely algorithmic inference might be insufficient. We find that sampling strategies that account for firm and product-level characteristics can improve the external validity of predicted linkages. We demonstrate this by showing that our preferred mappings, inferred from Chinese data, can reconcile product-level import and export dynamics also for (arguably) similar Asian economies. Fourth, using firm-product-level transactions data for 20 economies and multiple years, from the *Panjiva Global Trade Insights* database, we find that recovered mappings evolve smoothly over time within countries, and cluster sensibly across countries. Manufacturing-based economies exhibit similar upstream–downstream relationships. This evidence suggests that our method recovers persistent and gradually evolving production structures, while attention should be paid to sample size (i.e., number of observed firms per product) and the overall level of export diversification when implementing our algorithm.

We then use our method in two applications to quantify how changes in product-level export growth translate into changes in product-level imports of associated intermediate inputs. Using China's WTO accession as a motivating episode, we find that a 1 percentage point increase in export growth is accompanied by a 0.21-0.49 percentage point increase in growth of imports of inferred inputs. Countries with a revealed comparative advantage in these inputs benefited from China's demand shock through new and economically meaningful export opportunities, over and above the market-access effects operating through other post-accession channels (such as tariff reductions). In a second application, we focus on the 2018-2019 U.S.-China trade war, which resulted in sharp decline of China's exports to the US in products subject to additional US tariffs. Our mapping traces this contraction upstream to lower imports of the corresponding inputs, as well as to core supplier economies (Mexico, India, Vietnam), who reduced their exports to China and also imported fewer inputs. Both applications illustrate how our mapping, and granular information on input-output linkages more generally, can enhance the empirical evaluation of (trade) policy shocks along cross-border supply chains.

Related literature. Our paper relates to a recent but fast growing literature that develops novel approaches to identify disaggregated (product-level) input-output relationships. One strand of this literature draws on administrative data in case-study settings, such as Cajal-Grossi et al. (2023) or Cox (2025). While such research designs clearly benefit mitigating important identification problems, they remain highly contextualized to the industry (and country) they study. Another strand of this literature relies on expert knowledge and surveys to identify "essential" production inputs (Pichler et al., 2020). A similar approach is followed in the compilation of critical material lists administered by national or supranational government institutions. Fetzer et al. (2024) take this logic further and deploy a large language model (LLM) to construct, with the help of AI, a mapping of disaggregated product-level input-output linkages for all internationally traded goods. While valuable and informative, limitations of such approaches are their potential bias and "blind-spots" as well as limited reproducibility. A third strand of this literature, pioneered by Vandenbussche & Viegelahn (2018) and recently adapted by Ma et al. (2025), relies on data-driven inference. These papers rely on observed input purchases by single-product firms to create granular IO mappings. Berlingieri et al. (2021) derive industry-level input-output linkages by combining balance-sheet and trade transactions information of French firms. Our approach expands on these strategies by exploring simple pattern recognition techniques that can be flexibly employed on widely available customs statistics with minimal data requirements.

Doing so, our paper relates also more broadly to an economics literature that uses network-inference tools to describe product and firm characteristics. While the product-space and economic complexity index of Hidalgo et al. (2009) pioneered such approaches, more recent studies used data-driven inference to describe the relationship stickiness of traded goods as a proxy for supplier substitutability (Martin et al., 2025), firm-level buyer-supplier networks (Inoue & Todo, 2025), or the connection between firms and the exposure of their establishments to trade shocks (Boehm et al., 2025). Conceptually and methodologically most closely related to our paper are Macedoni et al. (2024), who infer the product-space of firms via systematically co-occurring products in their export basket, as well as Karbevska & Hidalgo (2025) and O'Clery et al. (2025), who identify product-level import-export co-occurrences to describe the global production network at an aggregated level. Our work expands on these methods, but accounts for product characteristics and well-known stylized facts about firms in international trade, to obtain improved link predictions which enable us to empirically evaluate how trade policy shocks propagate via supply-chain linkages.

Finally, through our applications, we make novel contributions to the empirical trade literature that seeks to understand and quantify the propagation of shocks along cross-border supply chains. In this context, our first application offers a more nuanced interpretation of

the so-called "China shock" (e.g., Autor et al., 2013, 2016; Pierce & Schott, 2016), providing new evidence of economically significant export opportunities that arose. While such channels have been noticed before (e.g., Greenaway et al., 2008; Dauth et al., 2014), data limitations prevented systematic analyses of the underlying mechanisms. Our granular mappings enable us to revisit this question, where our findings suggest that a considerable share of the arising export opportunities to China consisted of supplying inputs for their exports. Our second application contributes to ongoing research on the supply-chain effects of the U.S.-China trade war. Flaaen & Pierce (2024) use US input-output tables to investigate how the tariff shock affects downstream and upstream industries, while Utar et al. (2025) study the impact on reshoring and sourcing patterns of (indirectly) exposed Mexican firms. Most closely related to our paper, Handley et al. (2025) find that US exporters are negatively affected by the trade war when facing higher tariffs on imported inputs (from China). Our evidence documents the upstream propagation of the U.S.-China trade war, which results in lower input demand and fewer export opportunities for third countries. Besides the specific insights gained in the context of these applications, our work showcases how information on granular input-output linkages can be employed in the empirical evaluation of trade (policy) shocks. This will be important to complement quantitative studies that typically rely on much more aggregated input-output relationships, or evaluations that use firm-to-firm networks to bypass an explicit matching of inputs and outputs (e.g., Barrot & Sauvagnat, 2016; Boehm et al., 2019; Carvalho et al., 2020; Borin et al., 2025).

The rest of the paper is organized as follows. Section 2 characterizes the methodological challenge underlying the inference of granular input-output linkages and introduces our method to tackle this problem. Section 3 describes the data we use for our baseline mapping procedure and explains the implementation of our method, including potential identification challenges that are specific to our data. Section 4 describes and evaluates the obtained linkages through a set of validation exercises, emphasizing their predictive performance across settings. In section 5 we apply our mapping to trace and quantify how trade shocks transmit upstream through the supply chain. Section 6 concludes.

2 Background and methodology

This section defines the object of interest, the product-level input requirements for each output, and explains how we recover it from trade microdata. Because these linkages are rarely observed, we infer them from systematic co-occurrence patterns in firms' imported inputs and exported outputs, using auxiliary information in the customs records to discipline the mapping. We begin by stating the inference problem and its key identification challenges,

and then describe the methods we use to construct and validate the correspondence.

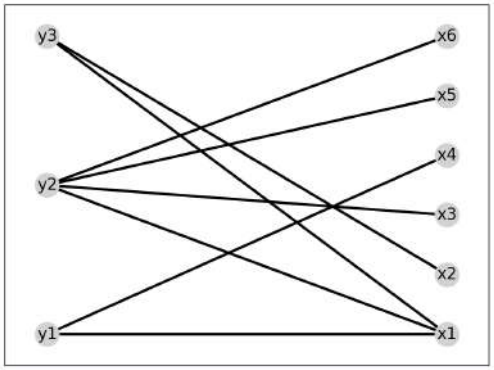
2.1 The problem

To fix ideas, consider a generic production technology for output good i,

$$y_i = f(X_i, \cdot). \tag{1}$$

The production of an output good i requires combinations of raw, processed and manufactured intermediate inputs (X_i), next to other factors like capital and labor. The main question we ask in this paper is: what are the specific inputs in vector $X_i = (x_{i1}, ..., x_{ij}, ..., x_{iJ})$ that are needed to produce y_i? Answering this question is difficult. Production takes place within firms and can typically not be directly observed. Moreover, it is common that firms produce multiple products (Irlacher, 2024), so that their sourced inputs cannot readily be linked to outputs. Figure 1 illustrates this by displaying a stylized firm that purchases inputs x_j to produce multiple outputs $y_i \neq x_j$. While we can reasonably assume that connecting lines $x_j \rightarrow y_i$ do exist, they cannot be observed.[6]

Figure 1: Input-output relations within a firm as a bipartite network representation



Note: Figure displays relationships between output goods y_i (left) of a firm and the inputs x_j used for their production (right) as a bipartite network. Edges indicate which inputs j are used for a respective output good i.

Any attempt to make such granular input-output linkages visible, inevitably faces a trade-off between accuracy and comprehensiveness. High levels of accuracy can be obtained through in-depth case studies and the exploitation of detailed contextual knowledge (e.g., Pichler

[6] An additional problem arises from the fact that the production of y_i often involves a sequence of steps, implying that x_{ij} itself is the output of a production process (i.e., $x_{ij} = g(Z_{ij}, \cdot)$) that can take place either within or outside of firms' boundaries (e.g., Alfaro et al., 2016, 2019). Our figure, hence, already incorporates the simplifying assumption that combinations of x_j are directly converted into y_i.

et al., 2020; Cajal-Grossi et al., 2023; Cox, 2025). However, this often comes at the expense of comprehensiveness, as the focus typically lies on a narrow set of products or firms. The scope for replication and external validation is limited, and the value of such mappings is mainly their strong identification in the context of a specific research question. Alternatively, researchers can choose to sacrifice some amount of accuracy when inferring product-level input-output linkages. Vandenbussche & Viegelahn (2018), for example, observe raw materials purchases of single product firms, which allows them to generate a fairly comprehensive mapping to study upstream and downstream effects of a trade liberalization. Ma et al. (2025) also follow such an approach, which nevertheless comes with limitations originating from the uniqueness of the underlying data and a focus on specific types of inputs or firms.

Our approach deliberately takes multiproduct firms into account, and uses a method that enables us to detect patterns in widely available and harmonized firm-level customs data, providing ample scope for external validation and replication. Given our reliance on a routinized pattern recognition (i.e., machine-learning) technique, we devote some space to discuss important properties of our method, the challenges that come with the underlying data we use, and how these can be addressed. We start by introducing our methodology.

2.2 Association Rule Mining (ARM)

Our approach is based on Association Rule Mining (ARM), a simple pattern recognition technique for bipartite network inference.[7] Based on our stylized representation in Figure 1 above, we argue that ARM can, in principle, also be used to infer the type of input-output linkages we are interested in.

Key statistics. ARM involves a statistic that indicates how likely it is that a specific observation — in our case, an input-output relationship x_{ij} — appears in the data. This statistic is based on two measures: the so-called *Support* (S_k) of an item and the *Lift* (L_{ij}) of the bundle. $S_k = \frac{N_k}{N}$ measures how frequently a product $k = \{i, j\}$ (or product bundle ij) appears in the data. The numerator N_k represents the number of cases in which k is observed, while N is the total number of observations. In the context of our paper, N represents the total number of firms, while N_k measures the number of firms, for which k is observed. The support S_k, hence, simply measures the prevalence of k among firms.

To detect systematic combinations of purchases j and sales i across firms, we need to decide whether or not the joint observation of j and i within a firm is random. To indicate

[7] ARM was introduced by Agrawal et al. (1993) to identify recurring patterns or relationships between items in transaction databases or other data repositories (Zhao & Bhowmick, 2003). Since then, it has been used in a range of different settings, including marketing research, computer networks, recommendation systems, and healthcare (e.g., Solanki & Patel, 2015; Telikani et al., 2020; Kaushik et al., 2021).

this, the *Lift* of the bundle ij expresses how the co-occurrence of j and i relates to the combined individual occurrences of each good across firms:

$$L_{ij} = \frac{S_{ij}}{S_j \times S_i} \tag{2}$$

Whenever the resulting measure is $L_{ij} > 1$, the co-occurrence ij is a more than random appearance. Conversely, whenever $L_{ij} = 1$ or $L_{ij} < 1$, the occurrence of this bundle cannot be distinguished from a random co-occurrence or even be unlikely.

Refinements. The simplicity and logic of ARM are appealing, but has an important limitation in the context of inferring granular input-output linkages. As a comprehensive mapping describes the existence and strength of backward and forward linkages between hundreds or thousands of different products that are used and produced by different industries, many (or even most) combinations $j \to i$ are irrelevant. The network we aim to infer is sparse, which requires a large sample to discover the relevant linkages for all products. At the same time, even though multi-product firms are common, firms specialize in producing a limited range of products. This implies that the denominator in Equation (2) quickly converges to zero, and inflates L_{ij} exponentially, if a good is typically produced or purchased by only a few firms.[8] A score $L_{ij} > 1$ can consequently mean two things: either, there is a systematic co-occurrence between i and j that correctly reflects an input-output relationship between these goods, or only very few firms (compared to the rest of the sample) produce i or purchase j. The latter would wrongly suggest a systematic relationship but actually reflect a property of the underlying data.

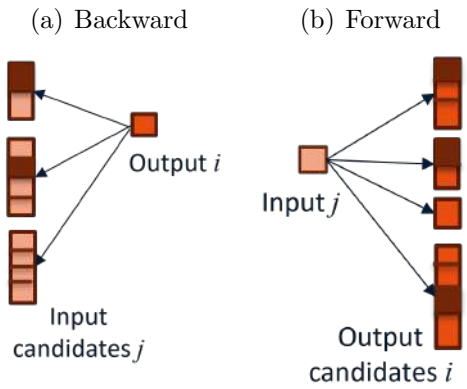
To address this potential source of misinterpretation, we refine the ARM process and use Bidirectional Frequency Association (BFA), which is conceptually reminiscent of the backward-forward algorithm used by Karbevska & Hidalgo (2025).[9] The procedure becomes immediately apparent, while the trick is to use product-specific denominators when S_ks are calculated.

BFA is implemented in two steps, starting, for example, with the set of *input candidates* j for the first item in the list of observed output goods $i \in I$. Only firms that produce i are taken into consideration, so that N_i denotes the total number of such firms. $S_{ij}^{\text{back}} = N_i(j)/N_i$ then denotes the fraction of all i-producing firms that purchase j. The example displayed in

[8] Economic realities, such as the existence of monopolies and heterogeneous market structures across sectors underscore this concern.

[9] Bidirectional (i.e., two-way) associative search and modeling is widely used in other disciplines that model networks and interdependent mechanisms. See Kosko (1987) for an early description of bidirectional association models.

Figure 2: Stylized representation of Bidirectional Frequency Association (BFA)



Note: Author's representation. Each square represents product that is either produced (i) or purchased (j). Stacked squares indicate observational units (i.e., firms) that use input candidates j to produce output i, in Panel (a) or produce output candidates i while sourcing j, in Panel (b).

Figure 2(a) suggests that two out three firms that sell output i also purchase the highlighted input j. To account for the possibility that j is a ubiquitous input or that i is typically produced by firms next to many other goods, forward linkages are computed in an analogous way. As shown in Panel (b) of Figure 2, $S_{ij}^{\text{fwd}} = N_j(i)/N_j$ expresses how many j-sourcing firms also sell i. Intuitively, inputs with very specific use will score higher, while widely used inputs score lower. The final step combines the two scores into a single measure:

$$S_{ij}^{\text{BFA}} = S_{ij}^{\text{back}} \times S_{ij}^{\text{fwd}} = \frac{N_i(j)}{N_i} \times \frac{N_j(i)}{N_j} \tag{3}$$

By multiplying S_{ij}^{back} and S_{ij}^{fwd}, we penalize high scores for input candidates that are also used in a wide range of other production contexts, while we "reward" those that are very specific to one or a small set of activities. In other words, $S_{ij}^{\text{BFA}} \in [0, 1]$ captures the likeliness that shock observed in j propagates to i an vice versa. At the same time, while maintaining the logic and simplicity of ARM, the BFA approach is independent of the overall sample size and instead puts greater emphasis on the actual number of co-occurrences, N_{ij}.[10]

To sum up, our approach builds on the assumption that a genuine input-output linkage should leave a clear footprint in firms' trade bundles: firms that sells a given output tends to purchase the relevant input, and firms that purchase that input tend to sell the output. ARM formalizes this footprint by identifying pairs that co-occur more often than their overall

[10]To see this, note that $N_i(j) = N_j(i) = N_{ij}$ by definition. This implies that Equation (3) can also be expressed as $S_{ij}^{\text{BFA}} = \frac{N_{ij}}{N_i N_j} N_{ij}$, whereas Equation (2) can be written as $L_{ij} = \frac{N_{ij}}{N_i N_j} N$. A high ARM score is, hence, likely to be obtained if N increases faster than N_i or N_j. This is not the case for our refined measure, which increases only with N_{ij} but ensures that $S_{ij}^{\text{BFA}} \in [0, 1]$.

prevalence would predict. Our BFA refinement tailors this approach to the underlying data structure and requires the pattern to hold in both directions. Doing so, we also reduce chance overlap, similar to O'Clery et al. (2025), by down-weighting inputs that are used broadly across many unrelated activities.

3 Data and implementation

This section describes how we implement BFA to infer product-level input output linkages from firm-level trade data. We first highlight and discuss some general relevant features of such data, and how they may impact the accuracy of BFA. We then describe and detail the construction of a "training sample" from the China Customs Database, and lay out alternative pre- and post-processing steps for our data to mitigate measurement error.

3.1 Properties of firm-level trade data

Applying BFA to firm-level trade transactions data has many advantages, but also comes with challenges that are important to understand as they inform decisions on appropriate sampling strategies and data processing steps before and after implementation.

3.1.1 Advantages

Harmonized product codes. A first and major advantage is that product codes in international trade datasets and customs statistics are standardized and internationally comparable. Indeed, different products and their detailed descriptions are systematically classified according to the *Harmonized System* (HS) nomenclature, which is administered and regularly updated by the World Customs Organization (WCO). Correspondence tables further enable comparability over extended periods of time. National customs offices adopt these classifications and sometimes introduce additional, more granular differentiations between products. Nevertheless, about 5,000 products can be distinguished the 6-digit HS level.[11]

Relevance. The second advantage is that the input-output linkages inferred from international trade data can be almost immediately applied to study important questions related to supply-chain disruptions and international interdependence (Fetzer et al., 2024). While

[11]The HS has a hierarchical structure. Six-digit codes are nested within roughly 1,500 four-digit product categories, which in turn fall under 96 two-digit HS chapters and broader HS sections (groups of chapters denoted by Roman numerals). The HS covers physical goods that are manufactured, grown, or extracted; it does not cover services. This is a limitation, although services are typically less traded internationally and are not recorded through customs.

previous studies have to rely on more or less accurate crosswalks to combine information that is based on different classifications and underlying concepts, trade (policy) shocks are typically observed at the HS-product level, so that their upstream and downstream propagation should also be studied along these dimensions (Handley et al., 2025; Cox, 2025).

Standardization and scope for refinements. Finally, even though firm-level customs records are collected and maintained by authorities from different countries, they are administered and reported in a relatively simple and consistent format, which makes our method transparent and easy to replicate. In its most basic form, the approach requires only a firm identifier, the list of products the firm exports, and the list of products it imports. Additional information, such as the customs regime or the firm's primary activity, are optional and can be used to refine the sample or the inferred mapping. We discuss these refinements alongside the construction of our training sample below.

3.1.2 Potential caveats

Besides these advantages, firm-level trade data also has characteristics that are important to understand and may have repercussions for the scope and type of research questions inferred mappings can address.

Domestic transactions. The most obvious challenge arises from the fact that customs statistics do not report domestic transactions. The product range of a firm's export or import basket might therefore represent only a subset of the actual basket. As a consequence, if inputs are sourced to produce a good that is sold only domestically, BFA produces false positive associations. Conversely, if firms source specific inputs primarily from domestic suppliers, false negative associations might be the result. There is no obvious solution to this potential caveat,[12] but the distortionary effects of domestic transactions are likely to be more relevant if trade data is used from a country with large domestic markets and, more generally, in industries in which the country has a comparative advantage.[13]

We furthermore stress that systematic mismeasurement due to unobserved domestic transactions arises only, if these take place systematically across firms. Yet, by focusing on firms that trade internationally, we extract information from the subset of an economy that is most

[12] Even if supplementary data on the scope of firms' domestic activities was available, detailed information on which products they purchase and sell is typically not available.

[13] Intuitively, if an economy is abundant in a particular input, firms may source it domestically and specialize in exporting downstream goods that use it intensively. For instance, a wood-abundant country may export furniture or paper while importing little raw wood. In that case, the input is central to production but largely absent from firms' lists of imported intermediates.

productive (Bernard & Jensen, 1995; Bernard et al., 2007) and sources and sells globally (Kugler & Verhoogen, 2011; Topalova & Khandelwal, 2011; Manova & Yu, 2017). As such firms tend to specialize in exporting their globally most competitive goods (Bernard et al., 2011; Mayer et al., 2014; Eckel et al., 2016), it is likely that their import and export transactions carry the footprint of actual input-output relationships, which we aim to detect.[14]

Carry-along trade and trade intermediation. A second challenge arises from the practice of some firms to sell goods that they did not produce. This phenomenon, known as "Carry-along Trade" (CAT) or trade intermediation, has been documented for several countries (e.g., Bernard et al., 2019; Eckel & Riezman, 2020; Erbahar & Rebeyrol, 2023), while firms' underlying rationales and mechanisms are still not fully understood. In our context, such behavior is prone to produce false-positive linkages, especially if the trade of certain product categories is predominantly handled by such firms. Although we are unaware of any evidence of such regularities, we do note that trade intermediaries are typically more diversified than similarly-sized firms in terms of their product portfolios, while being more specialized in terms of their export destinations and source countries. Such features, together with other information in the data (e.g., on customs regimes or firms' industry affiliation) might help to identify suspect intermediaries in the sample and should be removed before deploying BFA.

3.2 Sample

Selection and basic cleaning. We deploy our algorithm on firm-level trade transactions reported in the *China Customs Database*.[15] We focus on a single cross-section for the year 2004, which falls into the episode of China's integration into the global economy. We chose this sample for two reasons. First, China's sheer economic size makes it more likely to observe sufficient numbers of firms that sell and source a wide range of products internationally. Second, China's economic expansion in the early 2000s was substantially based on inputs sourced from abroad (Amiti & Freund, 2010). Both of these features should work in favor of our efforts to obtain a comprehensive and reasonably accurate mapping.

The raw data provides information on the value of each firm's exports and imports, by product, trade partner, as well as the customs "regime" under which imports and exports have been administered by the Chinese authorities. It can include assembly and processing

[14] As noted our description of Figure 1 in Section 2, we make the simplifying assumption that inputs are directly converted into outputs and that firms share similar production technologies.

[15] This data has been used extensively in the international trade and development literature. See, for example Manova & Zhang (2012); Manova & Yu (2016); Chor et al. (2021).

trade, trade through bonded area warehouses, or ordinary trade, among others. To reduce the risk of obtaining false positive mappings, transactions that can be considered as unrelated to a firm's production processes are excluded.[16] Moreover, to ensure we can relate our findings to product characteristics in other data sets, we remove observations of trade in goods that are unknown to the HS Rev. 2002 nomenclature at the 4-digit level (HS4).[17] Finally, we aggregate the data to the firm-product-flow level and keep only those firms that exported and imported at least one good. Our resulting baseline sample consists of 66,688 firms trading 1,361 different HS4 products, which we consider an appropriate level of detail.[18]

Descriptive statistics. Table 1 describes the structure of our sample. Panel A shows that most firms import and export a relatively small range of products. The number of goods imported tends to exceed the number of goods exported by a factor of about two. A substantial number of firms (17,922 or 26.87% of the sample) export only a single HS4 product. Moreover, Panel B shows that most HS4 products are exported and imported by several firms (more than 500 on average and around 170 for the median product). We consider these features as favorable for our purposes, as they suggest that firms are rather specialized in terms of their product range, while the linkages we aim to infer will largely rely on a reasonable amount of observations.

Table A.1 provides further insights on the type of goods Chinese firms traded in 2004. Columns (1)-(3) indicate that both exports and imports are relatively concentrated in a few sectors.[19] Column (4) indicates that these products are also traded by large numbers of firms. On the import side, concentration is somewhat weaker but still pronounced, with many imports falling into end-use categories typically classified as intermediates or capital goods.[20] In light of these patterns, we expect that the accuracy of inferred input-output linkages might vary across sectors. At the same time, we obtain guidance on how to control for potentially confounding factors, for example, by taking into account China's own revealed comparative advantage in imported products. Similarly, the prevailing end-use category of

[16] These include trade of bonded warehouse inbound and outbound goods; bonded area warehousing re-export goods; border microtrade; compensation; aid and gifts from countries and international organizations; consignment or consignment trade; duty-free foreign exchange goods; barter and leasing trade.

[17] To identify these, we use the list codes available at the United Nations Statistics Division (UNSD), https://unstats.un.org/unsd/classifications/Econ, which we accessed on 20 September 2024.

[18] Working at the HS4 level furthermore ensures sufficient observations per product, while retaining economically meaningful detail. More disaggregated HS6 codes refer to so-called sub-headings that typically differentiate between varieties of a good.

[19] About one quarter of all traded HS4 product lines account for almost two-thirds of China's export revenues in 2004. These are, namely, Textiles, Clothing and Apparel (HS Chapters 50-67) as well as Machinery and Electronic products (Chapters 85-86).

[20] End-use categories are based on concordances to Broad Economic Categories (BEC Rev. 5) from the UNSD, which distinguish capital goods, intermediate inputs, consumption goods, and combinations thereof.

Table 1: Number of HS4 products imported and exported by firms in 2004

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Mean	Std	Min	25%	50%	75%	Max
Panel A - Product diversification of firms							
HS4 products exported	10.83	29.42	1	1	3	7	556
HS4 products imported	14.65	25.03	1	2	6	17	811
Panel B - Product ubiquity							
Firms exporting a HS4 product line	544.56	985.30	1	41	174	580	11,676
Firms importing a HS4 product line	734.47	1,568.26	1	35	166	619	18,074

Note: Calculations use 2004 firm-level China Customs transactions after restricting to the trade modes described in the text and aggregating products to HS4. The sample includes 66,688 firms and 1,361 unique HS4 product lines. Panel A reports the number of distinct HS4 products exported and imported by firms. Panel B reports, for each HS4 line, the number of distinct firms that export or import it.

imported goods will help us to determine whether inferred relationships can plausibly be interpreted as input-output linkages, or whether they pertain to auxiliary transactions that are unrelated to production processes.

3.3 Implementation

Based on the above discussion and features of our data, we implemement BFA alongside further pre- and post-processing stages in three different ways.

Baseline mapping. We start with a baseline mapping that we obtain by implementing BFA without further refinements. Hence, we make use of the transactions reported for 66,688 firms to obtain a matrix with 1,361×1,361 products, where each cell reports a score $S_{ij}^{\text{BFA}} \in [0, 1]$. Our only restriction is that $i \neq j$, which means that a product cannot be used as an input for itself. We can subsequently inspect the distribution of these scores to identify positive or negative outliers to decide which input candidates are most likely essential for the production of an output good. As an *ad-hoc* rule, we label an input–output pair (i, j) as "essential" based on the relative magnitude of its BFA score within relevant groups. Specifically, a link (i, j) is classified as essential if $S_{ij}^{\text{BFA}} > 2\ \bar{S}_i^{\text{BFA}}$ *and* $S_{ij}^{\text{BFA}} > 2\ \bar{S}_j^{\text{BFA}}$ where $\bar{S}_i^{\text{BFA}} \equiv \frac{1}{N(i)} \sum_{j \in \mathcal{J}(i)} S_{ij}^{\text{BFA}}$, $N(i) \equiv |\mathcal{J}(i)|$ and $\bar{S}_j^{\text{BFA}} \equiv \frac{1}{M(j)} \sum_{i \in \mathcal{I}(j)} S_{ij}^{\text{BFA}}$, $M(j) \equiv |\mathcal{I}(j)|$. Here, $\mathcal{J}(i)$ denotes the set of all candidate input-output links involving output i, while $\mathcal{I}(j)$ denotes the set of all candidate links involving input j. Essentially, for each output, most candidates receive little support, while a few stand out. Comparing a candidate's score to the mean among non-zero candidates isolates inputs that are unusually strong for that output,

without diluting the benchmark with zeros. This threshold is best viewed as a transparent way to impose sparsity, not a tuned parameter.[21]

Post-processing: capital vs. intermediate inputs. Since our data includes different types of products we submit the baseline mapping to a number of minor, but potentially important post-processing steps. Doing so, we consider the set of product i's input candidates from our baseline mapping and inspect their BEC end-use category. We then require that an input must never represent a final consumption good and set $S_j(i)^{\text{BFA}} = 0$ whenever this is the case. Next, we apply our threshold for strong linkages separately for capital and intermediate inputs. We thereby acknowledge that capital goods are purchased much less frequently than material inputs. While the former are used for a sustained period of time and depreciate slowly, the latter are effectively consumed in the production process and have to be constantly "renewed".

Pre-processing: suspect trade intermediaries. Although we cannot directly observe and identify non-producing firms, such as retailers and trade intermediaries in our training sample, we generate an third mapping after undertaking additional efforts to clean our data before deploying BFA. Doing so, we inspect the level of diversification of firms in terms of their number of goods exported and remove potential suspect cases from our sample. We set a deliberately tough threshold at the 90th percentile, which still leaves us with more than 60,000 firms, but implies that no firm exports more than 19 HS4-level product lines. We do not *a priori* put a limit on the number of imported products, but conduct the same post-processing steps as outlined in the previous paragraph.

Figure A.1 summarizes our end-to-end pipeline. Starting from firm-level import and export transactions, we clean our data and form product "bundles" to execute AI-style mining of co-occurrences, scoring candidate input-output links. We apply additional economic filters to rule out implausible inputs and retain the strongest links, which yield a granular HS4-by-HS4 mapping. The next section describes our obtained mappings.

[21] At the HS4 level the candidate set is extremely large and the score distribution is typically heavy-tailed, with many near-zero candidates and a small set of inputs that stand out for a given output. Comparing each candidate to the mean among *non-zero* candidates preserves this sparsity and yields an interpretable "essential-input" set, while avoiding mechanically lower cutoffs that would arise from averaging in masses of zeros. Importantly, the continuous S_{ij}^{BFA} scores remain available for exercises that use the full weighting structure rather than a binary selection.

4 Main findings

The BFA mapping delivers an intuitive picture of production linkages: upstream inputs funnel into downstream outputs in patterns that are stable across implementations. We then show that these linkages align with external benchmarks and predict systematic co-movement between upstream imports and downstream exports.

4.1 Descriptive patterns

Figure A.8 summarizes the input-output mappings obtained from the three implementation variants. For ease of exposition, we aggregate the HS4-by-HS4 matrices to the HS2 chapter level. The resulting patterns are highly stable: the baseline and post-processed mappings in Panels (a) and (b) have a correlation of 0.97, while the baseline and fully processed mapping in Panels (a) and (c) have a correlation of 0.85. At the same time, the matrix becomes progressively sparser as we add restrictions, consistent with the removal of potential false positives in (b) and (c). By construction, these steps are less informative about false negatives, but a suggestive pattern is the more prominent row associated with capital goods like industrial machinery (HS Chapter 84) in Panels (b) and (c), which is precisely the type of distinction our BEC-based post-processing is designed to sharpen.

Figure 3: Inferred input-output linkages with alternative implementations of BFA

(a) Baseline (b) Post-processed (c) Pre- and post-processed



Note: Authors' calculations based on BFA applied to HS4-level firm transactions for China in 2004. Panels (a)-(c) differ in the pre- and post-processing steps described in Section 3.3. Scores are aggregated to the HS2 chapter level. Color intensity ranges from zero (white) to the sample average (light blue) and the 95th percentile and above (pink).

Figure A.2 presents some concrete examples of inferred inputs for goods from different sectors. The reported BFA score indicates how confident we can be that these (imported) inputs are actually used in the production of the exported good. Most of these inputs

are intuitively plausible, but we also observe some false positive and negative linkages. For example, for unwrought aluminum (HS4 7601), BFA assigns high scores not only to aluminum scrap but also to other metal scraps, which may reflect technological complementarities in production processes that result in bundled sourcing of inputs. By contrast, under the most conservative specification, aluminum ore is not selected as an input, even though it appears under less restrictive variants. A plausible explanation is that aluminum ore is largely sourced domestically in China, so it rarely shows up in firms' imported input portfolios. This is consistent our previous discussion on potential false-negative linkages for inputs from comparative advantage industries.

Overall, we find substantial variation in the inferred "complexity" of exported goods, as measured by the number of inputs we inferred for a given output. The same holds for the "ubiquity" of intermediate or capital inputs, which we measure by the number of outputs that use a given input. Table A.2 presents an overview of the inferred complexity and ubiquity of outputs and inputs. Depending on the degree of pre- and post-processing in our data, the number of inputs used ranges from less than 5 to more than 100, with a median of 10-25 inputs depending on which method we used. The downstream utilization of inputs varies in a similar range, indicating that some inputs are more widely used (and potentially more generic), whereas others are tailored for specific uses.

4.2 Benchmarks and empirical validation

In this subsection we validate our obtained mappings by comparing them to other data on input-output linkages and by submitting them to a number of out-of-sample empirical validation exercises.

4.2.1 Comparison with other Input-Output mappings

Granular mappings. Because no "ground-truth" of disaggregated input-output mappings like ours exists, we compare our results to the LLM-based AIPNET mapping by Fetzer et al. (2024), which is most similar to ours in terms of scope and detail. Overall, our methods produce similar network densities: while the number of linkages according to our baseline method exceeds that in AIPNET, it is lower in the more sparse mappings we obtain after pre- and post-processing our data.[22] Yet, despite the similar densities, the share of jointly identified linkages is only 16.7% for our baseline mapping and increases to 32.7% in our most sparse mapping. Although the pre- and post-processing steps executed on our data seem to

[22]This is different from O'Clery et al. (2025), who also deploy link-prediction methods, but obtain a mapping with a substantially larger number of linkages than AIPNET; even 4-5 times as many linkages in several industries.

move us closer to AIPNET, two out of three linkages identified with our BFA-based method do not appear in the LLM-based mapping. We draw similar conclusions in terms of the inferred complexity of outputs and ubiquity of inputs: while Figure A.3 shows similar distributions across products, more formal assessments in Table A.3 reveal that — despite a positive and significant correlation — the overall match is quite low. A potential explanation for these discrepancies can be found in the selected examples of Figure A.2, where AIPNET seems to omit auxiliary production materials and equipment (e.g., buttons for clothes, chemicals for rubber tires, paints and varnishes for metal containers, or machinery for material processing), even if they score high according to our method.

These differences clarify what each method is best at capturing. LLM-based mappings are anchored in text descriptions of production technologies and therefore tend to emphasize canonical, engineering-style input requirements. In contrast, our customs-based BFA mapping emphasizes inputs that are *revealed in procurement*, including auxiliary materials, packaging, and equipment that are routinely bundled with production in practice and that appear in trade. Conversely, BFA can understate inputs that are central but predominantly sourced domestically. For applied work, the two mappings are therefore complements: agreement across them signals linkages that are both technologically plausible and transaction-salient, while systematic gaps help diagnose whether a sector is constrained by domestic sourcing, auxiliary-input bundling, or differences in how "inputs" are defined. Altogether, our BFA method captures more broadly defined production inputs, besides being more sensitive to those that are internationally traded.

Aggregated IO tables. We compare our findings also to "traditional" input-output and supply-use tables, which requires us to aggregate our mappings to the corresponding industry classification systems used in these data. Doing so, we check to what extent the HS4-level linkages we identified within given industry pairs matches with the combined sales and expenditure shares in the respective IO tables. Table A.4 summarizes the results of log-linear regressions of the combined input and output coefficients on our aggregated BFA scores.[23] Across all panels in columns (1)-(4), we see statistically significant and positive correlations. The model fit varies, depending on the country and year to which an IO table relates. Yet, the 2004 IO-coefficients reported for China in the WIOD-2016 database (Timmer et al., 2015) produce the highest fit with an R-squared of 0.403 in Panel C of column (1). In more disaggregated Supply-and-Use tables for China (2017) and France (2006),[24] we obtain a lower

[23]We also include input and output fixed effects into our regression to normalize coefficients. Figure A.4 displays the corresponding scatterplots of these correlations.

[24]China's 2017 SUT were downloaded from China's National Bureau of Statistics. French IO for 2006 linkages come from Berlingieri et al. (2021). To map HS4 codes into alternative industry classifications, we

fit but overall comparable results. The match between our mappings and US Make-and-Use tables for 2007 is much lower, despite a persistent positive and significant correlation.

To inspect whether these correlations are co-determined by sectoral characteristics, we re-estimate the correlations between our BFA mapping and China's IO coefficients reported in WIOD-2016 separately for each industry. The resulting R-squares of these regressions are summarized in Figure A.5. They are very low in mining and in miscellaneous manufacturing, but substantially higher in many other sectors–reaching 0.85 for the manufacture of wood and cork products (ISIC C16). Results in Table A.5 suggest that the correlation between our mapping and WIOD-16 is co-determined by China's revealed comparative advantage (RCA). This seems to be the case especially in column (2) where a higher Chinese RCA in the production of inferred inputs weakens the correlation between our mapping and the IO coefficients reported in WIOD-2016. This finding suggests that our methodology performs well in tracing cross-border production chains, while the performance my drop in sectors with (globally) competitive domestic supply capacity.

4.2.2 Validation using import-export comovement

A further check on our mapping is whether it recovers a basic implication of supply chains: faster export growth in downstream products should be accompanied by faster import growth of the upstream inputs they use. We test this prediction with product-level trade data from CEPII BACI for 2003-2008 (Gaulier & Zignago, 2010).[25] Specifically, we regress the average annual growth rate of China's imports in product j ($\%\Delta x_j$) on a mapping-implied measure of downstream export growth:

$$\%\Delta x_j = \alpha + \beta \%\Delta \bar{y}_j^m + \varepsilon_j \tag{4}$$

To measure downstream export growth, we construct our variable of interest as a weighted average of average annual export growth rates in product categories i that are related to imports j:

$$\%\Delta \bar{y}_j^m \equiv \sum_i (w_{ij}^m \times \%\Delta y_i).$$

The weights w_{ij}^m reflect the BFA-scores obtained from our different implementation methods m, which we normalize such that $\sum_i w_{ij}^m = 1$ for each j.

use correspondence tables from the World Integrated Trade Solutions (WITS), available at `https://wits.worldbank.org/product_concordance.html`, and additional crosswalks based on Pierce & Schott (2012) and Brandt et al. (2017). Because these mappings are imperfect, some manual adjustments are required. The limited number of HS4 headings allows us to review ambiguous cases and construct unique matches from HS4 products to the relevant industry classifications.

[25] We use the 2023 release `BACI_HS96_V202301`.

Table 2: Supply-chain linkages in China's imports, 2003–2008

Dep. var.: Import growth ($\%\Delta x_j$)	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	BFA-based mappings						LLM-based	
Link prediction method:	Baseline		Post-processed		Pre-/post-processed		AIPNET	
	HS2	HS4	HS2	HS4	HS2	HS4	HS2	HS4
Downstream export growth ($\%\Delta\bar{y}_j^m$)	0.914**	0.333**	0.915**	0.339**	0.832**	0.355**	0.796**	0.218**
	(0.304)	(0.062)	(0.264)	(0.067)	(0.176)	(0.062)	(0.229)	(0.064)
Observations (# products)	96	1,161	84	894	84	888	92	1,087
R-squared (adjusted)	0.078	0.024	0.117	0.027	0.204	0.035	0.109	0.010

Note: Table reports results of regressing China's average annual import growth rate on the weighted average of corresponding growth rates in related export products. Related goods have been identified with alternative link-prediction methods, as indicated in column titles. Standard errors are reported in parentheses below the estimated coefficients. Statistical significance is indicated as follows: ^{a} $p < 0.1$, * $p < 0.05$, ** $p < 0.01$

Table 2 presents our results for different versions of our BFA-based mappings, as well as the LLM-based mapping from AIPNET for comparison. Although we find statistically significant positive correlations for all methods and at different levels of disaggregation, there are some noticeable differences.[26] Focusing first on columns (1)-(6), we see that the precision of the point estimates increases with additional pre- and post-processing steps carried out on our underlying data. The same holds for the overall fit of the models, measured by the adjusted R-squared at the bottom of each column. While the fit is naturally higher in the more aggregated (and less noisy) HS2-level growth rates, the performance of our granular HS4-level mappings is remarkable.[27] Similar conclusions also hold for the LLM-based mapping in columns (7)-(8), which produces somewhat lower point estimates and a lower R-squared, especially for the HS4-level mapping. As discussed above, we attribute this difference to AIPNET's potentially narrower concept of production inputs, whereas our mappings capture prevalent cross-border supply-chain linkages of a broader concept of inputs, which includes production materials and capital goods.

Finally, we also test if we can trace supply-chain linkages in other countries. Doing so, we focus on a set of eleven Southeast Asian economies and re-estimate Equation (4) for each of these countries.[28] Figure A.7 shows how our BFA-based mappings produce posi-

[26] HS4-level import and export growth rates are winsorized by trimming the top and bottom 1% of their respective distributions to limit the influence of extreme outliers, for which growth rates can differ from the mean by orders of magnitude.

[27] To check the robustness of these findings, Figure A.6 summarizes the results of 1,000 placebo estimates with randomized mappings of our preferred specification in Column (6) of Table 2. In 94.1% of draws, the placebo estimate satisfies $\hat{\beta} < 0.355$, and 79.2% of draws produce negative coefficients.

[28] Southeast Asian economies include: Brunei Darussalam, Cambodia, Indonesia, Lao People's Democratic Republic, Malaysia, Myanmar, Philippines, Singapore, Thailand, East Timor, and Vietnam.

tive and significant correlations for economies that are integrated in global value chains and have production technologies that are potentially similar to China's at that time. The very low-income and least industrialized countries in the sample do not reveal systematic linkages between their export growth and imports of upstream inputs. These results are very consistent across the three versions of our BFA-based mappings and seem to outperform the LLM-based mapping by Fetzer et al. (2024).

4.3 Cross-country correlations and persistence of mappings

To better understand how our mappings correlate across countries and over time, we re-run our algorithm on firm-level trade transactions reported in the *Panjiva Global Trade Insights* database. It reports transaction-level import and export shipments with identified trading parties and detailed product information for a broad set of emerging economies in Asia and Latin America, as well as the United States.[29] Figure A.9 displays the aggregated HS2 level input-output mappings for the year 2017 for our 20 economies, where we can observe partly similar patterns to those shown before for China in 2004 (Figure A.8).[30]

The correlations between the inferred IO-linkages for different countries, which we report in detail in Table A.6, reveal clusters of countries with similar structures. As displayed in Figure A.10(a), countries specialized in exporting raw materials and other primary sector outputs cluster together, while another cluster consists of more manufacturing-based economies. Some countries are outside of these clusters and reside between them.[31] Figure A.10(b) further illustrates that countries' sourcing patterns and global supply-chain integration changes over time. For each country, the correlation between the first available mapping and sub-

[29] *Panjiva Global Trade Insights* is a commercial provider that compiles shipment records from official customs and port documentation. Coverage differs across countries and years because it reflects the availability of underlying administrative sources and reporting regimes. A practical advantage of *Panjiva* for our purposes is that it provides a harmonized data structure across countries, with consistent identifiers for trading parties and comparable shipment-level fields (e.g., product descriptions/codes, quantities, and origin–destination information), which facilitates cross-country benchmarking at a common level of granularity. Appendix B describes sources and coverage in detail; Appendix Table B.1 summarizes the countries covered and the number of firms in each sample. For an external coverage check, Appendix Figure B.1 compares *Panjiva*-based annual import and export totals (raw and cleaned) to UN Comtrade aggregates, showing broadly similar trends for most countries but weaker agreement for a few cases.

[30] For countries where 2017 was not available, we chose the closest alternative year, which is 2020 for Indonesia, 2018 for Pakistan, 2019 for the Philippines, 2021 for Turkey, and 2018 for Vietnam.

[31] The mapping for the United States is most closely linked to other more manufacturing-oriented economies, but is not part of that cluster. A limitation of the U.S. data is that transactions reported in the *Panjiva* database are restricted to ocean-borne shipments, thereby omitting a potentially important share of air freight. In addition, U.S. firms may request the redaction of their own names and those of their trading partners from publicly released bill-of-lading records under a provision known as manifest confidentiality (informally referred to as "opt-out"). *Panjiva* estimates that approximately 20% of U.S. trade records are affected by such name redactions.

sequent years declines, as the time gap widens, indicating gradual changes in supply-chain organization. Even so, a set of economies that form a "global manufacturing base" (i.e, China, India, Mexico, Turkey, and Vietnam) stands out for having both higher correlations and greater persistence over time than other countries in the sample.

Altogether, we conclude that our BFA method produces plausible results for granular input-output mappings. Careful pre- and post-processing of the underlying data can significantly impact and improve their precision. Yet, their validity also hinges on country-specific characteristics, which might require caution in their applicability to third countries.

5 Applications to trade (policy) shocks

In our final section we illustrate how our inferred input-output mappings can be used to study the transmission of trade-policy shocks through cross-border supply chains. We focus on settings where shocks originate in one set of products but plausibly affect other products through upstream input demand and downstream production linkages.

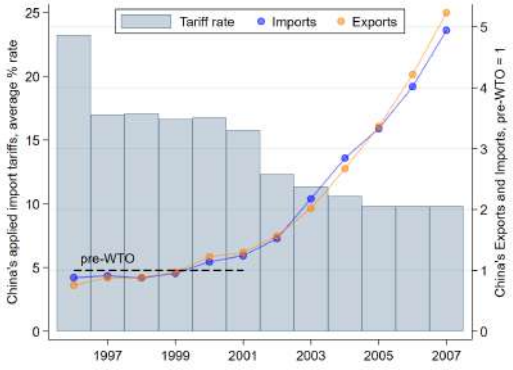
5.1 China's WTO entry in 2001

Our first application revisits the "China shock" literature (e.g., Autor et al., 2013, 2016; Utar & Torres-Ruiz, 2013; Utar, 2014, 2018; Pierce & Schott, 2016; Handley & Limão, 2017; Mau, 2017; Ciani & Mau, 2023). While a central theme in this work is the impact of China's export surge on foreign producers and workers, less is known about the demand-side implications of China's expansion and the role played by input-output linkages. This is striking given the widely documented importance of processing trade during the early 2000s (Amiti & Freund, 2010), which mechanically ties export growth to imported input demand. Figure 4(a) motivates our approach by showing the near-parallel acceleration of China's exports and imports following its WTO accession in December 2001.

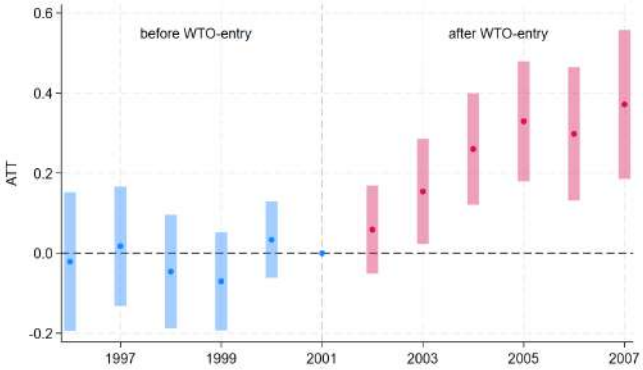
Panel (b) of Figure 4 provides direct evidence of the supply-chain channel linking China's export expansion to its import demand. Inputs mapped to China's fastest-growing exports exhibit a pronounced post-accession rise in imports, relative to inputs mapped to the slowest-growing exports. We construct this comparison in four steps. First, for each HS4 product, we compute China's average annual export growth before and after accession (1996-2001 versus 2002-2007) and take the difference to obtain the *acceleration* in export growth. Second, we trim the top and bottom one percent of this distribution to limit the influence of extreme outliers. Third, we translate each downstream export product's growth acceleration into an input-specific exposure measure using our BFA scores as weights (normalized to sum to one

Figure 4: China's exports and imports before and after WTO entry

(a) Aggregate trends



(b) Exports-related imports



Note: Panel (a) summarizes China's trade expansion around WTO accession. Bars (left axis) report average applied import tariff rates for China from the *World Integrated Trade Solutions* (WITS) database. The lines and markers plot total annual exports and imports of goods, normalized by their mean over the pre-accession period 1996–2001, using CEPII BACI data. Panel (b) reports event-study estimates from the extended two-way fixed effects estimator of Wooldridge (2023). "Treated" inputs are those linked to the fastest-growing exports (top quartile), while "controls" are inputs linked to the slowest-growing exports (bottom quartile).

for each input; see Section 4.2.2). Finally, we split inputs into the top and bottom quartiles of this exposure distribution and trace the evolution of their import growth in an event-study design. Table A.7 shows that the implied pass-through from downstream export growth to upstream import growth is economically meaningful and robust across specifications, with estimates ranging from 21 to 49 percent. Table A.8 then turns to the foreign supply response. We estimate how exports to China in product j from country n changed after China's WTO entry and find that the largest export gains are concentrated precisely in the products our mapping identifies as inputs to China's most rapidly expanding exports. These gains are especially pronounced for countries with revealed comparative advantage in those products ($\text{RCA}_n(j) > 1$), consistent with China's import demand reinforcing existing patterns of specialization.[32]

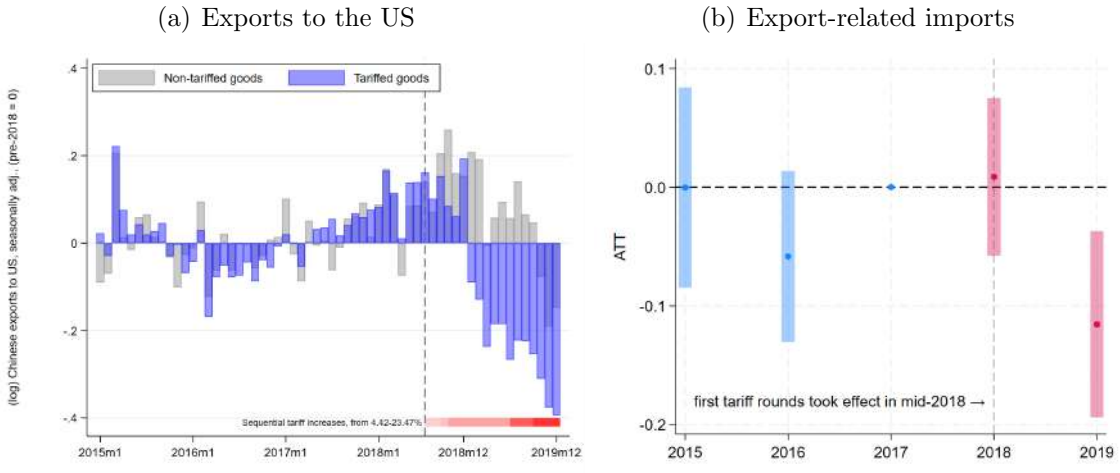
5.2 US–China trade war, 2018-2019

Direct effects of tariffs on exports and imports. Our second application considers the sharp reversal in the trade-policy environment that began in 2018. Unlike the WTO accession episode, which reduced trade barriers and supported a broad-based expansion of China's trade, the 2018-2019 US-China trade war introduced sizable, product-specific tariff

[32]Figure A.11 summarizes, by country and based on RCA, which exporters were most exposed to the rise in China's input demand following its export surge.

increases and triggered rapid adjustments along global supply chains. Panel (a) of Figure 5 shows that China's exports to the United States fell most strongly in the product categories targeted by additional tariffs. While the first tariff rounds took effect in mid-2018, the main escalation — and the bulk of the trade response — occurred in 2019. Panel (b) traces the upstream margin using our BFA-based mapping to imports of products identified as inputs into the tariff-exposed exports. Consistent with the supply-chain mechanism that weaker downstream exports translate into lower demand for upstream inputs, they decline sharply in 2019.

Figure 5: China's US exports and imports during the 2018-19 trade war



Note: Panel (a) traces China's exports to the United States for products placed on the additional U.S. tariff lists versus products that were not. We use the tariff lists and product-level trade data compiled by He et al. (2021), and plot the evolution of (log) exports over time to highlight how the export decline is concentrated in tariffed goods after the tariff rounds begin. Panel (b) asks whether this downstream shock propagates upstream to China's imported inputs. Using our BFA mapping, we link each export product to its imported inputs and construct an input-level exposure measure based on the U.S. tariff exposure of the downstream export goods that use that input (as defined in the main text). We then estimate an event-study difference-in-differences using the extended two-way fixed effects estimator of Wooldridge (2023). "Treated" inputs are those in the top quartile of the exposure distribution, while "control" inputs are those in the bottom quartile. The figure reports the estimated ATT by year (with 95% confidence intervals). In these regressions, we also control for whether China imposed retaliatory tariffs on the input.

Table A.9 quantifies these adjustments. Column (1) formalizes the descriptive patterns above and confirms a sharp decline in China's shipments to the United States after the onset of the trade war. On average, shipments in tariffed product categories fall by 29 percent relative to non-tariffed goods. Column (2) shows that the same pattern emerges in more aggregated data at annual frequency and the HS4 level: the relative decline in China's exports to the United States is 21.8 percent in 2019. Column (3) turns to China's worldwide exports at the HS4 level. Here, the estimated decline is smaller, at 5.4 percent, and less precisely

measured. This attenuation is consistent with two forces. First, exposure to the U.S. market varies widely across products.[33] Second, some of the shock may have been absorbed through trade diversion toward third markets, a margin that aggregate export data cannot cleanly disentangle from an overall export contraction (e.g., Nicita, 2019; Cigna et al., 2022).[34]

To study the upstream response in Chinese imports, we build an input-level measure of exposure to U.S. tariffs that operates through downstream exports. For each imported input j, the measure combines three ingredients. The first is our BFA linkage strength, S_{ij}^{BFA}, which captures how strongly input j is connected to downstream export product i.[35] The second is an indicator $\tau_i^{\text{US}} \in \{0, 1\}$ that equals one if China's exports of product i were subject to additional U.S. tariffs during the trade war. The third is the pre-trade-war importance of the U.S. market for product i, measured by the average U.S. share in China's exports of i over 2015-2017, $\bar{s}_{i,US}$. Weighting by $\bar{s}_{i,US}$ assigns larger exposure to downstream products that relied more heavily on U.S. demand. We define input j's downstream-tariff exposure as:

$$\text{Exp}_{\tau(\text{US})\to j} \equiv \sum_i \left(score_{ij}^{\text{BFA}} \times \tau_i^{\text{US}} \right) \times \bar{s}_{i,US}$$

This variable summarizes how much input j's expected downstream use is concentrated in export products hit by U.S. tariffs. We scale it by 100 for ease of exposition. In the regressions, we further control for whether China imposed retaliatory tariffs on the input, along with product and year fixed effects and product-specific linear trends. Column (4) of Table A.9 shows a clear negative response: a one-unit increase in exposure is associated with about a 2.8 percent decline in imports. Taken together with the descriptive patterns, the estimates indicate that the U.S.–China trade war propagated upstream through China's cross-border supply chains, complementing the downstream (forward-linkage) effects documented for the United States by Handley et al. (2025).[36]

Upstream propagation to third-country suppliers. We now move one step further upstream. If U.S. tariffs depress China's export production, the contraction should be visible not only in China's own input imports, but also in the exports of third-country suppliers

[33] Products subject to U.S. tariffs account for roughly 11 percent of China's total exports. Across HS4 products, the U.S. share ranges from below 0.25 percent to about 30 percent between the 5th and 95th percentiles.

[34] In this section, we compute annual trade flows using the HS2012 classification in the 2025 release of `BACI_HS12_V202501`, to maintain consistent product definitions with the tariff lists.

[35] We use the post-processed BFA scores estimated from China's 2017 shipment data (Section 4). Holding the mapping fixed in a pre-trade-war year avoids mechanically loading the tariff shock into the estimated links, so any subsequent change in imports reflects responses to tariffs rather than re-estimation of the network.

[36] Unlike their earlier working-paper version (Handley et al., 2020), Handley et al. (2025) map imports to downstream export exposure in a way that is close in spirit to our approach.

(TCs) that sell those intermediates to China. To shed light on this channel, we construct a product-level exposure measure for TC exports to China that mirrors our approach above: a TC product is more exposed when our BFA mapping links it more strongly to China's downstream export goods that face U.S. tariffs and that were, pre-trade-war, more reliant on the U.S. market. Among the economies covered in our *Panjiva* data, Mexico, Vietnam, and India stand out as sizable suppliers to China, particularly in the input products that our mapping associates with China's exports to the United States (see Table A.10).[37]

Specifically, we define, for each product j, a binary indicator that equals one if China's imports of product j are exposed to U.S. tariffs through downstream linkages ($\text{Exp}_{\tau(\text{US})\to j} > 0$), and zero otherwise. The output exposure of product j for a third country exporting to China is then obtained by scaling this indicator by the third country's market share in China for product j, denoted $\bar{s}_{j,CN}$:[38]

$$\text{Exp}_{\tau(\text{US})\to j(\text{TC})} \equiv (\text{Exp}_{\tau(\text{US})\to j} > 0) \times \bar{s}_{j,CN}$$

By construction, exposure is high only when product j lies on the tariff-affected supply-chain path and the TC relies heavily on China as a destination for that product.

Figure 6 summarizes how the U.S. tariff shock propagates upstream to China's foreign suppliers. In each country, we sort products by our downstream-tariff exposure measure and compare "high-exposure" products (top quartile) to "low-exposure" products (bottom quartile). The plotted points are event-study ATTs by year, with vertical bars denoting confidence intervals. Panel (a) shows a clear demand-side transmission to third-country exports. For Mexico, Vietnam, and India, the estimates are close to zero in the pre-trade-war period. After additional tariffs are imposed, export growth in highly exposed products falls relative to low-exposure products. The decline is most pronounced in 2019, when also China's exports to the United States contracted most sharply. Columns (1), (4), and (7) of Table A.11 present point estimates of these adjustments. For the average product, export declines for exposed products are notably larger in Vietnam than in Mexico and India, which is with Vietnam's greater orientation towards the Chinese market.[39]

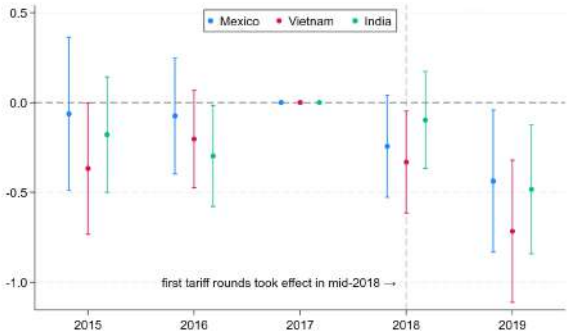
[37] Within the set of high-importance HS4 inputs (those most relevant for China's exports to the United States), Vietnam, India, and Mexico rank among China's largest sources: average imports per product equal 100,673, 47,944, and 30,149 (USD thousands), respectively. For the low-importance group, the corresponding means are markedly smaller at 39,773 for Vietnam, 13,402 for India, and 16,272 for Mexico, indicating that these suppliers are disproportionately represented in the inputs tied to China's U.S.-oriented export production. These three countries also have sufficiently large samples to infer input–output linkages (Table B.1) and exhibit relatively persistent mappings over time (Figure A.10).

[38] We compute $\bar{s}_{j,CN}$ as the average share of the TC's exports of product j shipped to China over 2015–2017 using CEPII BACI.

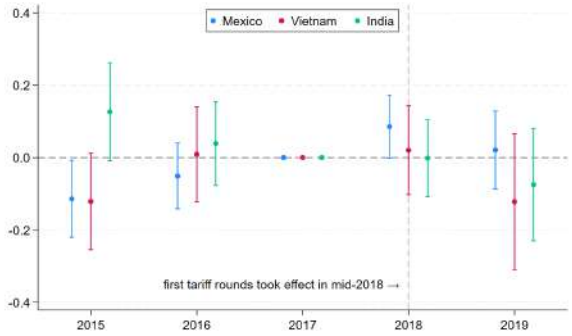
[39] The Chinese market share in total exports is 2.4 percent for Mexico, 14.7 percent for Vietnam, and 4.4 percent for India, based on CEPII Baci data for 2017.

Figure 6: Input suppliers' exports and upstream imports during the 2018-19 trade war

(a) Exports to China



(b) Export-related imports



Note: The figure reports event-study difference-in-differences estimates based on the extended two-way fixed effects estimator of Wooldridge (2023). Panel (a) plots the effects on third-country exports to China; Panel (b) plots the effects on the same countries' export-related imports. In each panel, products are split into a treated group (top quartile of our exposure measure) and a control group (bottom quartile). Exposure is constructed using the BFA mapping to link products to China's U.S.-tariffed export goods and is defined separately for outputs in Panel (a) and inputs in Panel (b). Coefficients are normalized to zero in the base period; bars denote 95% confidence intervals. See the main text for the construction of the exposure measures.

To examine whether these effects propagate further upstream, we use third-country input-output linkages to construct a measure of U.S. tariff exposure for input k, following an approach analogous to that used for Chinese imports. Specifically, we interact the jk-specific BFA score with a binary indicator that captures whether product j is exposed to U.S. tariffs.[40] Before aggregating over j, we weigh this term by the average share of the third country's exports of product j destined for China in the three years preceding the trade war. Formally, the exposure measure is defined as

$$\text{Exp}_{\tau(\text{US})\to k(\text{TC})} \equiv \sum_j score_{jk}^{\text{BFA}} \times (\text{Exp}_{\tau(\text{US})\to j} > 0) \times \bar{s}_{j,CN},$$

and scale this measure by 100 for ease of exposition. Intuitively, $\text{Exp}_{\tau(\text{US})\to k(\text{TC})}$ is higher when input k is predicted (by our mapping) to be used intensively in producing outputs j that are tariff-hit in the U.S., and when those outputs are disproportionately exported from the third country to China. Thus, it summarizes how strongly the third country's exports of input k are indirectly exposed to U.S. tariffs through downstream production chains.

Panel (b) of Figure 6 indicates to what extent suppliers' own imports contract in exposed products as their sales to China decline. Point estimates are generally negative around 2018-2019, which is the intuitive direction if reduced export demand leads firms to scale back

[40] We use post-processed BFA scores from 2017 for Mexico and India and from 2018 for Vietnam.

intermediate input purchases. At the same time, the confidence intervals are wider, so the effects are less precisely estimated. This is plausible because they operate one step further upstream, which gives scope to delayed (and potentially further margins of) adjustments that cannot be captured by our mapping. Columns (2)-(3), (5)-(6), and (8)-(9) of Table A.11 confirms these results, even if the period extends to later years. Noticeably, the less persistent decline of input-demand in Vietnam is consistent with recent studies emphasizing the rerouting of Chinese trade (e.g., Iyoha et al., 2024).

Altogether, our empirical evidence from China's WTO entry and the 2018-2019 trade war shows that the economically relevant adjustments to these shocks operate through product-level input-output linkages, which is why granular mappings matter. They allow us to take a shock defined on outputs — e.g., tariffs, demand shifts, or export booms — and translate it into predicted pressure on specific upstream inputs, and then into exposure for the countries that supply those inputs. The same mapping is useful well beyond trade wars. Whenever shocks are narrowly targeted but production relies on many inputs — e.g., sanctions, supply disruptions, industrial policy, environmental regulation, or technology shocks — a high-resolution I-O structure allows researchers to locate where the shock bites and to trace how it propagates through global value chains.

6 Conclusion

Modern trade and industrial policy are implemented at narrow product lines, but empirical work typically measures input-output relationships for broad sectors. This paper shows how to recover a product-level input-output mappings from widely available customs microdata. The approach exploits a simple footprint of production in firms' trade bundles: exporters of a given output disproportionately import the inputs used to make it, and importers of a given input disproportionately export the outputs that use it. We capture this two-way pattern with a refined association-rule metric, Bidirectional Frequency Association, and apply transparent economic filters to reduce chance co-occurrences and exclude implausible inputs. Implemented on the universe of Chinese importers and exporters in 2004, the procedure delivers a sparse, interpretable HS4-by-HS4 mapping that is well suited for tracing cross-border supply chains.

The resulting mapping has three practical advantages. First, it is *granular*: it operates at the product level rather than at broad sectors, so it can be used with the tariff lines, export-control categories, and narrowly defined technologies that policymakers and researchers care about. Second, it is *replicable and scalable*: it relies on harmonized customs records with minimal data requirements, and we show that the inferred structures evolve smoothly over

time and cluster sensibly across countries when implemented on multi-country shipment data. Third, it is *empirically validated*: the inferred linkages line up with independent benchmarks once aggregated to the industry level, and they recover a basic supply-chain implication in the data: downstream export growth is accompanied by stronger growth in the imports of mapped upstream inputs.

Our paper illustrates how product-level linkages sharpen inference about propagation. Around China's WTO entry, faster export growth is matched by faster import growth in the mapped upstream inputs: a 1 percentage point increase in export growth is associated with a 0.21-0.49 percentage point increase in input-import growth. The foreign supply response is likewise concentrated in the input lines most tied to China's expanding exports, and it is strongest for countries with revealed comparative advantage in those inputs. In the 2018-2019 U.S.-China trade war, the collapse in tariffed exports is mirrored upstream in weaker Chinese imports of the most exposed inputs and in reduced shipments to China by specialized third-country suppliers. Therefore, when shocks hit product lines, the action runs through product lines, and sector averages can wash out the relevant margins.

More broadly, a product-level input-output mapping turns targeted policy into measurable exposure. Tariffs, export controls, and sanctions do not stop at the listed outputs; they work through a narrow set of upstream inputs and can shift incidence onto suppliers and countries outside the policy's line of sight. Mapping outputs into inputs therefore pinpoints where substitution is limited, where bottlenecks are likely, and which upstream nodes concentrate spillovers, precisely the information needed for resilience and de-risking agendas that hinge on a small set of choke-point intermediates rather than broad sectors. Two extensions would strengthen this toolkit further: combining customs-based inference with complementary data to capture essential inputs sourced domestically, and more systematically identifying and down-weighting carry-along trade so inferred links better reflect production. More generally, routine trade records contain far more usable input-output structure than has to date been exploited. Extracting it in a transparent, validated way can improve both empirical work on propagation and the design of trade and industrial policy at the product margins where it is applied.

References

Agrawal, R., Imieliński, T., & Swami, A. (1993). Mining association rules between sets of items in large databases. In *Proceedings of the 1993 ACM SIGMOD International Conference on Management of Data*, SIGMOD '93 (pp. 207–216). New York, NY, USA: Association for Computing Machinery.

Alfaro, L., Conconi, P., Fadinger, H., & Newman, A. F. (2016). Do Prices Determine Vertical Integration? *The Review of Economic Studies*, 83(3), 855–888.

Alfaro, L., s, P. A., Chor, D., & Conconi, P. (2019). Internalizing Global Value Chains: A Firm-Level Analysis. *Journal of Political Economy*, 127(2), 508–559.

Amiti, M. & Freund, C. (2010). The Anatomy of China's Export Growth. In *China's Growing Role in World Trade*, NBER Chapters (pp. 35–56). National Bureau of Economic Research, Inc.

Antràs, P. & Chor, D. (2021). *Global Value Chains*. NBER Working Papers 28549, National Bureau of Economic Research, Inc.

Autor, D. H., Dorn, D., & Hanson, G. H. (2013). The china syndrome: Local labor market effects of import competition in the united states. *American Economic Review*, 103(6), 2121–2168.

Autor, D. H., Dorn, D., & Hanson, G. H. (2016). The china shock: Learning from labor-market adjustment to large changes in trade. *Annual Review of Economics*, 8(1), 205–240.

Barrot, J.-N. & Sauvagnat, J. (2016). Input Specificity and the Propagation of Idiosyncratic Shocks in Production Networks. *The Quarterly Journal of Economics*, 131(3), 1543–1592.

Berlingieri, G., Pisch, F., & Steinwender, C. (2021). Organizing global supply chains: Input-output linkages and vertical integration. *Journal of the European Economic Association*, 19(3), 1816–1852.

Bernard, A. B., Blanchard, E. J., Van Beveren, I., & Vandenbussche, H. (2019). Carry-Along Trade. *The Review of Economic Studies*, 86(2), 526–563.

Bernard, A. B. & Jensen, J. B. (1995). Exporters, jobs, and wages in u.s. manufacturing: 1976-1987. *Brookings Papers on Economic Activity*, 26(1995 Micr), 67–119.

Bernard, A. B., Jensen, J. B., Redding, S. J., & Schott, P. K. (2007). Firms in International Trade. *Journal of Economic Perspectives*, 21(3), 105–130.

Bernard, A. B., Redding, S. J., & Schott, P. K. (2011). Multiproduct Firms and Trade Liberalization. *The Quarterly Journal of Economics*, 126(3), 1271–1318.

Boehm, C., Flaaen, A. B., Pandalai-Nayar, N., & Schlupp, J. (2025). *The Local-Area Incidence of Exporting*. Working Paper 34508, National Bureau of Economic Research.

Boehm, C. E., Flaaen, A., & Pandalai-Nayar, N. (2019). Input Linkages and the Transmission of Shocks: Firm-Level Evidence from the 2011 Tōhoku Earthquake. *The Review of Economics and Statistics*, 101(1), 60–75.

Boehm, J., Dhingra, S., & Morrow, J. (2022). The comparative advantage of firms. *Journal of Political Economy*, 130(12), 3025–3100.

Bonadio, B., Huo, Z., Levchenko, A. A., & Pandalai-Nayar, N. (2021). Global Supply Chains in the Pandemic. *Journal of International Economics*, 133(C).

Borin, A., Conteduca, F., Leone, F., Mancini, M., & Zoi, P. (2025). *How Global Are Local Value Chains?* CESifo Working Paper Series 12271, CESifo.

Brandt, L., Biesebroeck, J. V., Wang, L., & Zhang, Y. (2017). Wto accession and performance of chinese manufacturing firms. *American Economic Review*, 107(9), 2784–2820.

Cajal-Grossi, J., Macchiavello, R., & Noguera, G. (2023). Buyers' Sourcing Strategies and Suppliers' Markups in Bangladeshi Garments. *The Quarterly Journal of Economics*, 138(4), 2391–2450.

Carvalho, V. M., Nirei, M., Saito, Y. U., & Tahbaz-Salehi, A. (2020). Supply Chain Disruptions: Evidence from the Great East Japan Earthquake*. *The Quarterly Journal of Economics*, 136(2), 1255–1321.

Chor, D., Manova, K., & Yu, Z. (2021). Growing like China: Firm performance and global production line position. *Journal of International Economics*, 130(C).

Ciani, A. & Mau, K. (2023). Delivery Times in International Competition: An Empirical Investigation. *Journal of Development Economics*, 161(C), None.

Cigna, S., Meinen, P., Schulte, P., & Steinhoff, N. (2022). The Impact of US Tariffs against China on US imports: Evidence for Trade Diversion? *Economic Inquiry*, 60(1), 162–173.

Cox, L. (2025). The Long-Term Impact of Steel Tariffs on U.S. Manufacturing. *mimeo*.

Dauth, W., Findeisen, S., & Suedekum, J. (2014). The rise of the east and the far east: German labor markets and trade integration. *Journal of the European Economic Association*, 12(6), 1643–1675.

De Backer, K. & Miroudot, S. (2014). *Mapping Global Value Chains*. Working Paper Series 1677, European Central Bank.

Diem, C., Borsos, A., Reisch, T., Kertész, J., & Thurner, S. (2024). Estimating the loss of economic predictability from aggregating firm-level production networks. *PNAS Nexus*, 3(3), pgae064.

Eckel, C., Iacovone, L., Javorcik, B., & Neary, J. P. (2016). Testing the core-competency model of multi-product exporters. *Review of International Economics*, 24(4), 699–716.

Eckel, C. & Neary, J. P. (2010). Multi-Product Firms and Flexible Manufacturing in the Global Economy. *The Review of Economic Studies*, 77(1), 188–217.

Eckel, C. & Riezman, R. (2020). CATs and DOGs. *Journal of International Economics*, 126(C).

Erbahar, A. & Rebeyrol, V. (2023). Trade intermediation by producers. *Journal of International Economics*, 140(C).

Fetzer, T., Garg, P., Lambert, P., & Feld, B. (2024). *AI-Generated Production Networks: Measurement and Applications to Global Trade*. Technical report, mimeo.

Flaaen, A. & Pierce, J. (2024). Disentangling the effects of the 2018-2019 tariffs on a globally connected u.s. manufacturing sector. *The Review of Economics and Statistics*, (pp. 1–45).

Gaulier, G. & Zignago, S. (2010). *BACI: International Trade Database at the Product-Level. The 1994-2007 Version*. Working Papers 2010-23, CEPII.

Greenaway, D., Mahabir, A., & Milner, C. (2008). Has china displaced other asian countries' exports? *China Economic Review*, 19(2), 152–169.

Handley, K., Kamal, F., & Monarch, R. (2020). *Rising Import Tariffs, Falling Export Growth: When Modern Supply Chains Meet Old-Style Protectionism*. NBER Working Papers 26611, National Bureau of Economic Research, Inc.

Handley, K., Kamal, F., & Monarch, R. (2025). Rising import tariffs, falling exports: When modern supply chains meet old-style protectionism. *American Economic Journal: Applied Economics*, 17(1), 208–38.

Handley, K. & Limão, N. (2017). Policy uncertainty, trade, and welfare: Theory and evidence for china and the united states. *American Economic Review*, 107(9), 2731–2783.

He, C., Mau, K., & Xu, M. (2021). Trade Shocks and Firms Hiring Decisions: Evidence from Vacancy Postings of Chinese Firms in the Trade War. *Labour Economics*, 71(C).

Hidalgo, C. A., Hausmann, R., & Dasgupta, P. S. (2009). The building blocks of economic complexity. *Proceedings of the National Academy of Sciences of the United States of America*, 106(26), 10570–10575.

Inoue, H. & Todo, Y. (2025). Disruption risk evaluation on large-scale production network with establishments and products.

Irlacher, M. (2024). *Multiproduct Firms in International Trade*. Oxford Research Encyclopedia of Economics and Finance.

Iyoha, E., Malesky, E., Wen, J., Wu, S.-J., & Feng, B. (2024). *Exports in Disguise?: Trade Rerouting during the US-China Trade War*. Harvard Business School Boston, MA, USA.

Johnson, R. C. (2014). Five facts about value-added exports and implications for macroeconomics and trade research. *Journal of Economic Perspectives*, 28(2), 119–42.

Karbevska, L. & Hidalgo, C. A. (2025). Mapping global value chains at the product level. *EPJ Data Science*, 14(21).

Kaushik, D., Mishra, N., & Udgata, S. K. (2021). A Systematic Assessment of Numerical Association Rule Mining Methods. *SN Computer Science*, 2(5), 357.

Kosko, B. (1987). Adaptive bidirectional associative memories. *Applied Optics*, 26(23), 4947–4960.

Kugler, M. & Verhoogen, E. (2011). Prices, plant size, and product quality. *The Review of Economic Studies*, 79(1), 307–339.

Ma, H., Ning, J., & Xu, Y. (2025). *Reshaping Factory Asia: How Did the US-China Trade War Affect the Asia-Pacific Supply Chain?* Ssrn working paper, SSRN. Posted June 13, 2025.

Macedoni, L., Morrow, J., & Tyazhelnikov, V. (2024). *Firms in Product Space: Adoption, Growth and Competition*. CEPR Discussion Paper 18800, CEPR Press, Paris & London.

Manova, K. & Yu, Z. (2016). How firms export: Processing vs. ordinary trade with financial frictions. *Journal of International Economics*, 100(C), 120–137.

Manova, K. & Yu, Z. (2017). Multi-product firms and product quality. *Journal of International Economics*, 109(C), 116–137.

Manova, K. & Zhang, Z. (2012). Export Prices Across Firms and Destinations. *The Quarterly Journal of Economics*, 127(1), 379–436.

Martin, J., Méjean, I., & Parenti, M. (2025). Relationship stickiness, international trade, and economic uncertainty. *The Review of Economics and Statistics*, (pp. 1–15).

Mau, K. (2017). US Policy Spillover(?) — China's Accession to the WTO and Rising Exports to the EU. *European Economic Review*, 98(C), 169–188.

Mayer, T., Melitz, M. J., & Ottaviano, G. I. P. (2014). Market size, competition, and the product mix of exporters. *American Economic Review*, 104(2), 495–536.

Nicita, A. (2019). *Trade and Trade Diversion Effects of United States Tariffs on China*. Technical report, UNCTAD.

O'Clery, N., Radcliffe-Brown, B., Spencer, T., & Tarling-Hunter, D. (2025). Deciphering the global production network from cross-border firm transactions.

Pichler, A., Pangallo, M., del Rio-Chanona, R. M., Lafond, F., & Farmer, J. D. (2020). *In and out of lockdown: Propagation of supply and demand shocks in a dynamic input-output model*. INET Oxford Working Papers 2021-18, Institute for New Economic Thinking at the Oxford Martin School, University of Oxford.

Pierce, J. R. & Schott, P. K. (2012). A concordance between ten-digit u.s. harmonized system codes and sic/naics product classes and industries. *Journal of Economic and Social Measurement*, 37(1-2), 61–96.

Pierce, J. R. & Schott, P. K. (2016). The surprisingly swift decline of us manufacturing employment. *American Economic Review*, 106(7), 1632–1662.

Solanki, V. K. & Patel, K. R. (2015). *A Survey on Association Rule Mining*. Technical report, Fifth International Conference on Advanced Computing & Communication Technologies.

Sutter, K. M. (2025). *U.S. Export Controls and China: Advanced Semiconductors*. Technical Report R48642, Congressional Research Service. Publication date: September 19, 2025. Accessed December 15, 2025.

Telikani, A., Nadali, A., & Sulaiman, R. (2020). A Survey of Evolutionary Computation for Association Rule Mining. *Information Sciences*, 524, 318–352.

Timmer, M. P., Dietzenbacher, E., Los, B., Stehrer, R., & Vries, G. J. (2015). An Illustrated User Guide to the World Input–Output Database: the Case of Global Automotive Production. *Review of International Economics*, 23(3), 575–605.

Topalova, P. & Khandelwal, A. (2011). Trade Liberalization and Firm Productivity: The Case of India. *The Review of Economics and Statistics*, 93(3), 995–1009.

Utar, H. (2014). When the floodgates open: "northern" firms' response to removal of trade quotas on chinese goods. *American Economic Journal: Applied Economics*, 6(4), 226–250.

Utar, H. (2018). Workers beneath the Floodgates: Low-Wage Import Competition and Workers' Adjustment. *The Review of Economics and Statistics*, 100(4), 631–647.

Utar, H. & Torres-Ruiz, L. B. (2013). International competition and industrial evolution: Evidence from the impact of chinese competition on mexican maquiladoras. *Journal of Development Economics*, 105(C), 267–287.

Utar, H., Zurita, A. C., & Torres, L. (2025). *The US-China Trade War and the Relocation of Global Value Chains to Mexico*. RFBerlin Discussion Paper Series 2546, ROCKWOOL Foundation Berlin (RFBerlin).

Vandenbussche, H. & Viegelahn, C. (2018). Input reallocation within multi-product firms. *Journal of International Economics*, 114, 63–79.

Wooldridge, J. M. (2023). Simple approaches to nonlinear difference-in-differences with panel data. *The Econometrics Journal*, 26(3), 31–66.

Zhao, Y. & Bhowmick, S. S. (2003). Association Rule Mining: A Survey. *Scientific Research Publishing*.

A Online Appendix

Table A.1: Sectoral patterns of Chinese trade in 2004

			(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
				% of total		# firms per HS4			% of HS4	
HS sections		HS chapters	# HS4	export	import	exports	imports	avg. RCA	interm. inputs	non-cons.
I-III	Animal and vegetable products	01-15	72	0.35	1.66	69	52	0.78	52.68	53.38
IV	Food, beverages, tobacco	16-24	80	1.60	0.45	135	79	0.70	37.81	37.81
V	Mineral products	25-27	66	2.31	9.58	88	157	0.89	89.24	89.24
VI-VII	Chemical products, plastic and rubber	28-40	275	7.21	11.21	384	816	0.85	89.22	89.22
VIII	Raw hides and skins, leather products	41-43	22	2.12	0.93	666	479	1.85	65.00	65.00
IX-X	Wood products, paper and paperboard	44-49	68	1.52	2.18	697	927	1.09	73.88	73.88
XI-XII	Textile, clothing and apparel	50-67	199	17.18	4.10	923	705	2.73	58.58	60.33
XIII-XIV	Stone, plaster, cement, ceram.; pearls, ...	68-71	102	2.35	1.00	412	231	1.21	78.21	78.96
XV	Metal and metal products	72-83	195	7.03	7.24	472	666	1.34	79.94	90.48
XVI-XVII	Machinery and electronics	84-85	132	45.23	36.29	941	2,018	1.11	32.78	87.79
XVIII	Vehicles and other transport equipment	86-89	37	3.73	2.16	322	147	0.98	20.11	83.18
XIX	Precision instr. and measurement app.	90-92	70	3.48	6.53	418	870	1.81	31.14	71.94
XX	Other manufactured goods	93-97	43	5.89	0.32	1,160	743	2.67	22.17	42.63

Note: Calculations based on 2004 firm level trade data from China Customs, after selecting relevant trade modes described in the text and aggregating products to 4-digit HS product lines (HS4). Sample features 66,688 firms and 1,361 unique HS4 product. RCA calculated with data from CEPII Baci version 202401b (HS02). Intermediate inputs and non-consumption goods are inferred from BEC Rev. 5 correspondence tables, available at the UNSD.

Figure A.1: Workflow for inferring product-level input–output linkages

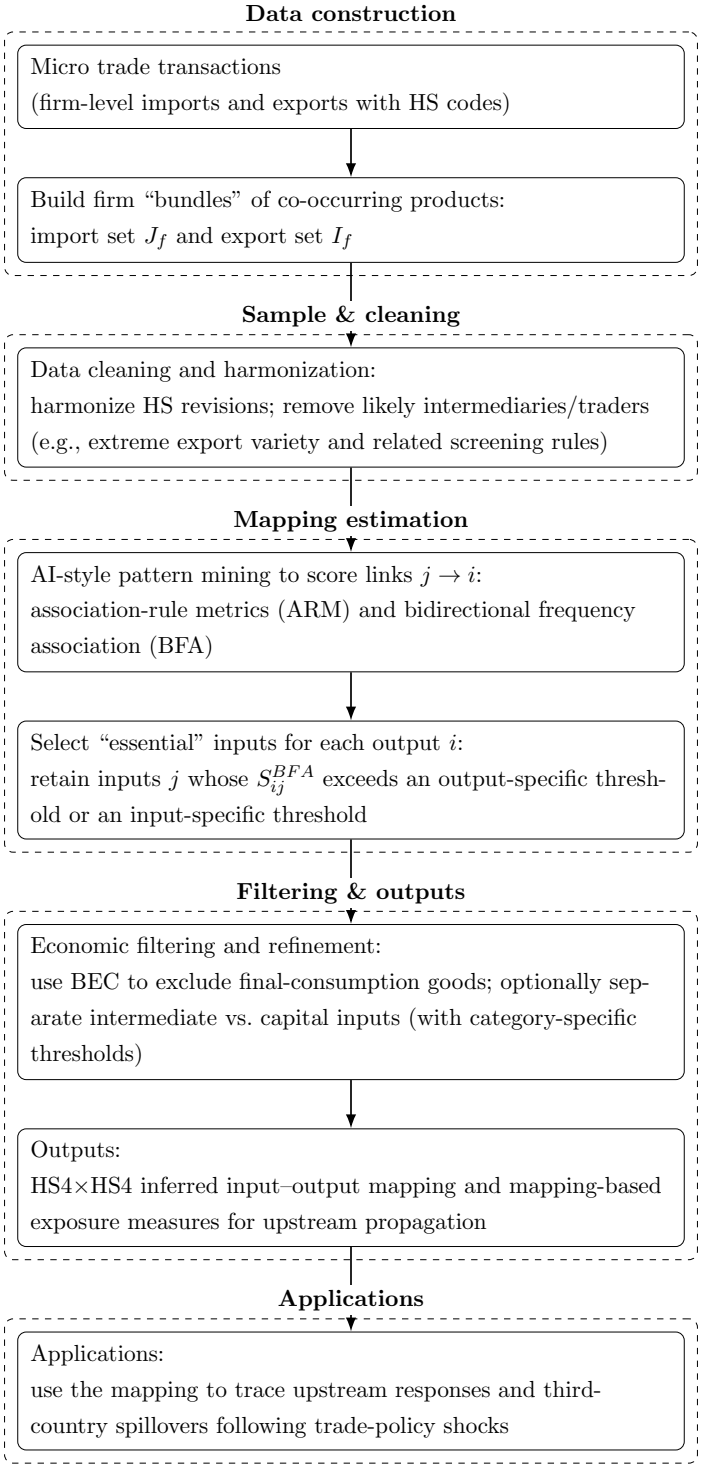


Figure A.2: BFA-based mappings examples

HS4: 8471	Automatic data processing machines and units thereof (Computers)	Share in exports (2004): 7.85%	# HS4 inputs: 54	
Top HS4 Inputs			BFA score	AIPNET
1	8473	Machinery; parts and accessories (other than covers, carrying cases and the like) suitable for use solely or principally with machines of headings 84.70 to 84.72	0.121	✓
2	8542	Electronic integrated circuits	0.089	✓
3	8541	Diodes, transistors, similar semiconductor devices; including photovoltaic cells assembled or not in modules or panels, light-emitting diodes (LED), mounted piezo-electric crystals	0.063	✓
4	8532	Electrical capacitors; fixed, variable or adjustable (pre-set)	0.056	✓
5	8534	Circuits; printed	0.053	✓
6	8533	Electrical resistors (including rheostats and potentiometers), excluding heating resistors	0.052	✓
7	8504	Electric transformers, static converters (e.g. rectifiers) and inductors	0.049	✓
8	9013	Liquid crystal devices not constituting articles provided for more specifically in other headings; lasers, not laser diodes; other optical appliances and instruments n.e.c. in this chapter	0.047	
9	8517	Telephone sets, including telephones for cellular networks or for other wireless networks; other apparatus for the transmission or reception of voice, images or other data (including wired/wireless networks), excluding items of 8443, 8525, 8527, or 8528	0.042	
10	8524	Flat panel display modules, whether or not incorporating touch-sensitive screens	0.040	

HS4: 7601	Unwrought aluminum	Share in exports (2004): 0.36%	# HS4 inputs: 7	
Top HS4 Inputs			BFA score	AIPNET
1	7602	Aluminium; waste and scrap	0.065	✓
2	2818	Aluminium oxide (including artificial corundum); aluminium hydroxide	0.023	✓
3	7404	Copper; waste and scrap	0.007	
4	7902	Zinc; waste and scrap	0.006	
5	8454	Converters, ladles, ingot moulds and casting machines; of a kind used metallurgy or in metal foundries	0.003	✓
6	8429	Bulldozers, graders, levellers, scrapers, angledozers, mechanical shovels, excavators, shovel loaders, tamping machines and road rollers, self-propelled	0.001	
7	8474	Machinery for sorting, screening, separating, washing, crushing, grinding, mixing or kneading earth, stone, ores in solid form, shaping, moulding machinery for solid mineral fuels	0.001	
8	*0909*	*Seeds of anise, badian, fennel, coriander, cumin, caraway or juniper*	*0.000**	
9	*7501*	*Nickel mattes; nickel oxide sinters and other intermediate products of nickel metallurgy*	*0.000**	
10	*2606*	*Aluminium ores and concentrates*	*0.000**	✓

HS4: 4011	New pneumatic tyres, of rubber	Share in exports (2004): 0.32%	# HS4 inputs: 40	
Top HS4 Inputs			BFA score	AIPNET
1	4001	Natural rubber, balata, gutta-percha, guayule, chicle and similar gums; in primary forms or in plates, sheets or strip	0.092	✓
2	5902	Textile fabrics; tyrecord of high tenacity yarn of nylon or other polyamides polyesters or viscose rayon	0.040	
3	8714	Vehicles; parts and accessories of heading no. 8711 to 8713	0.019	✓
4	2503	Sulphur of all kinds; other than sublimed, precipitated and colloidal sulphur	0.018	
5	4002	Synthetic rubber and factice derived from oils, in primary forms or in plates, sheets or strip; mixtures of heading no. 4001 and 4002, in primary forms or in plates, sheets or strip	0.018	✓
6	2802	Sulphur; sublimed or precipitated, colloidal sulphur	0.016	
7	4013	Inner tubes, of rubber	0.014	
8	4005	Compounded rubber, unvulcanised, in primary forms or in plates, sheets or strip	0.009	✓
9	3823	Industrial monocarboxylic fatty acids; acid oils from refining; industrial fatty alcohols	0.007	
10	3812	Prepared rubber accelerators; compound plasticisers for rubber or plastics, n.e.c. or included; anti-oxidising preparations and other compound stabilisers for rubber or plastics	0.007	✓

HS4: 6204	Women's Suits, jackets, dresses, skirts, trousers, overalls, breeches and shorts (not knitted or crocheted)	Share in exports (2004): 1.09%	# HS4 inputs: 46	
Top HS4 Inputs			BFA score	AIPNET
1	9606	Buttons, press-fasteners, snap-fasteners and press-studs, button moulds and other parts of these articles; button blanks	0.345	
2	9607	Slide fasteners and parts thereof	0.294	
3	5807	Labels, badges and similar articles; of textile materials, in the piece, in strips or cut to shape or size, not embroidered	0.288	✓
4	5407	Woven fabrics of synthetic filament yarn, including woven fabrics obtained from materials of heading no. 5404	0.277	✓
5	5208	Woven fabrics of cotton, containing 85% or more by weight of cotton, weighing not more than 200 g/m2	0.263	✓
6	5209	Woven fabrics of cotton, containing 85% or more by weight of cotton, weighing more than 200g/m2	0.232	✓
7	5210	Woven fabrics of cotton, containing less than 85% by weight of cotton, mixed mainly or solely with man-made fibres, weighing not more than 200 g/m2	0.210	✓
8	5309	Woven fabrics of flax	0.175	✓
9	5603	Nonwovens; whether or not impregnated, coated, covered or laminated	0.174	✓
10	5806	Fabrics; narrow woven, other than goods of heading no. 5807; narrow fabrics consisting of warp without weft assembled by means of an adhesive (bolducs)	0.171	✓

HS4: 8609	Containers; (including containers for transport of fluids) specially designed and equipped for carriage by one or more modes of transport	Share in exports (2004): 0.67%	# HS4 inputs: 14	
Top HS4 Inputs			BFA score	AIPNET
1	4412	Plywood, veneered panels and similar laminated wood	0.041	✓
2	7208	Iron or non-alloy steel; flat-rolled products of a width of 600mm or more, hot-rolled, not clad, plated or coated	0.033	✓
3	3208	Paints, varnishes; (enamels and lacquers) based on synthetic polymers or chemically modified natural polymers, dispersed or dissolved in a non-aqueous medium	0.014	
4	7225	Alloy steel flat-rolled products, of a width 600mm or more	0.009	✓
5	3814	Organic composite solvents and thinners, not elsewhere specified or included; prepared paint or varnish removers	0.009	
6	7219	Stainless steel; flat-rolled products of width of 600mm or more	0.005	✓
7	8467	Tools; for working in the hand, pneumatic, hydraulic or with self-contained electric or non-electric motor	0.003	
8	7216	Iron or non-alloy steel, angles, shapes and sections	0.003	✓
9	8481	Taps, cocks, valves and similar appliances for pipes, boiler shells, tanks, vats or the like, including pressure-reducing valves and thermostatically controlled valves	0.003	✓
10	8515	Electric (electrically heated gas) soldering, brazing, welding machines and apparatus, capable or not of cutting, electric machines and apparatus for hot spraying of metals or sintered carbides	0.002	

Note: Figure displays top-10 inputs for outputs inferred with BFA from Chinese firm-level trade flows in 2004. BFA score indicates relative co-occurrence of imports and exports, as described in Section 3. AIPNET indicates whether same inputs were identified also by Fetzer et al. (2024).

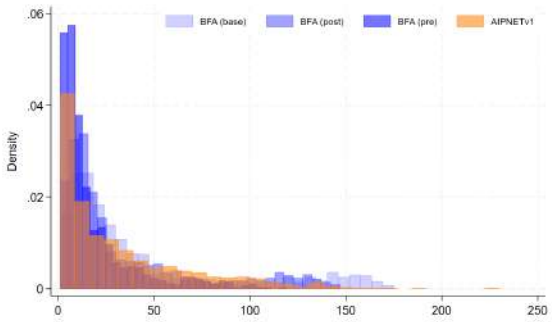
Table A.2: Summary statistics BFA-based mappings (HS4-level), China 2004

		Inputs per output (HS4 level)						
Method	count	avg.	sd	p5	p25	p50	p90	max
BFA baseline	1,218	47.4	50.8	5	13	25	144	197
BFA post	1,210	35.5	38.9	4	9	18	110	153
BFA pre- & post	1,168	17.7	20.8	2	5	10	68	121
		Outputs per input (HS4 level)						
Method	count	avg.	sd	p5	p25	p50	p90	max
BFA baseline	1,215	47.5	50.7	4	14	26	142	215
BFA post	935	46.0	49.3	4	13	25	135	213
BFA pre- & post	924	22.4	22.4	2	8	14	53	118

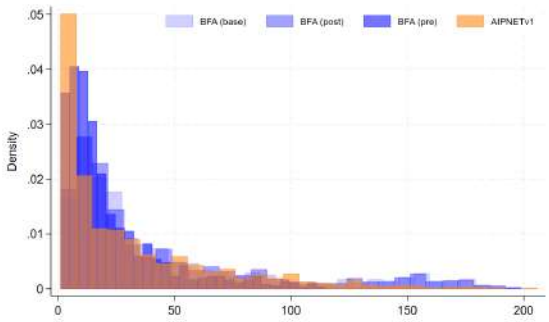
Note: Author's calculations based on Firm-level trade data from *China Customs* for the year 2004.

Figure A.3: Product complexity and input ubiquity in granular input-output mappings

(a) Inputs per output



(b) Outputs per input



Note: Figure compares HS4-level input-output linkages inferred with different BFA-based approaches with the LLM-based AIPNET by (Fetzer et al., 2024). Panel (a) displays distributions of the number of detected inputs per HS4 product. Panel (b) indicates the frequency distribution of outputs per input.

Table A.3: Output complexity and input ubiquity in BFA vs. LLM-based inference

	(1)	(2)	(3)	(4)	(5)	(6)
Dep. var.: AIPNET →	Output complexity (#*j*s per *i*)			Input ubiquity (#*i*s per *j*)		
BFA-mapping	baseline	post	pre & post	baseline	post	pre & post
Coefficient	0.361**	0.435**	0.714**	0.125**	0.244**	0.749**
	(0.019)	(0.024)	(0.045)	(0.023)	(0.023)	(0.045)
Observations	1,154	1,154	1,154	1,159	1,159	1,159
R-squared (adj.)	0.230	0.214	0.179	0.025	0.089	0.195

Note: Table displays coefficients obtained from complexity and ubiquity measures calculated from AIPNET on corresponding metrics obtained from BFA-based mappings. Standard errors in parentheses. Statistical significance: a $p < 0.1$, * $p < 0.05$, ** $p < 0.01$.

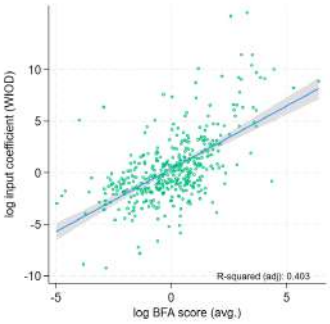
Table A.4: Correlation between BFA-based mappings and other Input-Output tables

Dep. var.: IO coefficients	(1)	(2)	(3)	(4)	(5)
Country	China 2004	China 2017	France 2006	USA 2007	
Source	WIOD-16	NBS SUT	Berlingieri et al. (2021)	BEA	AIPNET
Industry classifcation	ISIC Rev. 4	CIC Rev. 3.1	ISIC Rev. 2	NAICS	HS4 Rev.'02
Panel A: BFA mapping (baseline)					
BFA-score (ln)	1.557**	1.669**	1.150**	0.878**	0.213**
	(0.154)	(0.070)	(0.038)	(0.077)	(0.004)
Observations	457	3,871	3,638	3,743	50,618
R-squared(within)	0.321	0.171	0.255	0.047	0.069
Panel B: BFA mapping (post-processed)					
BFA-score (ln)	1.477**	1.776**	1.117**	0.935**	0.232**
	(0.148)	(0.074)	(0.041)	(0.080)	(0.005)
Observations	446	3,209	3,113	3,572	39,281
R-squared(within)	0.312	0.192	0.244	0.053	0.074
Panel C: BFA mapping (pre- and post-processed)					
BFA-score (ln)	1.224**	1.545**	0.931**	0.804**	0.166**
	(0.102)	(0.063)	(0.033)	(0.072)	(0.005)
Observations	398	2,381	2,355	2,665	19,384
R-squared(within)	0.403	0.233	0.321	0.061	0.063

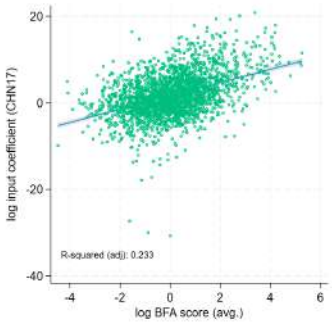
Note: All BFA-based mappings are inferred from Chinese firm-level transactions data, as described in Section 3. Input-output data obtained from sources indicated in column headings, where WIOD-16 reflects the 2016 release of the World Input Output Database described in Timmer et al. (2015); NBS SUT refers to Chinese Supply and Use Tables made available by the National Bureau of Statistics of China; French IO linkages reflect the linkages inferred and published by Berlingieri et al. (2021); while BEA reflects the 2007 Make-and-Use tables, published by the US Bureau of Economic Analysis. Column (5) summarizes results of estimating linear probability model regressing binary linkages in AIPNET on (log) BFA-scores. Standard errors in parentheses. Statistical significance: ^{a} $p < 0.1$, * $p < 0.05$, ** $p < 0.01$.

Figure A.4: Correlation between BFA-based mappings and other Input-Output tables

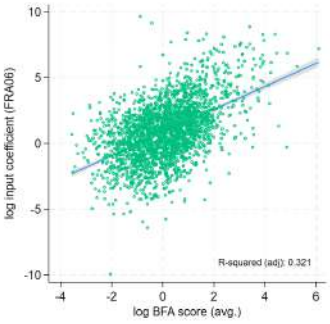
(a) China 2004 (WIOD-16)



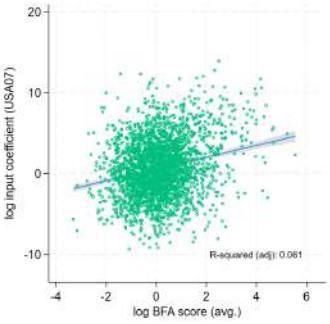
(b) China 2017 (NBS SUT)



(c) France 2006 (Berlingieri et al., 2021)

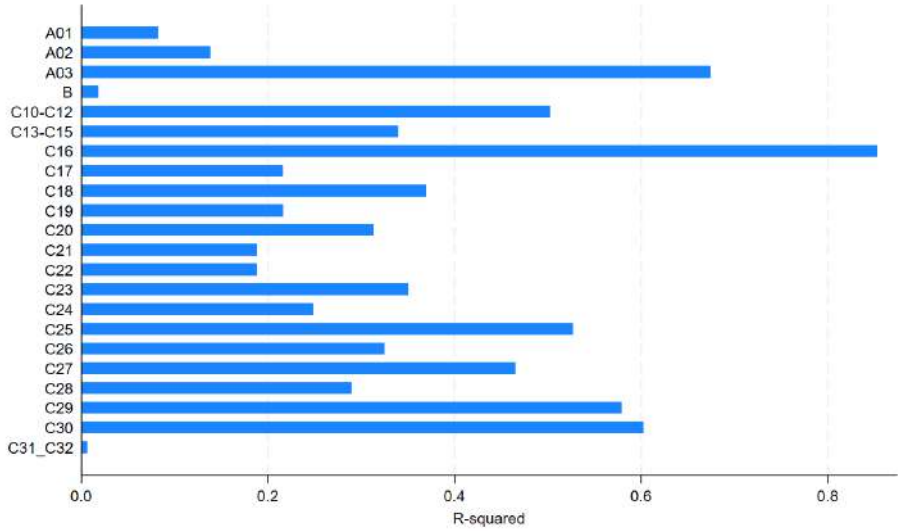


(d) USA 2007 (BEA)



Note: Figures display (centered) correlations between aggregated BFA-based IO linkages and combined coefficients from other IO tables. BFA mappings reflect linkages obtained after pre- and post-processing underlying firm-level transactions data (as described in Section 3).

Figure A.5: Sector-wise correlation between BFA-based mapping and WIOD-2016



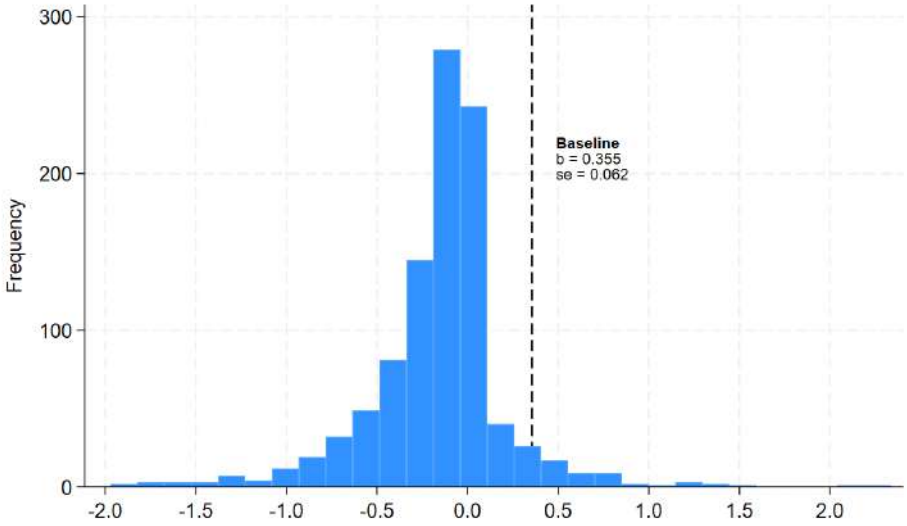
Note: Figure displays R-squares obtained for sector-wise log-linear regressions of WIOD IO coefficients for China on BFA-based mapping (with pre- and post-processing) inferred from Chinese firm-level transactions in 2004.

Table A.5: Conditional correlation between BFA and WIOD-2016, China 2004

	(1)	(2)
Dep. var.: IO coefficient	RCA output	RCA input
BFA-score (log)	1.568**	1.549**
	(0.229)	(0.186)
$\times RCA(output)$	-0.238^{a}	
	(0.138)	
$\times RCA(input)$		-0.252*
		(0.113)
Observations	398	398
R-squared (within)	0.413	0.412

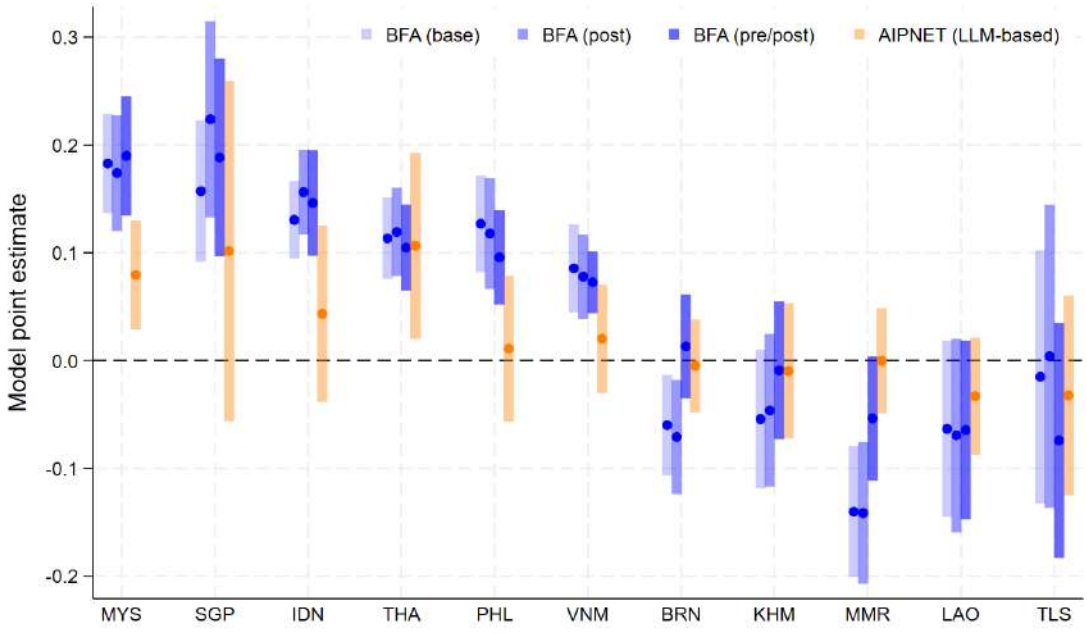
Note: Table displays estimated conditional correlation between between (log) BFA-scores and WIOD-2016 IO coefficients for China 2004. RCA refers to China's average revealed comparative advantage in respective output-producing and input-sourcing industries, as computed from CEPII Baci data for the year 2004. BFA-based mapping (with pre- and post-processing) inferred from Chinese firm-level transactions in 2004. Standard errors in parentheses. Statistical significance: ^{a} $p < 0.1$, * $p < 0.05$, ** $p < 0.01$

Figure A.6: Placebo estimates with random mappings, 1000 iterations



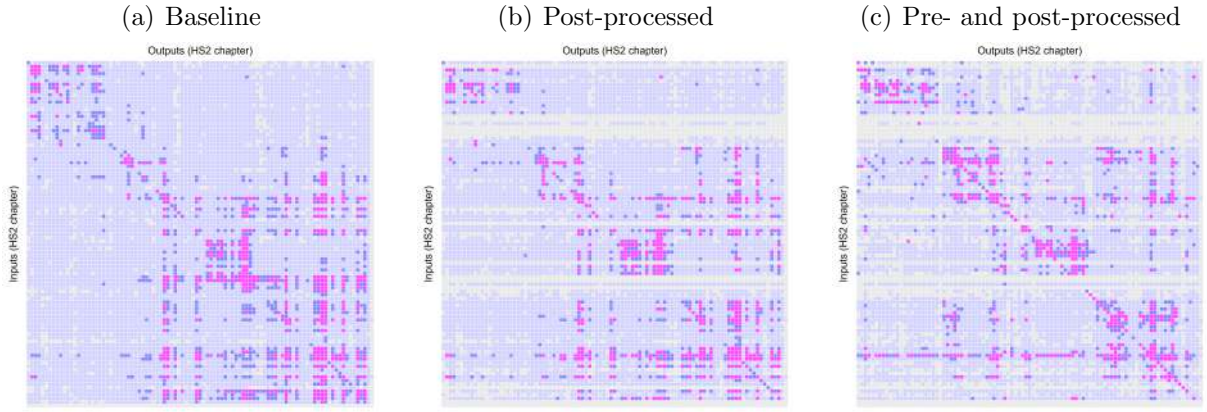
Note: Figure displays coefficients obtained from estimating Equation (4) with randomized BFA-scores across input-output pairs ij. Indicated baseline result corresponds to coefficient reported in Column (6) of Table 2, Panel A. In 79.2 percent of the iterations the placebo point estimate is $\hat{b} \leq 0$. In 94.1 percent of the iterations the placebo point estimate is $\hat{b} < 0.355$.

Figure A.7: Import response to downstream export growth in South-East Asia, 2003-2008



Note: Figure displays point estimates and 95 percent confidence intervals of the import response to downstream export growth. Blue point estimates and intervals reflect results from HS4-level BFA mappings inferred from Chinese firm-level transactions in 2004. Orange estimates and confidence intervals reflect the HS4-level LLM-based by Fetzer et al. (2024).

Figure A.8: Inferred input-output linkages with alternative implementations of BFA



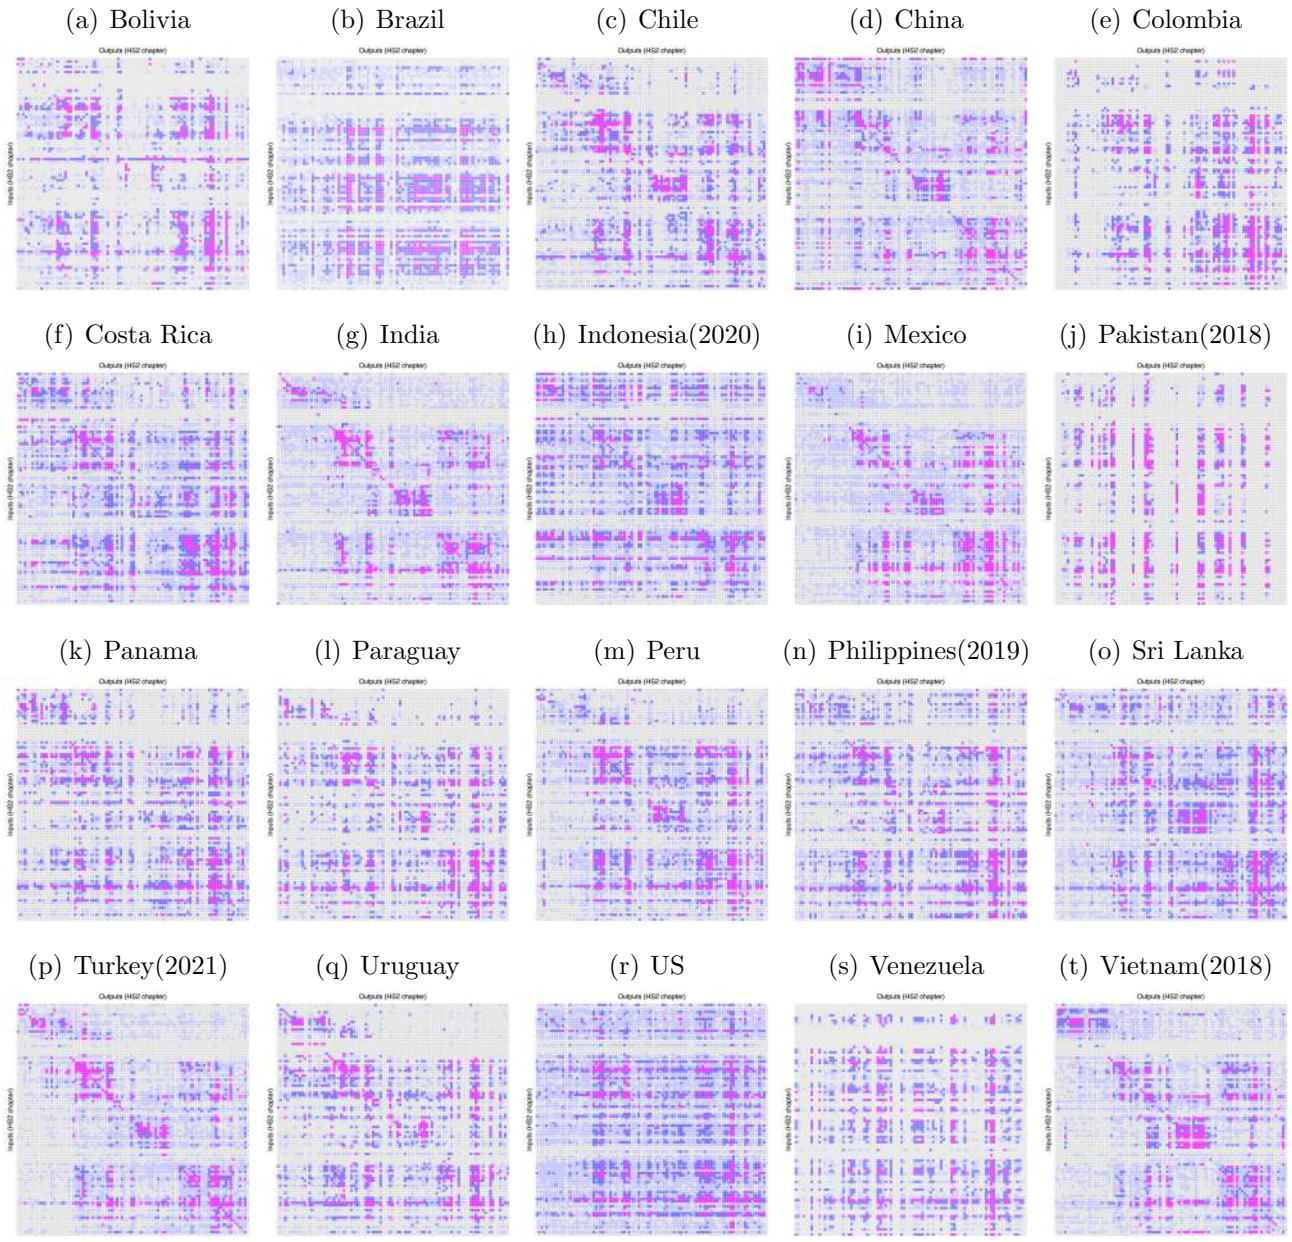
Note: Authors' calculations based on BFA applied to HS6-level firm transactions for China in 2013–2018. Panels (a)-(c) differ in the pre- and post-processing steps described in Section 3.3. Scores are aggregated to the HS2 chapter level. Color intensity ranges from zero (white) to the sample average (light blue) and the 95th percentile and above (pink).

Table A.6: Cross-country correlation of mapping results

	BO	BR	CL	CN	CO	CR	ID	IN	LK	MX	PA	PE	PH	PK	PY	TR	US	UY	VE	VN
BO																				
BR	0.0810																			
CL	0.4304	0.1435																		
CN	0.0768	-0.0471	-0.0189																	
CO	0.2165	0.0146	0.3300	-0.0027																
CR	0.2766	0.1513	0.2819	0.2846	0.0486															
ID	0.2329	0.0759	0.1770	0.3781	0.1503	0.3155														
IN	0.1612	-0.0930	0.1157	0.4749	0.0055	0.2865	0.3902													
LK	0.2433	0.0854	0.1862	0.3310	0.1324	0.3220	0.5461	0.3356												
MX	-0.0212	-0.1042	-0.1306	0.6118	-0.1539	0.2843	0.1629	0.3585	0.1122											
PA	0.3636	0.1592	0.3062	-0.0368	0.2715	0.1792	0.1624	0.0542	0.1449	-0.1206										
PE	0.2890	0.0892	0.3429	0.2369	0.1035	0.3463	0.3248	0.2590	0.2466	0.1978	0.1656									
PH	0.3652	0.1238	0.3438	0.0163	0.2417	0.1855	0.1993	0.0833	0.2020	-0.0957	0.1866	0.2041								
PK	0.2657	-0.0609	0.0688	-0.0301	0.1027	0.2044	0.1224	0.1956	0.1115	-0.0862	0.1462	0.1691	0.2026							
PY	0.4654	0.0868	0.3088	0.0570	0.2415	0.2022	0.2485	0.1146	0.2430	-0.0885	0.3788	0.2094	0.3524	0.1697						
TR	0.1782	-0.0719	0.0893	0.5175	0.0180	0.3432	0.3511	0.5240	0.2096	0.4063	0.0808	0.3388	0.1156	0.1491	0.1461					
US	0.0824	-0.1159	0.0680	0.1437	-0.0737	0.1379	0.0846	0.2818	0.0746	0.1338	0.0487	0.1507	-0.0592	0.0402	0.0117	0.2463				
UY	0.5145	0.1681	0.4364	0.0221	0.3310	0.1803	0.2645	0.1407	0.2269	-0.1340	0.3845	0.2017	0.4898	0.2210	0.4558	0.1692	0.0331			
VE	0.3745	0.3074	0.2919	-0.1687	0.4160	-0.0658	0.1570	-0.0784	0.0273	-0.2801	0.3063	0.0352	0.2958	0.1335	0.3074	-0.0438	-0.1565	0.4364		
VN	0.1665	0.0434	0.0447	0.4596	0.0751	0.2465	0.5305	0.3242	0.4785	0.3244	-0.0004	0.2057	0.0766	0.0502	0.1377	0.3036	0.1012	0.0859	0.0057	

Note: Table reports cross-economy correlations of inferred I–O linkages, identified by ISO two-digit codes, using 2017 or the closest available year. Linkages are inferred using BFA based on post-processed HS4-level firm transactions.

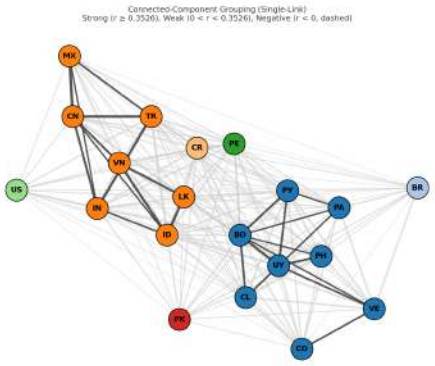
Figure A.9: Inferred I-O linkages across countries using BFA



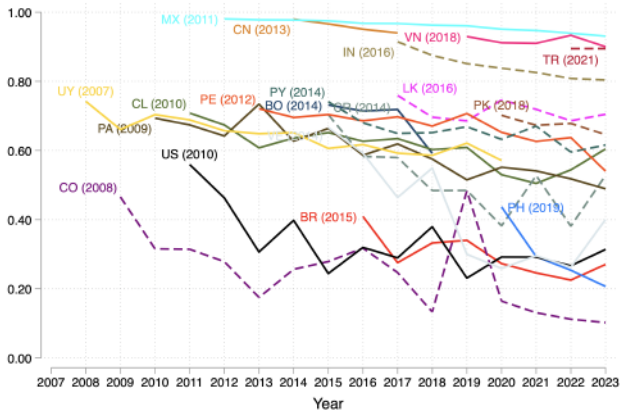
Note: Authors' calculations based on the BFA applied to HS4-level firm transactions for 20 economies, using the post-processed dataset. The reference year is 2017; for countries where 2017 data are unavailable, the closest alternative year is used (2020 for Indonesia, 2018 for Pakistan, 2019 for Philippines, 2021 for Turkey, and 2018 for Vietnam). Resulting scores have been aggregated to HS2-level chapter. Color intensity indicates scores ranging from scores equal to zero (white), the respective sample average (light blue), and to scores equal and above 95th percentile (pink).

Figure A.10: Similarity of HS4-level input-output mappings across countries and over time

(a) Across countries

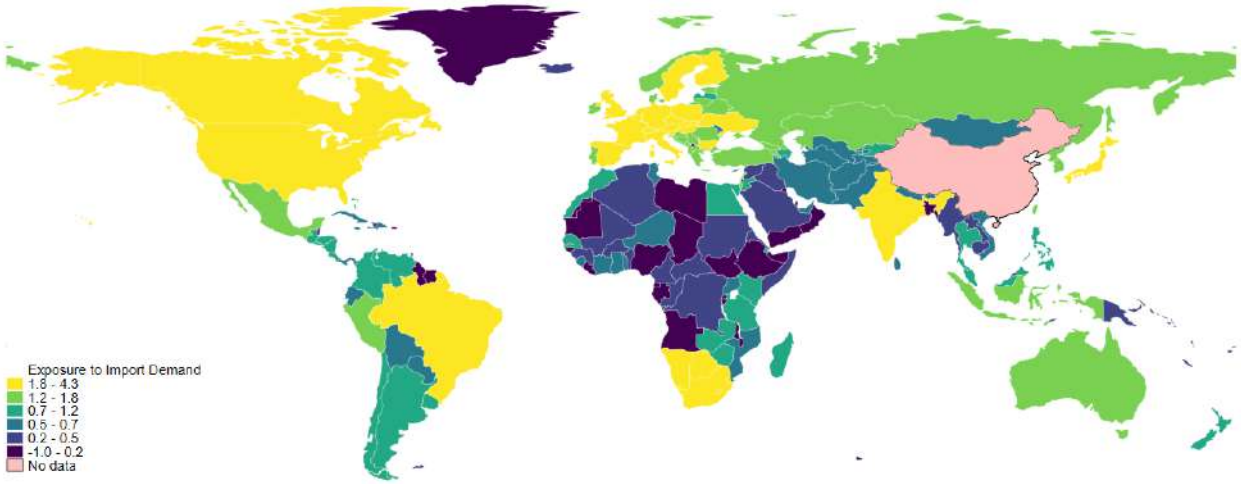


(b) Overtime



Note: Figure displays similarity of IO linkages (inferred from HS4-level firm transactions in 2017 (2020 for Indonesia, 2018 for Pakistan, 2019 for the Philippines, 2021 for Turkey, and 2018 for Vietnam), using BFA and subsequent post-processing. Panel (a) displays clusters of economies based on the cross-country correlation of inferred linkages. Economies are assigned to the same cluster (indicated by the same color) if their pairwise correlation with any economy in the group satisfies $r \geq 0.3526$, equal to twice the mean correlation across countries. Cross-country correlation values are reported in Table A.6. Light grey lines indicate weak correlations($0 < r < 0.3526$), while dashed lines represent negative correlations. Panel (b) plots the time-series correlation of I-O linkages for each economy relative to its earliest available year in the dataset (reported in parentheses).

Figure A.11: Exposure to the China's export-driven change in import demand after WTO entry



Note: Figure depicts countries' average exposure to China's export-driven import demand after WTO entry, based on observed change in product-level export growth and binary indicator for countries' $RCA > 1$ in the respective input goods.

Table A.7: Downstream export growth and upstream imports after China's WTO entry

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent variable:	China's imports (upstream)					
Measured as	Δ% growth rate		Levels (log-lin)		Levels (PPML)	
Processing of BFA-mapping	post	pre-/post	post	pre-/post	post	pre-/post
Downstream export growth (top-25%)	0.369**	0.366**	0.213**	0.237**	0.488**	0.451**
×WTO	(0.096)	(0.080)	(0.065)	(0.062)	(0.116)	(0.120)
HS4 fixed effects			✓	✓	✓	✓
Year fixed effects			✓	✓	✓	✓
Control variables			✓	✓	✓	✓
Cluster (HS4 products)			679	689	679	689
Observations	1,084	1,084	8,146	8,264	8,148	8,268

Note: Table reports estimated coefficients obtained from regressing Chinese imports on downstream export growth, using alternative BFA-based input-output mappings at the HS4 level. Columns (1)-(2) adopt a cross sectional regression of the *change* in product level import growth rates in a diff-in-diff setting. The treated imports correspond to inputs for which the change in export growth of related downstream goods resides in the upper quartile of the distribution. The control group consists of inputs with downstream export growth rates in the bottom quartile. The same treatment is used in the remaining columns, where the dependent variable is measured in (log) levels (columns (3)-(4)) and in absolute levels (columns (5)-(6)). Fixed effects and control variables are included as indicated at the bottom of the table. Standard errors in parentheses. Statistical significance: ^{a} $p < 0.1$, * $p < 0.05$, ** $p < 0.01$

Table A.8: Chinese demand for inputs and supplier RCA

	(1)	(2)	(3)	(4)
Dependent variable:	(log) Exports to China			
BFA-mapping	Post-processed		Pre/post-processed	
Downstream growth (China) ×WTO	0.104^{a}	0.029	0.218**	0.140*
	(0.062)	(0.063)	(0.060)	(0.062)
× RCA$_n(j) > 1$		0.227**		0.234**
		(0.042)		(0.041)
(log) Chinese tariff$_{jt}$	-0.103	-0.104	-0.175*	-0.175*
	(0.083)	(0.083)	(0.083)	(0.083)
HS4 fixed effects	✓	✓	✓	✓
Year fixed effects	✓	✓	✓	✓
Control variables	✓	✓	✓	✓
Cluster (HS4 products)	678	678	687	687
Observations	207,868	207,868	215,547	215,547

Note: Table reports results of HS4-country pair exports to China in response to Chinese downstream export growth. The interaction term is a HS4-country specific dummy variable, which indicates whether exporter n has a comparative advantage in exporting goods inferred as inputs to produce Chinese exports. All specifications include exporter-HS4, exporter-year, and HS2-year fixed effects. Standard errors in parentheses adjusted for clustering at the HS4 level. Statistical significance: ^{a} $p < 0.1$, * $p < 0.05$, ** $p < 0.01$

Table A.9: China's trade adjustments during the 2018–2019 trade war

	(1)	(2)	(3)	(4)
Frequency and aggregation:	Monthly (HS6 level)	Annual (HS4 level)		
Dep. var. (in logs) and trade partner:	Exports to US	Exports to US	Exports to World	Imports fr. World
US tariff $(k,t) = 1$	-0.290**			
	(0.014)			
US tariff $(k) = 1 \times$ (year = 2019)		-0.218**	-0.054	
		(0.080)	(0.045)	
US tariff exposure of inputs $\times$ (year = 2019)				–0.028*
				(0.011)
Fixed effects	HS6×month	HS4, year	HS4, year	HS4, year
Lin. trends (product-level)	✓	✓	✓	✓
Other controls				✓
Cluster	4,336	1,180	1,224	1,139
Observations	218,142	5,793	6,119	5,694

Note: The table reports regressions describing China's export and import adjustments during the 2018–2019 trade war. Column (1) uses monthly HS6 shipment data and the U.S. tariff lists compiled by He et al. (2021) for 2015–2019; the coefficient on US tariff(k,t) compares Chinese exports to the United States in tariffed versus non-tariffed HS6 products. Columns (2)–(4) use annual HS4 trade flows from CEPII BACI (2025 release `BACI_HS12_V202501`) over 2015–2019. Columns (2) and (3) estimate the differential change in 2019 for HS4 products subject to U.S. tariffs, for exports to the United States and for China's worldwide exports, respectively. Column (4) examines the upstream margin: it relates China's imports to the world to a measure of input exposure to U.S. tariffs, constructed by weighting U.S.-tariffed downstream export products using our post-processed BFA scores (see main text). All specifications include the fixed effects and product-specific linear trends indicated in the table; column (4) additionally includes the controls noted. Standard errors (in parentheses) are clustered at the product level. Statistical significance: ^{a} $p < 0.1$, * $p < 0.05$, ** $p < 0.01$.

Table A.10: China's imports by product importance and source country

	High-importance HS4 products (n = 569)				Low-importance HS4 products (n = 643)			
Import fr. (USD, thousands)	(1)Mean	(2)Std	(3)Min	(4)Max	(5)Mean	(6)Std	(7)Min	(8)Max
Vietnam	100,673	712,315	0	10,596,108	39,773	243,572	0	3,246,713
Philippines	79,689	860,878	0	18,746,698	14,186	154,789	0	3,481,374
India	47,944	273,072	0	3,817,426	13,402	111,301	0	2,519,331
Indonesia	37,535	173,205	0	2,493,964	67,867	486,056	0	6,415,805
Chile	36,613	802,752	0	19,139,114	50,413	710,743	0	17,124,736
Mexico	30,149	209,341	0	4,160,866	16,272	175,337	0	3,459,026
Brazil	13,388	78,464	0	1,161,945	177,498	2,325,127	0	50,539,688
Turkey	8,058	99,510	0	2,339,546	6,002	40,950	0	665,357
Pakistan	7,625	123,215	0	2,911,604	3,205	26,353	0	498,518
Peru	5,603	109,953	0	2,610,695	38,518	659,753	0	16,363,082
Venezuela	4,258	94,148	0	2,240,695	26,652	638,407	0	16,164,527
Costa Rica	1,743	21,562	0	364,427	349	3,293	0	52,619
Sri Lanka	890	6,603	0	127,743	990	12,375	0	300,757
Uruguay	384	5,096	0	103,032	8,401	106,358	0	1,582,396
Bolivia	322	5,457	0	126,363	1,968	25,907	0	498,196
Columbia	149	1,001	0	13,655	13,294	293,527	0	7,421,271
Panama	115	2,075	0	48,479	545	10,329	0	258,473
Paraguay	61	646	0	12,348	81	1,426	0	29,349

Note: Table reports China's imports during 2015–2017 from countries covered in the *Panjiva* dataset. HS4 products are classified into two groups based on their importance for China's exports to the United States relative to the average product. Product importance is measured by the average U.S. market share of a product's output, weighted by post-processed BFA scores. Columns (1)–(4) report summary statistics (mean, standard deviation, minimum, and maximum) for imports of high-importance HS4 products from each country, while columns (5)–(8) report the corresponding statistics for low-importance products. Data on imports and U.S. market shares are sourced from CEPII Baci.

Table A.11: Upstream propagation of U.S. trade-war tariffs to third-country trade with China

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Country:	Mexico			Vietnam			India		
Period of sample:	2015–2019		2015–2023	2015–2019		2015–2023	2015–2019		2015–2023
Dep. var. (in logs) and trade partner:	Exports to China	Imports fr. World	Imports fr. World	Exports to China	Imports fr. World	Imports fr. World	Exports to China	Imports fr. World	Imports fr. World
Output exposure to US tariff	-1.298*			-1.239**			-1.015*		
× (year≥2019)	(0.583)			(0.392)			(0.419)		
Input exposure to US tariff		-0.181	-0.169[a]		-0.076	-0.027		-0.055	-0.129*
× (year≥2019)		(0.111)	(0.100)		(0.049)	(0.049)		(0.067)	(0.065)
Fixed effects	HS4, year	HS4, year	HS4, year	HS4, year	HS4, year	HS4, year	HS4, year	HS4, year	HS4, year
Lin. trends (product-level)	✓	✓	✓	✓	✓	✓	✓	✓	✓
Other controls	✓	✓	✓	✓	✓	✓	✓	✓	✓
Control for Covid-19			✓			✓			✓
Cluster	768	1,214	1,214	983	1,210	1,210	1,063	1,211	1,211
Observations	3,369	6,056	10,895	4,419	6,036	10,847	4,911	6,041	10,863

Note: Table presents estimates of how additional U.S. tariffs on Chinese goods transmit upstream to third-country supply of intermediate inputs to China. The sample covers annual HS4-level trade flows from CEPII BACI (2025 release `BACI_HS12_V202501`). The dependent variables, trade partner and sample periods are indicated in column headings. Explanatory variable *Output exposure* captures exposure of the country's exports to China through downstream products j that face U.S. tariffs. *Input exposure* captures exposure of the country's upstream inputs k through their predicted use (BFA scores) in producing tariff-hit outputs; the measure is weighted by the country's pre-trade-war export share of product j shipped to China and is scaled by 100 (see main text). All specifications, include HS4 and year fixed effects and HS4-specific linear trends. "Other controls" include controls for U.S. tariffs on China and China's retaliatory tariffs on U.S. goods. For regressions spanning 2015-2023, we additionally include indicator variables for 2020-2022 to absorb COVID-19 disruptions. Standard errors in parentheses are clustered at the HS4 level. Statistical significance: ^{a} $p < 0.1$, * $p < 0.05$, ** $p < 0.01$.

B Details on Dataset (for online publication)

The *Panjiva* dataset covers trade flows for Bolivia, Brazil, Chile, China, Colombia, Costa Rica, Ecuador, India, Indonesia, Mexico, Pakistan, Panama, Paraguay, Peru, the Philippines, Sri Lanka, Turkey, the United States, Uruguay, Venezuela, and Vietnam. *Panjiva* sources its raw data from government agencies as well as private sources, with the underlying data structure depending on country-specific disclosure regimes. In most cases, the database relies on bill-of-lading (BoL) data or customs declaration data.

For example, *Panjiva* obtains shipment-level information for the United States by collecting bills of lading from U.S. Customs and Border Protection (CBP), which are legally required to be publicly disclosed. In contrast, trade data for China are sourced from China Customs and are available only in monthly aggregated form rather than at the shipment level. Due to heterogeneity in national disclosure regimes, the coverage of transport modes varies across countries, with some datasets limited to specific modes such as maritime or air freight. Specifically, the datasets for Brazil, Ecuador, and the United States cover only maritime shipments, while Sri Lanka includes both air and maritime transport; the remaining countries report imports and exports across all transport modes. Across countries, trade data typically exclude transactions between domestic firms as well as transactions that do not involve the exchange of physical goods, such as tourism or financial services.

Although some country datasets provide shipment-level disaggregation while others do not, after harmonizing data from different sources, the *Panjiva* database consistently includes information on shipment or arrival time (at the date or month level), the exporter (shipper) in export records, the importer (consignee) in import records, and the HS-6–level product classification.[41] Shipment values are also available, although for Brazil, the United States, and Vietnam these values are estimated. Finally, *Panjiva* assigns a unique firm identifier based on company information in the dataset, which facilitates the matching of exporters and importers across trade records.

Figure B.1 compares annual export and import values for each country constructed from the Panjiva dataset with those reported in UN Comtrade. Overall, Panjiva covers the majority of trade flows for most countries. Although gaps between Panjiva and UN Comtrade arise in some cases—most notably for the United States due to differences in underlying data collection—the overall trends in trade flows are broadly comparable across the two sources. For Brazil, Costa Rica, Pakistan, and Panama, however, the correspondence with UN Comtrade is weaker, and results for these countries should therefore be interpreted with caution.

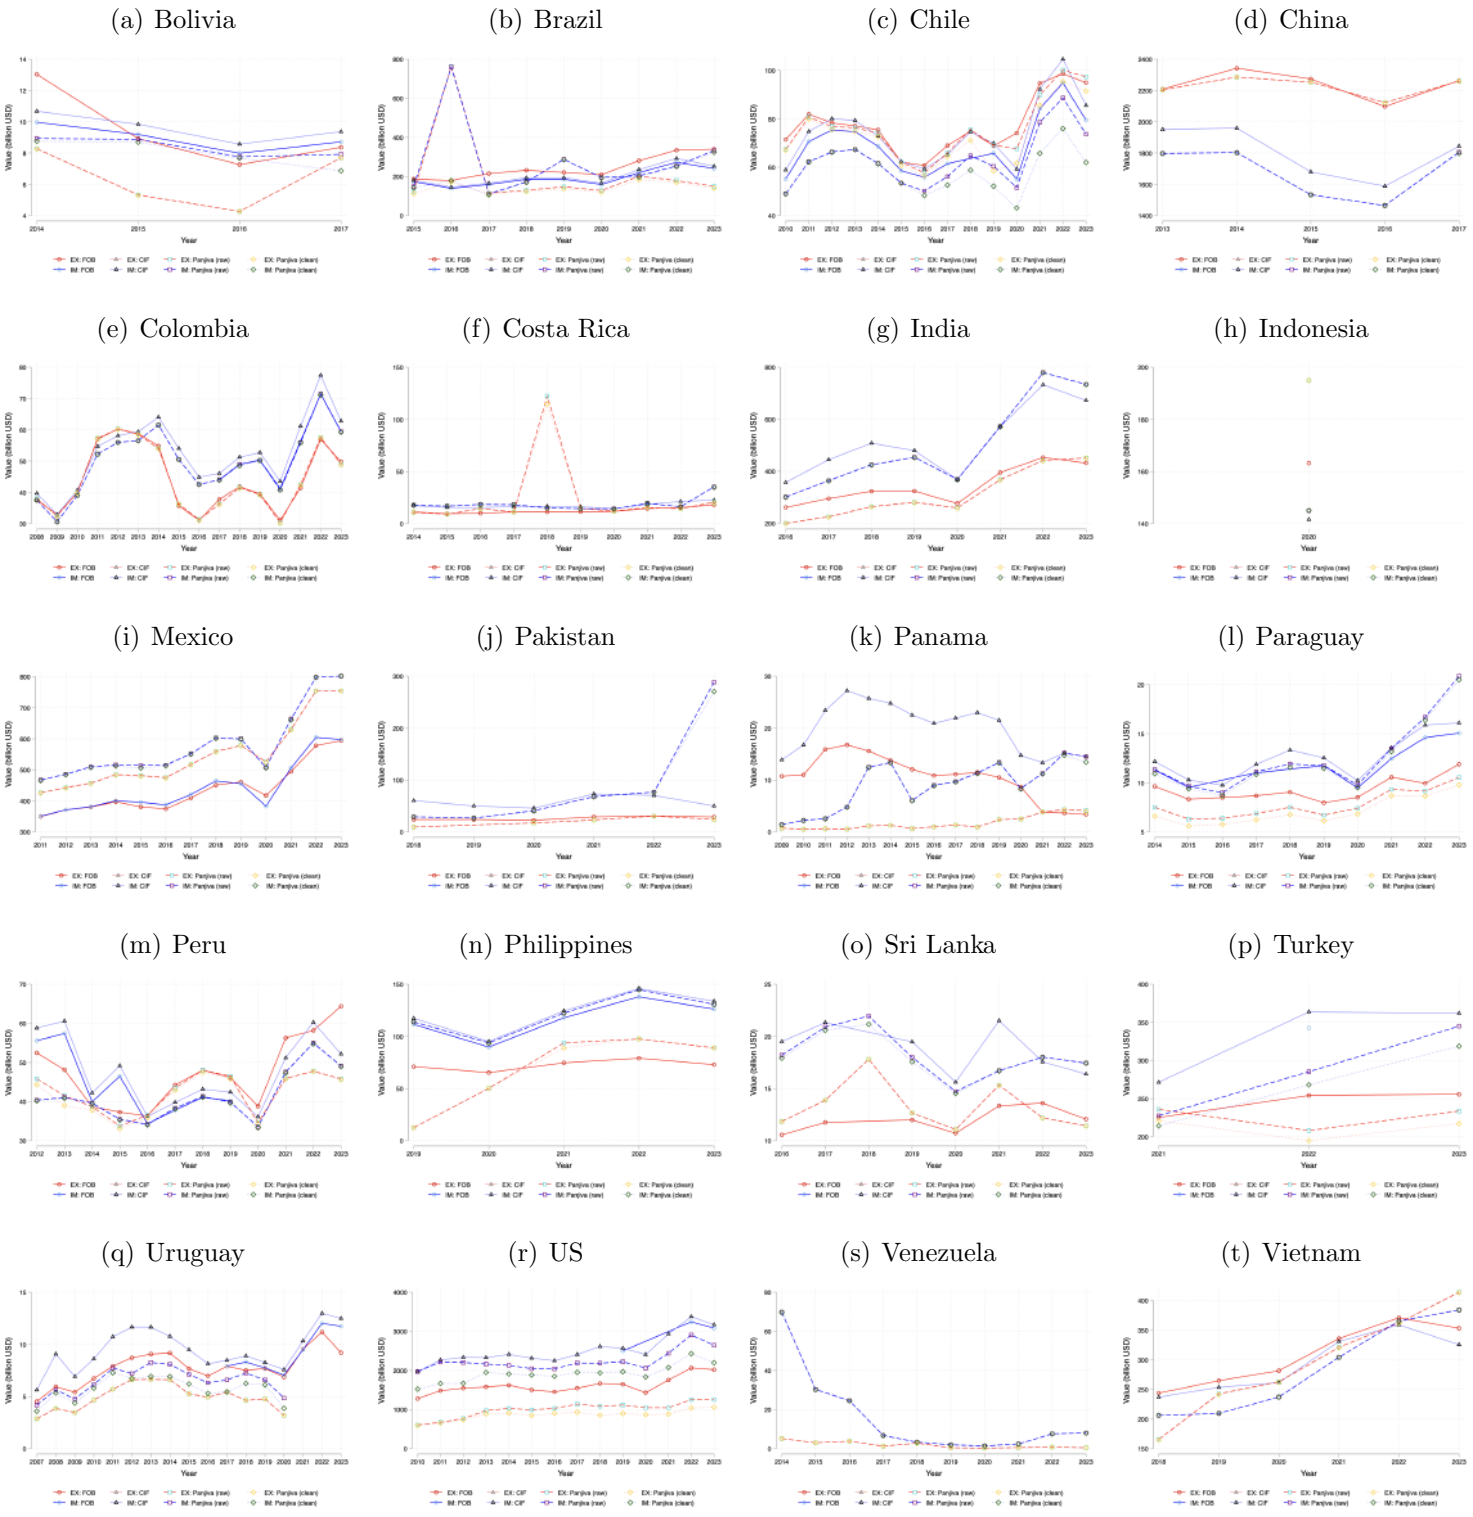
Prior to analysis, we clean the Panjiva dataset by removing irregular HS codes and obser-

[41] The Ecuador dataset does not report HS codes and is therefore excluded from the analysis.

vations lacking firm-level identifiers. As shown in Figure B.1, the differences between the raw and cleaned Panjiva series are modest. In some cases, data quality issues related to HS code reporting reduce coverage, particularly for the United States. By contrast, for Brazil, the cleaning procedure improves the alignment between Panjiva-based totals and those reported in UN Comtrade.

Table B.1 summarizes the final analytical sample, reporting country coverage, sample periods, and the number of firms that appear as both exporters and importers and are used in the linkage prediction analysis. China, India, Mexico, Turkey, the United States, and Vietnam have relatively large samples, each with more than 10,000 firms available for input–output linkage prediction. For countries with smaller samples—due either to country size or data collection limitations—including Bolivia, Colombia, Panama, Paraguay, Uruguay, and Venezuela, the empirical analysis should be interpreted more conservatively.

Figure B.1: Trade data coverage: *Panjiva* vs. *UN Comtrade*



Note: Figure displays aggregated export (red) and import (blue) values from the *Panjiva* dataset (dashed lines) and those from *UN Comtrade* (solid lines). Using *Panjiva*, annual export and import totals are computed based on shipment or arrival time. Trade values reported in *Panjiva* are recorded on either an FOB or CIF basis. For comparison, the figure also reports FOB and CIF values, where available, from *UN Comtrade*. Both the raw *Panjiva* data and the cleaned *Panjiva* data used in the analysis are shown.

Table B.1: Annual importer and exporter counts by country

(1) Year	(2) Total	(3) Exporter	(4) Importer	(5) Both	(6) B/T (%)	(7) B/E (%)	(8) B/I (%)
Bolivia							
2014	84,654	2,231	83,230	807	0.95	36.17	0.97
2015	66,040	2,196	64,631	787	1.19	35.84	1.22
2016	61,018	1,851	59,882	715	1.17	38.63	1.19
2017	12,267	1,832	11,055	620	5.05	33.84	5.61
Brazil							
2015	26,638	4,624	9,377	1,388	5.21	30.02	14.80
2016	30,954	5,489	11,371	1,824	5.89	33.23	16.04
2017	29,123	5,356	11,399	1,830	6.28	34.17	16.05
2018	27,506	5,080	10,757	1,674	6.09	32.95	15.56
2019	27,662	5,523	10,953	1,844	6.67	33.39	16.84
2020	26,561	5,621	10,636	1,643	6.19	29.23	15.45
2021	27,806	5,823	11,251	1,714	6.16	29.43	15.23
2022	24,498	5,595	10,331	1,623	6.63	29.01	15.71
2023	21,802	5,072	9,329	1,370	6.28	27.01	14.69
Chile							
2010	40,334	7,610	35,629	2,905	7.20	38.17	8.15
2011	43,293	7,845	38,482	3,034	7.01	38.67	7.88
2012	44,487	7,675	39,781	2,969	6.67	38.68	7.46
2013	46,550	7,664	41,783	2,897	6.22	37.80	6.93
2014	48,894	8,044	43,799	2,949	6.03	36.66	6.73
2015	47,571	7,852	42,290	2,571	5.40	32.74	6.08
2016	37,672	7,278	32,800	2,406	6.39	33.06	7.34
2017	18,213	5,594	14,820	2,201	12.08	39.35	14.85
2018	19,587	5,418	16,226	2,057	10.50	37.97	12.68
2019	17,781	5,212	14,219	1,650	9.28	31.66	11.60
2020	10,412	3,125	8,504	1,217	11.69	38.94	14.31
2021	8,927	2,905	7,082	1,060	11.87	36.49	14.97
2022	7,816	2,741	6,005	930	11.90	33.93	15.49
2023	11,447	3,302	9,195	1,050	9.17	31.80	11.42
China							
2013	333,532	270,828	159,219	96,516	28.94	35.64	60.62
2014	351,389	285,160	162,978	96,749	27.53	33.93	59.36
2015	355,708	290,188	162,240	96,720	27.19	33.33	59.62
2016	366,938	298,956	165,284	97,302	26.52	32.55	58.87
2017	365,239	297,852	164,993	97,606	26.72	32.77	59.16
Colombia							
2008	39,007	12,234	29,734	2,980	7.64	24.36	10.02
2009	36,342	9,121	27,922	737	2.03	8.08	2.64
2010	37,487	7,761	29,979	285	0.76	3.67	0.95
2011	41,598	7,682	34,075	191	0.46	2.49	0.56
2012	44,193	7,744	36,585	160	0.36	2.07	0.44
2013	43,233	7,816	35,514	120	0.28	1.54	0.34
2014	43,486	7,861	35,705	116	0.27	1.48	0.32
2015	41,705	8,042	33,727	107	0.26	1.33	0.32
2016	41,847	8,244	33,661	94	0.22	1.14	0.28
2017	41,041	8,532	32,596	115	0.28	1.35	0.35
2018	40,248	8,877	31,456	107	0.27	1.21	0.34
2019	44,141	10,107	34,876	875	1.98	8.66	2.51
2020	41,861	8,440	33,488	95	0.23	1.13	0.28

Continued on next page

Table B.1 (continued)

(1) Year	(2) Total	(3) Exporter	(4) Importer	(5) Both	(6) B/T (%)	(7) B/E (%)	(8) B/I (%)
2021	36,234	8,942	27,360	91	0.25	1.02	0.33
2022	36,655	9,288	27,428	81	0.22	0.87	0.30
2023	37,860	9,684	28,253	101	0.27	1.04	0.36
Costa Rica							
2014	13,244	2,987	12,296	2,041	15.41	68.33	16.60
2015	13,459	2,988	12,559	2,088	15.51	69.88	16.63
2016	13,945	3,134	12,928	2,117	15.18	67.55	16.38
2017	15,348	3,856	13,823	2,333	15.20	60.50	16.88
2018	14,462	3,038	13,237	1,813	12.54	59.68	13.70
2019	14,471	3,114	13,197	1,840	12.72	59.09	13.94
2020	13,819	3,534	11,887	1,604	11.61	45.39	13.49
2021	15,354	4,108	13,078	1,836	11.96	44.69	14.04
2022	13,232	2,844	11,969	1,583	11.96	55.66	13.23
2023	18,527	3,881	15,748	1,335	7.21	34.40	8.48
India							
2016	201,922	111,731	129,330	39,979	19.80	35.78	30.91
2017	199,719	111,586	128,754	40,638	20.35	36.42	31.56
2018	209,242	121,626	131,652	44,091	21.07	36.25	33.49
2019	240,923	141,468	147,503	48,223	20.02	34.09	32.69
2020	228,460	140,240	133,188	45,069	19.73	32.14	33.84
2021	248,762	153,757	142,121	47,263	19.00	30.74	33.26
2022	255,553	168,256	137,041	49,989	19.56	29.71	36.48
2023	244,130	156,655	138,951	51,648	21.16	32.97	37.17
Indonesia							
2020	41,007	15,358	33,636	8,137	19.84	52.98	24.19
Mexico							
2013	64,488	31,008	53,095	19,616	30.42	63.26	36.95
2014	62,292	30,469	51,652	19,829	31.83	65.08	38.39
2015	63,463	30,791	52,652	19,980	31.48	64.89	37.95
2016	67,809	32,961	55,285	20,437	30.14	62.00	36.97
2017	68,471	33,220	55,706	20,455	29.87	61.57	36.72
2018	68,497	32,816	56,026	20,346	29.70	62.00	36.32
2019	69,932	33,196	57,275	20,540	29.37	61.87	35.86
2020	67,288	32,666	54,344	19,722	29.31	60.37	36.29
2021	67,974	33,021	54,420	19,468	28.64	58.96	35.77
2022	67,617	33,084	54,069	19,536	28.89	59.05	36.13
2023	76,182	36,116	58,863	18,797	24.67	52.05	31.93
Pakistan							
2018	29,292	7,801	22,819	2,522	8.61	32.33	11.05
2019	27,636	0	20,668	0	0.00		0.00
2020	35,286	13,519	26,302	4,562	12.93	33.75	17.34
2021	39,049	15,978	27,967	4,903	12.56	30.69	17.53
2022	41,434	17,400	28,859	5,089	12.28	29.25	17.63
2023	42,248	16,608	28,997	3,456	8.18	20.81	11.92
Panama							
2009	7,980	2,134	6,611	766	9.60	35.90	11.59
2010	8,869	2,420	7,280	831	9.37	34.34	11.41
2011	9,022	2,353	7,507	838	9.29	35.61	11.16
2012	8,963	1,436	8,095	569	6.35	39.62	7.03
2013	12,680	2,706	10,950	976	7.70	36.07	8.91
2014	13,282	3,162	11,254	1,134	8.54	35.86	10.08

Continued on next page

Table B.1 (continued)

(1) Year	(2) Total	(3) Exporter	(4) Importer	(5) Both	(6) B/T (%)	(7) B/E (%)	(8) B/I (%)
2015	11,685	1,893	10,689	897	7.68	47.39	8.39
2016	12,815	2,352	11,413	951	7.42	40.43	8.33
2017	12,670	2,744	10,916	990	7.81	36.08	9.07
2018	12,515	2,210	11,273	969	7.74	43.85	8.60
2019	12,985	2,882	11,010	907	6.98	31.47	8.24
2020	10,931	2,541	9,188	798	7.30	31.40	8.69
2021	10,888	2,318	9,470	900	8.27	38.83	9.50
2022	9,698	1,946	8,778	1,026	10.58	52.72	11.69
2023	7,026	1,287	6,626	887	12.62	68.92	13.39
Paraguay							
2014	6,320	1,025	6,039	745	11.79	72.68	12.34
2015	6,166	1,067	5,868	769	12.47	72.07	13.10
2016	6,134	1,105	5,811	782	12.75	70.77	13.46
2017	6,235	1,099	5,920	785	12.59	71.43	13.26
2018	6,895	1,108	6,585	808	11.72	72.92	12.27
2019	6,177	1,199	5,845	870	14.08	72.56	14.88
2020	6,083	1,260	5,725	902	14.83	71.59	15.76
2021	6,389	1,397	5,946	954	14.93	68.29	16.04
2022	6,351	1,368	5,902	919	14.47	67.18	15.57
2023	6,365	1,316	5,914	865	13.59	65.73	14.63
Peru							
2012	27,743	7,311	21,760	1,329	4.79	18.18	6.11
2013	28,374	7,218	22,475	1,319	4.65	18.27	5.87
2014	28,799	7,428	22,711	1,340	4.65	18.04	5.90
2015	28,650	6,894	23,043	1,289	4.50	18.70	5.59
2016	29,333	6,804	23,733	1,205	4.11	17.71	5.08
2017	30,886	7,040	25,070	1,228	3.98	17.44	4.90
2018	32,493	7,213	26,448	1,171	3.60	16.23	4.43
2019	33,908	7,473	27,615	1,182	3.49	15.82	4.28
2020	35,572	6,624	29,884	941	2.65	14.21	3.15
2021	35,306	6,732	29,426	862	2.44	12.80	2.93
2022	37,123	7,354	30,613	866	2.33	11.78	2.83
2023	39,129	7,462	32,194	564	1.44	7.56	1.75
Philippines							
2019	17,610	2,447	16,298	1,135	6.45	46.38	6.96
2020	18,290	3,595	16,676	1,981	10.83	55.10	11.88
2021	39,801	10,537	31,342	2,078	5.22	19.72	6.63
2022	37,111	3,755	35,583	2,227	6.00	59.31	6.26
2023	36,677	3,836	35,068	2,227	6.07	58.06	6.35
Sri Lanka							
2016	16,608	5,274	13,330	2,002	12.05	37.96	15.02
2017	18,441	5,484	15,205	2,254	12.22	41.10	14.82
2018	20,996	5,855	17,515	2,375	11.31	40.56	13.56
2019	19,802	5,754	16,163	2,115	10.68	36.76	13.09
2020	20,620	5,723	17,537	2,640	12.80	46.13	15.05
2021	17,331	5,720	14,202	2,592	14.96	45.31	18.25
2022	15,417	4,982	12,755	2,397	15.55	48.11	18.79
2023	16,511	5,106	13,741	2,339	14.17	45.81	17.02
Turkey							
2021	42,775	36,825	28,052	22,102	51.67	60.02	78.79
2022	45,547	38,903	28,855	22,211	48.77	57.09	76.97

Continued on next page

Table B.1 (continued)

(1) Year	(2) Total	(3) Exporter	(4) Importer	(5) Both	(6) B/T (%)	(7) B/E (%)	(8) B/I (%)
2023	47,227	41,113	31,176	25,062	53.07	60.96	80.39
Uruguay							
2007	8,081	2,216	6,407	544	6.73	24.55	8.49
2008	8,667	2,250	6,963	548	6.32	24.36	7.87
2009	8,582	1,974	7,046	440	5.13	22.29	6.24
2010	9,142	2,027	7,634	521	5.70	25.70	6.82
2011	9,659	2,042	8,166	550	5.69	26.93	6.74
2012	9,910	2,076	8,385	552	5.57	26.59	6.58
2013	10,185	2,119	8,633	571	5.61	26.95	6.61
2014	11,171	1,936	9,884	653	5.85	33.73	6.61
2015	11,987	1,990	10,752	756	6.31	37.99	7.03
2016	11,775	1,939	10,589	758	6.44	39.09	7.16
2017	11,935	1,940	10,760	767	6.43	39.54	7.13
2018	10,891	1,757	9,923	793	7.28	45.13	7.99
2019	11,371	1,944	10,205	781	6.87	40.17	7.65
2020	10,676	1,659	7,998	618	5.79	37.25	7.73
US							
2010	361,838	36,151	334,728	9,287	2.57	25.69	2.77
2011	351,594	31,986	327,799	8,429	2.40	26.35	2.57
2012	365,484	43,434	333,144	11,631	3.18	26.78	3.49
2013	381,630	56,147	339,452	14,824	3.88	26.40	4.37
2014	387,006	58,945	342,084	15,107	3.90	25.63	4.42
2015	373,377	52,129	334,512	14,336	3.84	27.50	4.29
2016	387,348	50,528	350,222	14,280	3.69	28.26	4.08
2017	397,380	49,089	361,542	13,955	3.51	28.43	3.86
2018	414,185	48,201	378,892	13,581	3.28	28.18	3.58
2019	401,022	44,653	368,244	12,643	3.15	28.31	3.43
2020	431,952	42,234	400,671	12,026	2.78	28.47	3.00
2021	462,869	39,657	433,564	11,007	2.38	27.76	2.54
2022	469,035	37,897	441,429	11,049	2.36	29.16	2.50
2023	453,017	35,629	426,899	10,366	2.29	29.09	2.43
Venezuela							
2014	22,076	1,171	21,332	427	1.93	36.46	2.00
2015	16,062	1,129	15,327	394	2.45	34.90	2.57
2016	11,408	974	10,778	345	3.02	35.42	3.20
2017	8,522	689	8,107	274	3.22	39.77	3.38
2018	5,753	503	5,459	209	3.63	41.55	3.83
2019	4,577	598	4,118	139	3.04	23.24	3.38
2020	4,498	243	4,335	80	1.78	32.92	1.85
2021	5,566	383	5,298	117	2.10	30.55	2.21
2022	9,700	576	9,255	162	1.67	28.13	1.75
2023	9,685	549	9,312	177	1.83	32.24	1.90
Vietnam							
2018	86,320	40,376	71,401	25,856	29.95	64.04	36.21
2019	88,945	44,446	72,520	28,410	31.94	63.92	39.18
2020	92,616	46,306	74,505	28,795	31.09	62.18	38.65
2021	90,959	45,708	72,312	27,838	30.60	60.90	38.50
2022	105,236	52,712	82,939	30,853	29.32	58.53	37.20

2023	142,837	80,565	102,817	40,624	28.44	50.42	39.51

Note: Table reports the annual number of importers and exporters in each country recorded in the *Panjiva* dataset. Firm identities are based on the unique identifiers provided by *Panjiva*. An exception is the United States, where firm records are matched using fuzzy name-matching techniques. Column (5) reports the number of firms that appear as both importers and exporters. B/T, B/E, and B/I denote Both divided by Total, Exporter, and Importer, respectively.