



Universität St.Gallen

Harmonizing the Harmonized System

Piotr Lukaszuk and David Torun

October 2022 Discussion Paper no. 2022-12

Editor: Mirela Keuschnigg
University of St.Gallen
School of Economics and Political Science
Department of Economics
Müller-Friedberg-Strasse 6/8
CH-9000 St.Gallen
Email seps@unisg.ch

Publisher: School of Economics and Political Science
Department of Economics
University of St.Gallen
Müller-Friedberg-Strasse 6/8
CH-9000 St.Gallen

Electronic Publication: <http://www.seps.unisg.ch>

Harmonizing the Harmonized System¹

Piotr Lukaszuk and David Torun

Author's address:

Piotr Lukaszuk
State Secretariat for Economic Affairs (SECO)
and University of St.Gallen
Email piotr.lukaszuk@seco.admin.ch
Website <https://ideas.repec.org/f/plu427.html>

David Torun
SIAW-HSG, University of St.Gallen
Bodanstrasse 8
9000 St. Gallen, Switzerland
Email david.torun@unisg.ch
Website <https://sites.google.com/view/davidtorun/>

¹ We thank Emily Blanchard, Peter Egger, Simon Evenett, Reto Foellmi, Stefan Legge, Marc Muendler, Marcelo Olarreaga, Ralph Ossa, Frank Pisch, Ulrich Schetter, and the participants to several conferences and seminars for helpful comments and discussions. The views expressed in this study are the authors' and do not reflect those of SECO.

Abstract

International trade research relies heavily on data reported at the product level. Regular classification updates of the “Harmonized System” lead to intertemporal inconsistencies affecting up to 44% of world goods trade. Existing methods to standardize product vintages either drop numerous products over time or bulk updated codes into large synthetic categories. We largely overcome these issues by developing an algorithm that exploits the persistence of trade data to convert trade flows between vintages. Our conversion estimates are robust to year and sample choices. Provided a clear correspondence and persistent data, the algorithm can be applied to other classifications.

Keywords

international trade data; product classification; classification updates; Harmonized System (HS); concordance

JEL Classification

F1; C81

1 Introduction

Empirical research on international trade has long expanded from the aggregate analysis of trade flows to more granular data at the sector, product, firm, and even transaction level. Given that international trade figures by default include flows between two different jurisdictions, these granular data require an internationally harmonized classification of the shipped products. Since 1988, the World Customs Organization (WCO) has developed the Harmonized Commodity Description and Coding System (HS) to globally standardize the classification of products into over 5,000 6-digit HS codes.¹

Although the harmonization in 1988 created a standard allowing for cross-country comparisons of trade figures going beyond trade flows to, e.g., tariff lines, it was unable to avoid necessary product code adjustments over time. On average every five years, the WCO updates its list of HS product codes and publishes a new vintage.² The amendments have been necessary, inter alia, due to technological progress (creating new products and making some void), shifting trade patterns (increasing or declining relevance of specific products), adaptation to reflect trade practice (unclear product classifications or bundling of product parts by firms), as well as amendments related to social and environmental fields (e.g., to improve the tracking of endangered plants and animals).³

While these adjustments are meant to facilitate trade, they often result in products being split up (“1:n”), merged (“m:1”), or generally reassigned (“m:n”). This, in turn, complicates the comparison of trade flows both across time and across countries. The former comparison is difficult because it is unclear how to convert trade flows from one HS classification to another whenever a link is 1:n or m:n, while the fact that some jurisdictions adopt the new HS code version (“vintage”) with a delay hinders cross-sectional analyses.⁴ Below, we revisit commonly used conversion methods and develop an alternative to tackle several issues that arise when applying these existing techniques.

Given that some HS code adjustments are interlinked, they may not be viewed as individual changes but rather as a connected network or *group* of adjusted HS codes. Put differently, whenever two codes of an “initial” vintage are assigned to the same code of a “target” vintage, they are (indirectly) linked, and will thus belong to the same *group*. Figure 1 visualizes an exemplary group of HS codes in the form of a network chart, where bigger circles depict larger trade shares within this group. In this example, the “initial” codes *a1* and *a2* are reassigned to “target” codes *b1* and *b2*. Therefore, we want to know


¹Each country may determine and adjust its HS codes at a more granular level. For instance, the United States classifies its products at the 10-digit level in its so-called HTS classification. However, the HS codes remain globally harmonized at the 6-digit level. Before the introduction of HS codes, goods were largely classified under the so-called SITC system.

²More precisely, revised HS vintages were introduced in 1992, 1996, 2002, 2007, 2012, and 2017. We will refer to a vintage introduced in a specific year as, e.g., HS 2007.

³For further details see Appendix A.

⁴For instance, according to article 4(2) of the HS Convention of the WCO, developing countries—representing three-quarters of the WCO’s 183 members—may adopt the newest vintage with a delay of up to five years.

Figure 1 – Network Visualization of an Exemplary HS Code Group



Notes: This figure illustrates an exemplary HS code group in the form of a network diagram. In this example, code *a1* was reassigned to code *b1* (therefore receiving a conversion weight equal to one), while code *a2* was reassigned to codes *b1* and *b2*, which means that for product *a2* we do not know a priori how much trade should be allocated to which new code when converting the *a*-vintage into the *b*-vintage. Larger circles depict larger product trade shares. We refer to such networks of linked products as product groups.

how we can convert trade flows reported in the *a*-vintage into the *b*-vintage. Since product *a1* has only a link to product *b1*, all trade flows reported in product category *a1* will be added to product *b1* when converting to this new HS vintage. By contrast, because code *a2* has a link with both *b1* and *b2*, it is a priori unclear which part of trade flows classified in the initial vintage should be allocated to *b1* or to *b2*. We will refer to the share of *a2* flows being allocated to *b1* as the *conversion weight* from *a2* to *b1*, because overall trade in *b1* will be a weighted sum of trade in the linked initial products (here, *a1* and *a2*). The conversion weight from *a1* to *b1* is equal to 1, and, accordingly, equal to 0 for *a1* to *b2*, while the conversion weights from *a2* to *b1* and *b2* are unknown.

While in each amendment most HS codes remain unchanged, the non-1:1 adjustments accumulate when the conversion is performed over multiple vintages. As indicated in Table 1, the majority of HS codes remain the same and have a 1:1 assignment when converting the HS 2017 codes into HS 2012. However, the further one extends the conversion back to earlier vintages, the larger the share of conversions with undefined weights becomes. When converting HS 2017 codes into HS 1992 codes, the share of products with unknown conversion weights increases to about 16%.⁵ More generally, out of 6,539 HS codes applied between 1992 and 2017, 48% have faced a non-1:1 adjustment at least in

⁵Note that one may instead want to convert HS 1992 to HS 2017 codes. In that case, those HS codes summarized in Table 1 as “m:1” will have unknown conversion weights.

Table 1 – Distribution of Product Links between HS 2017 and Other Vintages

Relationship with HS 2017	HS 2012	HS 2002	HS 1992	Conversion Weight
1:1	4744	3966	3332	1
$m:1$	578	948	1332	1
$1:n$	36	262	360	?
$m:n$	60	316	538	?

Notes: This table depicts how many product links between HS 2017 codes and other vintages are of a specific type. As illustrated by the question marks, whenever a link is “ $1:n$ ” (i.e., the code has to be split into multiple codes of the other reporting standard) or “ $m:n$ ” (i.e., m HS 2017 codes are interlinked with n codes of the other HS vintage), the correct conversion is unknown.

Data source: UN Comtrade Database.

one of the classification updates. In terms of 2018 trade, these codes accounted for 43.7% of total trade.

Widely-used methods address the HS code conversion issue by either forcing a one-to-one mapping between product codes across vintages (UNSD, 2017) and, thus, ignoring remaining product links, or by pooling adjusted codes into synthetic codes (Pierce and Schott, 2012). The former method leads to intertemporal inconsistencies in HS vintage amendment years. For instance, when converting HS 2007 codes into the HS 2002 vintage, the number of non-traded products increases from 0 in the year 2006 to 299 in the year 2007, which represents 42% of codes where the assignment is not pre-determined. Meanwhile, the Pierce and Schott (2012) method increases the trade share of the largest product code in overall trade of the year 2018 from about 7% to 26%, when converting the HS 2017 codes into HS 1992 codes. Not only do we then obtain, by construction, more aggregated data, the distribution also appears to become more skewed due to the fact that unchanged codes are never lumped into these synthetic clusters. Researchers may circumvent these issues either by using only data at coarser levels, as in gravity regressions at the country level (cf. Head and Mayer (2014) for an overview), or by focusing exclusively on products with a 1:1 mapping (e.g., Boehm et al., 2020). While the former approach eliminates interesting variation, the latter one substantially reduces the sample size for longer time horizons.⁶

In light of the available methods, this paper aims to contribute to the existing liter-

⁶In the paper we exclusively present results based on converting backwards (i.e., from newer vintages to older ones), because Comtrade provides only backward conversions. However, our algorithm can also be applied to forward conversion. For illustrative purposes, we focus on the example of converting HS 2007 to HS 2002, since that change lies in the middle of our sample and happened before the Great Trade Collapse. Results for other conversions look very similar.

ature by proposing an alternative conversion technique between different vintages of HS product codes. The key feature of trade data we exploit is their persistence over time. Specifically, we provide an algorithm that through the choice of conversion weights minimizes the squared deviations between pre- and post-classification-change country-pair data. This specification is in line with the trends we observe in product-level trade data. Our minimization imposes a set of constraints which respects the product links, thereby taking into account indirect links within product groups. The algorithm largely solves the aforementioned issues of existing methods, and the calculated weights are robust to altering the set of years and country-pair observations.

In a simple descriptive application, we illustrate how our method may facilitate answering simple questions like: How many different varieties did the U.S. import each year over the past three decades? We show that using data converted to HS 1992 by Comtrade suggest that the number of products with positive U.S. imports has declined by about 10% from 1992 to 2018. By contrast, data converted into HS 1992 using our weights show that this number has been fairly stable over that time span. [Pierce and Schott’s \(2012\)](#) method has a level effect in that there are a quarter less product categories available in the sample, but the trend looks similar to that suggested by our conversion. Such differences matter, e.g., when we want to investigate the evolution of product varieties à la [Broda and Weinstein \(2006\)](#), where a variety is defined as an exporter-product combination. For instance, when looking at the aftermath of the Great Trade Collapse, data converted using our technique suggest that the number of varieties imported by the U.S. had recovered to its pre-crisis level by 2011. The Comtrade-converted data, however, hint at a more sluggish regeneration, as there the number of imported varieties reaches its pre-crisis level only by 2015.

The following section summarizes existing methods. [Section 3](#) then describes our conversion algorithm. [Section 4](#) reports the performance of our method, while [Section 5](#) contextualizes it in simple descriptive exercises. [Section 6](#) explores the heterogeneity and robustness of the conversion weights. [Section 7](#) concludes.

2 Related Literature

Thus far there have been two main approaches to address the HS code conversion issue: the method applied by the UN Statistics Department which generates the Comtrade trade database (“Comtrade method”), and an algorithm developed by [Pierce and Schott \(2012\)](#).

According to the Comtrade method, “the data conversions [...] assign one single code (subheading) of an earlier HS edition to each HS 2017 subheading” ([UNSD, 2017](#), p.1). This method thus always assigns a weight of 1 to a single target HS code, and ignores the remaining target HS codes. In terms of the conversion problem depicted in [Figure 1](#), this would, e.g., imply that the link between $a2$ and $b1$ is ignored with all $a2$ trade now converted to $b2$, thus imposing conversion weights of 0 and 1, respectively. In general,

the Comtrade approach therefore tends to assign (too) much weight to one single code, and thus results in some codes not receiving any trade at all. Furthermore, using this method may lead to significant jumps at the extensive margin in HS code amendment years (Cebeci, 2012, p.5).

In terms of data continuity, the Comtrade assignment also leads to jumps at the intensive margin of trade. As depicted in Figure 2, when we look at trade data provided in HS 2002, the relative trade shares within conversion-dependent groups jump in a conversion year (here, in 2007), as the Comtrade method simplifies product links. By contrast, these product hierarchies (i.e., the product trade shares within groups) are highly persistent over time, as Figure 2 shows. Accordingly, there is no a priori reason to expect (large) jumps in product shares within their groups which are not due to a misspecified conversion. The share of zeros jumps from 0% in 2006 to 42.4% in 2007. Note that this figure includes only groups (or, networks) with at least one 1: n or m : n assignment from HS 2007 to HS 2002.

Alternatively to this method, Pierce and Schott (2012, p.58) developed an algorithm which avoids the data continuity issues by pooling the adjusted HS codes together and “assigning a common synthetic code to all HS codes in a growing family tree.” A family tree is an HS product group in our notation. While this method has been widely applied (e.g., Wagner and Zahler (2015), Fort et al. (2018), Felbermayr et al. (2019), Piveteau and Smagghue (2019), Bellert and Fauceglia (2019)), pooling linked codes becomes an issue when analyzing trade data reported in multiple HS code vintages.

In fact, the longer the time series, the more HS codes are piled together. For instance, when converting 2018 data back to HS 1992 using their method, one synthetic cluster of HS codes ends up covering over a quarter of overall trade (cf. Table 2).⁷ Apart from creating synthetic codes not applied in practice, pooling that much trade across sectors may have significant consequences for what kind of variation is (not) picked up by product fixed effects. Moreover, Table 2 clearly illustrates that this method tends to artificially skew the distribution, especially when we consider that 1:1 links are never piled together.⁸

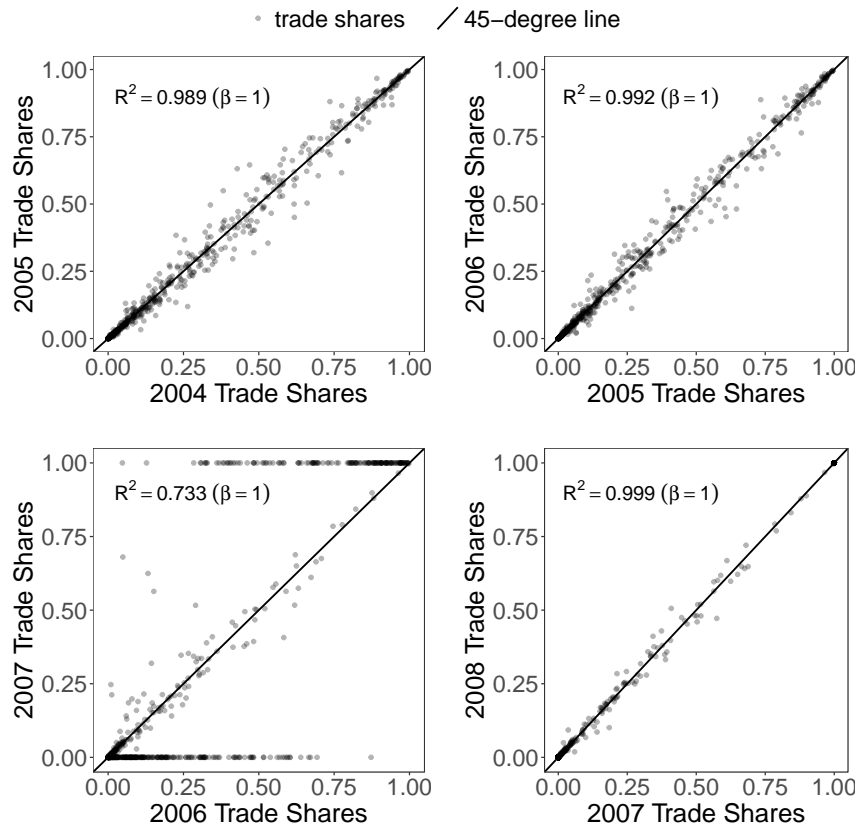
At the same time, one should point out that the Pierce and Schott (2012) method was initially developed to address monthly adjustments of individual 10-digit HTS codes (the U.S. version of the detailed HS codes), and was extended to the HS code conversion. For the initial purpose, their method seems adequate provided a small number of irregular code amendments. However, as indicated in the paper (p. 56), most of the HTS code adjustments are clustered around the general HS code vintage updates.

Similar methods to that of Pierce and Schott (2012) have been proposed by Cebeci

⁷This largest synthetic group covers 607 different HS 2017 codes and includes sectors ranging from chemicals (ca. 27% of the codes) through paper products (10%) to machinery and electrical equipment (40%). The second largest HS code group which accounts for almost exactly a fifteenth of 2018 trade is crude oil and the corresponding HS code 270900 remains unchanged from 1992 to 2017, as can be seen by the unchanged share of 6.66% in each row of Table 2.

⁸The bottom three rows in Table 2 also illustrate that the Comtrade conversion method leads to bigger product groups but to a much lesser extent.

Figure 2 – Product Trade Shares within HS Product Groups over Time (Comtrade Method)



Notes: This figure plots for each product its trade share within its specific group across several years. We include only groups which contain at least one 1:n or m:n assignment from HS 2007 to HS 2002. The figure uses Comtrade data provided in HS 2002, i.e., 2007 and 2008 trade data have been converted into HS 2002, while previous years were originally reported in HS 2002. The data are restricted to 70 importers reporting in the latest available HS code version in a given year. Product groups are explained in the main text. The provided R-squared assumes a fit along the 45-degree line.

Data source: [UN Comtrade Database](#).

(2012) and [Diodato \(2018\)](#). The former also pools product codes but applies the method directly to HS codes and accounts for the fact that some developing countries may implement the HS code adjustment at a later stage. [Diodato \(2018\)](#) depicts product links across vintages using a network-adjusted structure of codes.⁹

Below, we develop a simple conversion algorithm to alleviate the issues encountered when applying the above-mentioned approaches.

3 Methodology

This section briefly introduces the data we use, before explaining our conversion algorithm.

⁹As opposed to [Pierce and Schott \(2012\)](#), [Diodato \(2018\)](#) removes certain links using a community discovery algorithm. However, this method still creates large synthetic groups, and it may ignore certain dependencies between product codes across vintages that could be important for particular products.

Table 2 – Largest Trade Shares by Actual/Synthetic HS Code in 2018 Trade Data

Category	Covered vintages	Largest code	2nd largest	3rd largest	4th largest	5th largest
Actual	HS 2017	6.66%	2.17%	1.92%	1.61%	1.59%
Synthetic (P&S method)	HS 2012–HS 2017	13.43%	6.66%	4.17%	2.17%	1.92%
Synthetic (P&S method)	HS 2002–HS 2017	16.00%	6.66%	4.17%	3.43%	2.35%
Synthetic (P&S method)	HS 1992–HS 2017	25.74%	6.66%	4.17%	3.43%	1.92%
Comtrade converted	HS 2012	6.66%	2.17%	1.92%	1.61%	1.61%
Comtrade converted	HS 2002	6.66%	4.17%	2.46%	2.25%	1.92%
Comtrade converted	HS 1992	6.66%	4.19%	3.44%	2.46%	1.91%

Notes: This table depicts trade shares in overall trade of the largest products in the year 2018. The first row displays the trade share of the largest HS 2017 products. The following three rows depict the trade shares of the largest products after applying [Pierce and Schott’s \(2012\)](#) method to standardize the vintages from 2017 to those of 2012, 2002, and 1992, respectively. The bottom three rows summarize the shares of the top five products after the Comtrade conversion of 2018 trade data into HS 2012, HS 2002, and HS 1992.

Data source: [UN Comtrade Database](#).

Data. The data used in this paper have been extracted from UN Comtrade and include all import data reported by UN members between 1988 and 2018. This data source is standard to the international trade literature. The import figures are corrected for re-import numbers, as the raw import data include re-imports. The analysis focuses on those countries that in each year applied the latest available HS vintage, and are thus considered “compliers.”¹⁰ An exemplary phasing in of the HS 2007 vintage is shown in [Figure B.1](#). The bilateral country pairs include complying importers and all of the partners with which they had at least one non-zero trade flow in a given year.

Furthermore, the analysis includes import data that have been converted by Comtrade to a given HS vintage (e.g., the 2007 and 2008 numbers converted by Comtrade into HS 2002 used in [Figure 2](#)).

Persistence. The key feature of trade data we want to exploit is the persistence of trade flows over time. The data show that, e.g., if a specific product was relatively important for a country pair in one year, it is very likely to be important for that country pair in surrounding years. Similarly, if a country is a large importer of a specific product in one year, it is likely to also have relatively high imports of that product in other years. Accordingly, when we look at product hierarchies within groups over time (for a given HS vintage), they appear to be fairly stable (cf. [Figure 2](#)). Product-level trade flows at the bilateral (i.e., country pair) level also reveal a relatively high persistence over time. [Table 3](#) shows

¹⁰For instance, Vietnam reported its 2017 trade figures using the HS 2012 version and would thus not be included when estimating conversion weights between HS 2012 and HS 2017. However, it did change to the latest available version in 2012, so it would be included when converting HS 2012 to HS 2007.

Table 3 – Persistence of Bilateral Product-Level Trade Flows

Year	2005	2006	2007	2008	2009
All HS codes	0.874 (0.009)	0.867 (0.008)	<i>skipped</i> <i>due to</i>	0.854 (0.009)	0.862 (0.009)
Adjusted HS codes	0.876 (0.020)	0.889 (0.018)	<i>vintage</i> <i>change</i>	0.880 (0.024)	0.836 (0.045)

Notes: This table depicts first-order autoregressive coefficients for bilateral trade flows at the product level for the years 2005, 2006, 2008 and 2009 (2007 was left out due to the HS vintage change). The bilateral product-level flows are scaled by overall trade in the product’s group in a given year. The data include 70 countries which reported in the latest available HS code version in the years 2005–2009. The first row depicts the coefficients for all trade flows, whereas the second is restricted to HS codes facing a change in 2007 (i.e., which have a non-1:1 assignment between HS 2002 and HS 2007). We regress the bilateral product-level flows on their respective first lags without a constant. Robust standard errors are reported in parentheses.

Data source: [UN Comtrade Database](#).

this by regressing the share of bilateral product-level flows in their groups on their first lags (cf. also Figure B.2). Therefore, we use information on bilateral flows reported in one HS vintage to predict how trade flows should be allocated within a target vintage in cases where the allocation is not a priori clear (i.e., neither 1:1 nor $m:1$). The reasoning behind our method is that we do not expect (substantial) changes in product hierarchies within groups due to a change in classification standards. Moreover, using disaggregated rather than aggregated product-level trade flows is appealing for two reasons. Not only does this add information on how trade flows are distributed across country pairs, it also avoids encountering regressions where we have less observations than unknowns.

Algorithm. The algorithm can be split into three components. As a first step, we identify groups of HS codes that are linked together, and keep those groups which contain ex ante unspecified conversion weights.¹¹ Note that one group would be pooled to one synthetic cluster when applying [Pierce and Schott’s \(2012\)](#) method.

Second, in order to maximize the number of observations, we use trade flows at the country-pair level. Specifically, we prepare the trade data in form of an $(I \times J) \times K$ -matrix separately for each group and year, where an entry is equal to $X_{ij,k}$, i.e., the exports from country $i \in I$ to country $j \in J$ of product $k \in K$ in a specific group and year. For notational simplicity we suppress time and group subscripts. We restrict the sample to importers which actually switched reporting standards in the years we focus on. Moreover, in order to consider product hierarchies within groups, we scale each pair’s

¹¹Accordingly, groups that are left out will contain only conversion weights which are a priori known to be equal to zero or one.

product-level imports by the total imports within a group. Note that this also allows us to control for group-specific trends. This procedure yields our scaled flows, $x_{ij,k}$, defined as

$$x_{ij,k} \equiv \frac{X_{ij,k}}{\sum_{\hat{i}, \hat{j}} \sum_{\hat{k}} X_{\hat{i}, \hat{j}, \hat{k}}},$$

which will be the inputs in our regressions. Recall that the scaling is computed separately for each group and year. It is worth noting that this could be further adjusted to calculate, e.g., region- or country-group-specific conversion weights, as addressed in Section 6.

The third component of the algorithm tries to fit the trade flows which are reported in the “target” reporting standard by using flows reported in the “initial” reporting standard and a constrained least-squares procedure. Note that these flows will, by construction, not belong to the same year. Hence, we (implicitly) assume that product-level bilateral flows relative to overall group flows are fully persistent over time. That is why most figures report a fit along the 45-degree line. In Section 6 we relax this assumption.

As an example, consider “initial” HS codes $k \in K$ used in year t_0 , and “target” HS codes $s \in S$ used in year t_1 .¹² We need three types of constraints in our optimization. First, the conversion weights for an initial product k need to sum up to one across target products s , otherwise overall trade might artificially increase or decrease across vintages. Second, the weights should be non-negative, else we would subtract product flows, which might improve the fit but would make no intuitive sense.¹³ Third, some weights are known to be equal to zero or one from the correlation tables provided by the WCO.¹⁴ Our constrained minimization problem is outlined in Equation (1).

$$\begin{aligned} \min_{\{\beta_{k,s}\}} \quad & \sum_s \sum_{i,j} \left(x_{ij,s}^{t_1} - \sum_k x_{ij,k}^{t_0} \beta_{k,s} \right)^2 \\ \text{s.t.} \quad & \beta_{k,s} \geq 0 \quad \forall k, s \\ & \sum_s \beta_{k,s} = 1 \quad \forall k, \end{aligned} \tag{1}$$

where $\beta_{k,s}$ is the conversion weight from initial product k to target product s . Note that various $\beta_{k,s}$ are a priori known to be equal to zero or one. This problem can then be estimated separately for each group to reduce the size of the coefficient matrix, given that conversion weights across groups are, by construction, always equal to zero. This is the specification we use for our main results, and variations can be found in Section 6.

In order to apply our conversion weights to trade data, one has to multiply the trade flow reported in the initial HS vintage by the respective weight, and then sum up across initial codes.¹⁵ Or, put differently, one computes the fitted value, $\hat{x}_{ij,s}^{t_0} = \sum_k x_{ij,k}^{t_0} \hat{\beta}_{k,s}$.


¹²Note that we can also estimate conversion weights going, e.g., from 2017 to 2016 data, in which case $t_0 > t_1$.

¹³These two constraints together restrict conversion weights to be smaller than or equal to one.

¹⁴The correlation tables were downloaded from the UN Comtrade webpage.

¹⁵Also, we recommend rounding the weights to, e.g., the sixth digit after the comma, in order to set

Figure 3 – Product Trade Shares within HS Product Groups over Time (Our Conversion Weights)



Notes: This figure plots for each product its trade share within its specific group across several years. We include only groups which contain at least one 1:n or m:n assignment from HS 2007 to HS 2002, and trade data for 72 countries which used the latest available HS code classification in each year between 2004 and 2008. The figure uses Comtrade data provided in the latest available HS vintage, and transforms 2007 and 2008 data using our weights back into HS 2002. Product groups and the algorithm to obtain our conversion weights are explained in the main text. The provided R-squared assumes a fit along the 45-degree line.

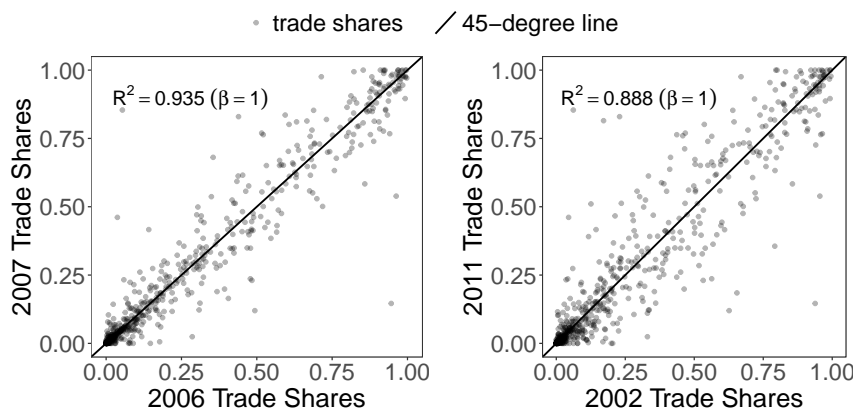
Data source: [UN Comtrade Database](#).

Furthermore, in order to convert one vintage to another that is not adjacent, we multiply the weights across conversions. For instance, if we want to convert HS 2017 codes $k \in K$ to an HS 2007 code l , which are linked through HS 2012 codes $s \in S$, we compute $\hat{x}_{ij,l}^{t_0} = \sum_s \sum_k x_{ij,k}^{t_0} \hat{\beta}_{k,s} \hat{\beta}_{s,l}$. The weights $\hat{\beta}_{k,l}$ ($= \sum_s \hat{\beta}_{k,s} \hat{\beta}_{s,l}$) have the same properties as those constructed based on adjacent vintages $\hat{\beta}_{k,s}$ and $\hat{\beta}_{s,l}$.

4 Results

In Section 1, or in Figure 2, we illustrate how product shares within groups jump in 2007 when applying Comtrade’s conversion method, despite product hierarchies being strikingly stable over time. When applying our estimated conversion weights, we receive weights which are numerically indistinguishable from zero to 0.

Figure 4 – Product Trade Shares within HS Product Groups (Our Conversion Weights)



Notes: This figure plots for each product its trade share within its specific group across several years. We include only groups which contain at least one 1:n or m:n assignment from HS 2007 to HS 2002, and trade data for 63 countries which used the latest available HS code classification in each of the years used in the plot. The figure uses Comtrade data provided in the latest available HS vintage, and transforms 2007 and 2011 data using our weights back into HS 2002. Product groups and the algorithm to obtain our conversion weights are explained in the main text. The conversion weights are calculated using 2006 and 2007 trade data. The provided R-squared assumes a fit along the 45-degree line.

Data source: [UN Comtrade Database](#).

a substantially smoother transition from 2006 to 2007, as depicted by Figure 3. The share of zeros in the graph drops heavily from 42.4% to only 5.6%. We thus manage to reduce the share of HS 2002 products (in total 5,224) with no trade from 6 to below 1%. Moreover, although we tend to slightly decrease the variation between the years 2007 and 2008, it is important to note that it was already very low in the original vintage. Results for other vintages look very similar. Hence, our method alleviates the aforementioned issues introduced by the Comtrade method. [Pierce and Schott’s \(2012\)](#) method would reduce the number of product categories in the graph by roughly three-quarters from 706 to 183, given that we identify 183 groups in our conversion.

We can verify how these estimated conversion weights perform with other year constellations. Specifically, in Figure 4, we convert 2011 data into HS 2002 using the conversion weights from above, and compare the resulting product hierarchies within groups with those in the year 2002. This is the widest year range we can take when comparing HS 2007 with HS 2002. We also re-include the comparison of 2006 shares and converted 2007 shares (left panel). The product shares closely follow the 45-degree line in both panels, with (intuitively) somewhat more variation in the wider year range. In order to compare the product shares in different years but with unconverted data, we can confront the two “initial” as well as the two “target” years. Figure B.3 in Appendix B illustrates that the similarity should be expected (and desirable) given the persistence of product hierarchies measured in the reported HS vintages. Overall, this suggests that, if our method is deemed appropriate to convert 2007 data into HS 2002, it also allows for a conversion of later years into that HS vintage.

Table 4 – Largest Trade Shares by Actual/Converted HS Code in 2018 Trade Data (Our Conversion Weights)

Category	Covered Vintages	Largest Code	2nd largest	3rd largest	4th largest	5th largest
Actual	HS 2017	6.66%	2.17%	1.92%	1.61%	1.59%
Converted	HS 2012	6.66%	2.17%	1.92%	1.73%	1.61%
Converted	HS 2002	6.66%	2.58%	2.25%	1.92%	1.73%
Converted	HS 1992	6.66%	3.43%	2.09%	1.92%	1.73%

Notes: This table depicts trade shares in overall trade of the largest products in the year 2018. The first row displays the trade share of the five largest HS 2017 products. The remaining rows depict the trade shares of the five largest products after applying the algorithm described in this paper and converting 2018 data from HS 2017 to HS 2012, HS 2002, and HS 1992, respectively.

Data source: [UN Comtrade Database](#).

Importantly, we obtain similar outcomes when comparing absolute (i.e., not scaled by group totals) values. In Figure B.4 we re-produce Figure 3 using the log of product-level trade flows. There are only a few products for which we get somewhat large deviations. However, these products tend to be relatively small. This in turn implies that exploiting product group hierarchies in our estimation routine seems to yield conversion weights which are suitable for the analysis of absolute trade flows.

In Section 3 we explain how to standardize HS vintages which are non-adjacent by multiplying the conversion weights for the intermediate vintages. In Figure B.5 we plot the log of product-level trade flows for products which have at least one non-1:1 assignment over our sample period. The figure shows that this procedure does not create significant inconsistencies when converting all the way back to HS 1992, even in recent years, although we do eliminate a little less than 1% of the HS 1992 products (5,020 in total) over the years. In comparison, the Comtrade conversion drops almost 10% of the HS 1992 products by 2017. Importantly, note that our technique does not artificially fix the importance of a product over longer time spans. The multiplication of the weights, which are recalculated whenever there is an amendment, implies that long-run trends should be captured by our method, as long as the HS vintage adjustments reflect such trends.

Additionally, as illustrated in Table 4, our method seems not to skew the distribution, as opposed to [Pierce and Schott’s \(2012\)](#) method (cf. Table 2). Accordingly, at least in terms of relative size, the converted product categories remain interpretable as such. The trade shares obtained using our conversion method perform slightly better than the Comtrade method also in this dimension (cf. bottom three rows in Table 2).

5 Descriptive Exercises

Having shown in the previous section that the proposed method addresses the trade data issues caused by prevalent methods, this section analyses the impact our method has on several trade applications. We provide illustrative examples related to both the demand and supply side. These are meant to demonstrate, for instance, how motivating facts may change depending on the conversion method that is applied.

Import demand trends. One key dimension of assessing welfare gains from trade continues to be the gains from product variety in trade. As shown in [Broda and Weinstein \(2006\)](#), product varieties have significantly increased over the past decades, substantially contributing to the increased role of trade in the global economy ([Hsieh et al., 2020](#)). In order to track product varieties correctly, one ought to have a continued standardised classification of products throughout the sample, requiring thus to correct for the HS code adjustments every five years.¹⁶

Taking U.S. imports as an example, we illustrate in [Figure 5](#) that using our method we obtain a stable flow of 6-digit HS products being imported each year into the United States. Applying instead the Comtrade method to convert all codes back to the HS 1992 vintage leads to artificial reductions of imported products in 2002 and especially in 2007 – years in which the United States updated its applied HS vintage. Therefore, over time the number of imported products appears to be shrinking, despite there being no obvious reason for that to be the case.

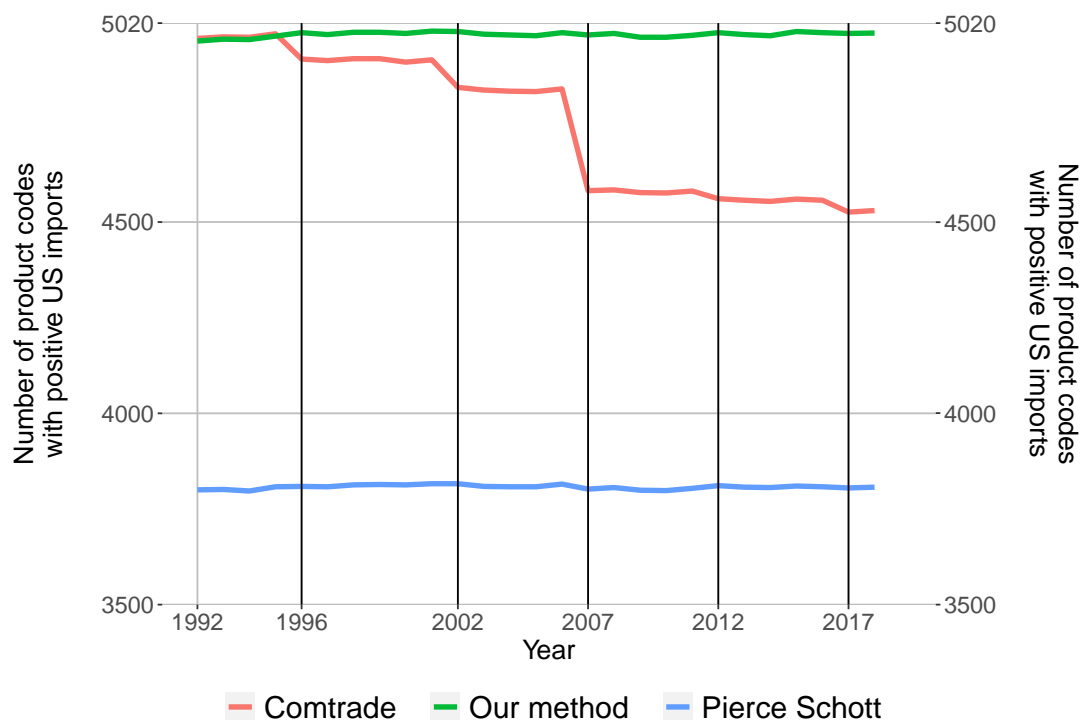
This in turn leads to different interpretations when looking at recent trends of U.S. imported product varieties after the Great Trade Collapse (cf. [Figure B.6](#) in [Appendix B](#)). Following [Broda and Weinstein \(2006\)](#), we define a variety as a product imported from a specific origin country. While data based on our method imply a relatively swift return to pre-crisis levels, data based on the Comtrade method would suggest that the number of imported product varieties did not recover until 2015.

Meanwhile, applying the [Pierce and Schott \(2012\)](#) method leads to nearly a quarter less of products (and thus product varieties). In terms of recovery patterns, the trends captured by the [Pierce and Schott \(2012\)](#) data resemble those suggested by our method, as both versions show a strong post-crisis upturn. This pattern is also in line with aggregate U.S. imports (depicted as bars in [Figure B.6](#)), which had recovered to pre-crisis values by 2011.

Supply-side trends. While the example above is meant to illustrate the benefits of our conversion method to study import demand trends, similar statements may be made on the supply side. For instance, [Figure 6](#) investigates entry and exit rates of product-

¹⁶In fact, the aggregate price index by [Broda and Weinstein \(2006\)](#) accounts for the possibility of codes being merged and split depending on whether their importance in trade decreases or increases. They do not account, however, for trade being reclassified due to other reasons (cf. [Appendix A](#)).

Figure 5 – Number of Products Imported by U.S. according to Different Conversion Methods



Notes: This figure plots for each year the number of 6-digit HS products (in the case of Pierce Schott: synthetic codes) imported by the United States. The three lines depict the different numbers of products based on the Comtrade, the [Pierce and Schott \(2012\)](#), and our conversion method; with the first and the last method used to convert all codes back to the HS 1992 vintage (containing 5,020 codes).

Data source: [UN Comtrade Database](#).


exporter combinations for 71 countries which reported data between 1996 and 2017. For these countries, we find an unwarranted jump nearly doubling the exit rate in the year 2007 when applying the Comtrade method, in line with findings of [Cebeci \(2012\)](#). Similar artificial jumps – while less pronounced – occur in each of the other years facing an HS code amendment. Otherwise, all three methods seem very much aligned with respect to the exit rate statistics.^{17,18}

With regard to the entry rate, all three methods seem more aligned; with the notable exception of using our method in 2007. The jump may indicate that our method generates “too many” positive product flows in that year. However, this increase – as well those in 1996 and 2002 – are not as pronounced as the ones noticeable for the Comtrade method regarding exit rates.

¹⁷Note that not all 71 countries each time adjusted the HS vintage in the same year, leading to the fact that the HS-conversion-related jumps may be lagged for some countries. Reducing the pool to the 35 complying countries – i.e., those which each year applied the latest available HS vintage – further sharpens the jumps seen in the vintage adjustment years. However, the overall picture does not change.

¹⁸The Pierce-Schott method results in very similar entry and exit rates as the two other methods (despite having a quarter fewer products) largely due to the share of unchanged HS codes. Restricting the products in Figure B.6 to those HS codes that change at least once results in the entry and exit rates of the Pierce-Schott method being roughly three percentage points lower than those based on the other approaches.

Figure 6 – Entry and Exit Rates (in Percent) for Product-Exporter Combinations using the Different Conversion Methods



Notes: This figure plots for each year the entry and exit rates (in percent) for product-exporter combinations. Products are defined as 6-digit HS codes (in the case of Pierce Schott: synthetic codes) and observations are restricted to those 71 countries which reported trade data for the years 1995 through 2017. The three lines depict the different rates based on the Comtrade, the [Pierce and Schott \(2012\)](#), and our conversion method; with the first and the last method used to convert all codes back to the HS 1992 vintage. The large 2000 increase in the entry rate seen in all three versions relates to a large wave of trade liberalisation, in particular to China.

Data source: [UN Comtrade Database](#).

6 Heterogeneity and Robustness of Weights

This section presents variations of our main specification outlined in [Section 3](#).

Product-specific trends. One potential concern is that we do not account for product-specific trends in the data. In order to tackle the issue of product categories potentially becoming more or less relevant over time, we can estimate product-specific autoregressive coefficients which are then used to augment our minimization problem. Specifically, using all years in which the “initial” HS vintage was the latest available (i.e., the years where this specific classification standard was officially in place), let us estimate the following regression

$$x_k^t = \rho_k x_k^{t-1} + \varepsilon_k^t,$$

where x_k^t is the (aggregate) trade flow of product k in year t , ρ_k is the product-specific autoregressive coefficient, and ε_k^t is an error term. We can then assign more weight to flows in t_0 where we measured a higher persistence over time. The objective function for

both $t_0 < t_1$ or $t_0 > t_1$ is illustrated in Equation (2).

$$\sum_s \sum_{i,j} \left(x_{ij,s}^{t_1} - \sum_k (x_{ij,k}^{t_0} \hat{\rho}_k^{t_1-t_0}) \gamma_{k,s} \right)^2, \quad (2)$$

where $\gamma_{k,s}$ is the new conversion weight.¹⁹

When applying the adjusted objective function, our results remain largely the same (cf. Figure B.7). Overall, we see very few weights receiving substantially different values. The majority of conversion weights remain very similar, as the correlation coefficient for not a priori determined weights is 0.97.

Multiple years. A further limitation might be the use of only one initial and one target year. We therefore re-estimate the conversion weights by using average product-level bilateral flows over the “initial” and “target” periods, each scaled by their respective total group flows. This averaging may be particularly important for products ordered irregularly such as airplanes or trains, as well as for goods subject to multi-annual seasonal patterns, especially in agriculture. Nonetheless, when applying the algorithm on conversion weights for the vintages HS 2007 to HS 2002, we obtain very similar results between averaged and single-year trade data. As Figure B.8 shows for averaged (2007–2009 and 2004–2006) and single-year (2007 and 2006) data, most weights are close to the 45-degree line, with a correlation coefficient of 0.96.

Sampling. To further verify the robustness of our weights, we re-estimate them using subsamples of our data. More precisely, we randomly divide the data into 20 equi-sized bins and look at the distribution of weights after running the regressions twenty times, where each regression leaves out one block (or, five percent of the observations). The conversion weights appear to be fairly stable. For instance, when we convert HS 2007 to HS 2002, the difference between the maximum and minimum estimated coefficient for a

¹⁹In order for trade shares to still sum up to one in year t_0 , we rescale the flows by total trade within groups after multiplying them with the product-specific autoregressive coefficients. Also, note that this reasoning can be extended to higher-order autoregressive processes. For instance, let \underline{t}_0 and \bar{t}_0 be the earliest and latest year, respectively, in which the “initial” reporting standard was mostly applied. Further, let T_0 be the set containing these two years as well as those in between. We can then estimate

$$x_{ij,k}^{\bar{t}_0} = \sum_{z=\underline{t}_0}^{\bar{t}_0-1} \rho_{k,(\bar{t}_0-z)} x_{ij,k}^z + \nu_{ij,k}^{\bar{t}_0},$$

where $\rho_{k,l}$ is the coefficient on the l^{th} lag, and $\nu_{ij,k}^{\bar{t}_0}$ is an error term. Correspondingly, if, e.g., $t_1 - \bar{t}_0 = 1$, the adjusted objective function looks as follows

$$\sum_s \sum_{i,j} \left(x_{ij,s}^{t_1} - \sum_k \left(\sum_{z \in T_0 \setminus \{\underline{t}_0\}} x_{ij,k}^z \hat{\rho}_{k,(t_1-z)} \right) \tilde{\gamma}_{k,s} \right)^2.$$

given initial and target product pair rarely surpasses 0.1, as illustrated in Figure B.9.²⁰ Moreover, in order to see whether large fluctuations happen for products which are relatively important, in Figure B.10 we plot this difference against the multiplication of the relative size of the initial and target products in their groups. The northeastern area of the graph depicts relatively important links which feature rather unstable conversion weights. As can be seen in Figure B.10, that area is almost empty. Results for other vintages look very similar.

Overall, this exercise suggests that our weights are generally not driven by only a few country pairs, and that if they are, it mostly concerns relatively unimportant products. Also, note that taking the difference between the maximum and minimum coefficients is the most conservative test we may choose.

Aggregation level. As discussed in Section 3, bilateral product-level flows tend to be persistent over time, although not perfectly so. Product-level imports, on the other hand, tend to fluctuate less over time for a given country, probably because a country’s demand structure changes more sluggishly than bilateral transactions at the product level. Hence, one may argue that using importer-level rather than country pair-level data – i.e., aggregating our input variables $x_{ij,k}$ across exporters $i \in I$ – better suits our optimization routine, despite reducing the number of observations substantially. In a similar vein as in the subsection above, we thus look at how conversion weights estimated using importer-level data compare with our baseline results. In Figure B.11 we do this for the conversion from 2007 to 2006 data (or, from HS 2007 to HS 2002).²¹ Most absolute differences between the coefficients are again below 0.1. Furthermore, in order to assess whether large fluctuations appear for products which are relatively important, in Figure B.12 we scale the coefficients by the relative size of the initial and target products in their groups, and plot them against each other. The outcome reveals that if there are large differences, we find them only for relatively small products. Results for other vintages look very similar.

Given that, although the results are similar, the importer-level data tend to estimate more coefficients (very) close to zero, which is likely due to the lower number of observations, we stick to the country pair-level data for our baseline results.

Regions. Our calculations thus far were focused on global trade patterns. As countries differ in their product mixes though, the estimated conversion weights may vary across jurisdictions. In Figure B.13 we compare how the conversion weights calculated using the

²⁰This figure excludes pre-specified weights (i.e., conversions known to be equal to zero or one), as these would, by definition, always yield a difference of zero. Also, keeping in mind that about 85% of the groups in the conversion from HS 2007 to HS 2002 have less than five weights which are not pre-specified, a difference of less than ten percentage points can be considered small.

²¹As above, this figure excludes pre-specified weights (i.e., conversions known to be equal to zero or one), as these would, by definition, always yield a difference of zero.

entire sample perform relative to country-group-specific weights. Panel B.13a compares the 2007 trade shares converted into HS 2002 using country-group-specific weights with the 2006 shares for both OECD and non-OECD countries. In panel B.13b we convert the 2007 figures using our “global” conversion weights (i.e., those we used in Figure 3). Intuitively, the country-group-specific conversion provides a better fit for both country groups, although the global weights also work fairly well. It should be expected that the estimates from the entire sample better suit the OECD countries, since most large pairs belong to that country group.

7 Conclusion

Product-level trade data often suffer from intertemporal inconsistencies due to classification changes. Existing methods to harmonize HS vintages either drop product codes over time, suggesting, e.g., a 10% decline in the number of products imported by the U.S. over the past three decades. Or, they bulk amended codes into synthetic categories, where the largest one covers a quarter of global goods trade.

This paper presents an approach to tackle the issues above by estimating conversion weights between linked HS codes. The weights may be readily applied to existing trade data and combined for multiple vintage changes. The algorithm can produce region- or country-group-specific conversions (e.g., by focusing on OECD vs. non-OECD countries). However, robustness checks indicate that the weights are not significantly affected by the set of year and country pair observations we include.

Whilst we focus on HS codes, the algorithm could be applied to other classification updates (e.g., NACE, NAICS, SITC). Our method has two requirements. First, it needs a clear correspondence between different vintages. Second, the underlying data have to be persistent, i.e., there should not be (unexpected) jumps in the data in years with classification changes.

Our proposed conversion method allows for a more sensible descriptive investigation of trade patterns. When looking at imported product varieties (defined as in Broda and Weinstein (2006)), data converted using our algorithm do not feature large jumps in amendment years and are in line with general trade developments. The same holds for the analysis of exporter entry and exit rates. This relates to a growing literature investigating firm-level exporter dynamics (cf. Alessandria et al. (2021) for an overview), in the sense that we should be able to compare the trends in micro and macro data.

References

- Alessandria, G., Arkolakis, C., and Ruhl, K. J. (2021). Firm Dynamics and Trade. *Annual Review of Economics*, 13:253–280.
- Bellert, N. and Fauceglia, D. (2019). A Practical Routine to Harmonize Product Classifications over Time. *International Economics*, 160:84–89.
- Boehm, C. E., Levchenko, A. A., and Pandalai-Nayar, N. (2020). The Long and Short (Run) of Trade Elasticities. *NBER Working Paper Series*, Working Paper 27064.
- Broda, C. and Weinstein, D. E. (2006). Globalization and the Gains from Variety. *Quarterly Journal of Economics*, 121(2):541–585.
- Cebeci, T. (2012). A Concordance among Harmonized System 1996, 2002 and 2007 Classifications. *World Bank Working Papers*, No. 74576. <http://documents.worldbank.org/curated/en/914591468335500796/A-concordance-among-harmonized-system-1996-2002-and-2007-classifications>.
- Diodato, D. (2018). A Network-based Method to Harmonize Data Classifications. *Papers in Evolutionary Economic Geography*. No 18.43.
- Felbermayr, G., Teti, F., and Yalcin, E. (2019). Rules of origin and the profitability of trade deflection. *Journal of International Economics*, 121:103–248.
- Fort, T. C., Pierce, J. R., and Schott, P. K. (2018). New Perspectives on the Decline of U.S. Manufacturing Employment. *Journal of Economic Perspectives*, 32(2):47–72.
- Head, K. and Mayer, T. (2014). Gravity Equations: Workhorse, Toolkit, and Cookbook. *Handbook of International Economics*, 4:131–195.
- Hsieh, C.-T., Li, N., Ossa, R., and Yang, M.-J. (2020). Accounting for the New Gains from Trade Liberalization. *Journal of International Economics*, 127:103370.
- Pierce, J. R. and Schott, P. K. (2012). Concording U.S. Harmonized System Codes over Time. *Journal of Official Statistics*, 28(1):53–68.
- Piveteau, P. and Smagghue, G. (2019). Estimating Firm Product Quality Using Trade Data. *Journal of International Economics*, 118:217–232.
- UNSD (2004). *International Merchandise Trade Statistics Compilers Manual*. United Nations, New York.
- UNSD (2017). Correlation and Conversion Tables Used in UN Comtrade. <https://unstats.un.org/unsd/trade/classifications/correlation/HS2017%20conversion%20to%20earlier%20HS%20versions%20and%20other%20classifications.pdf>, accessed on February 20, 2022.
- Wagner, R. and Zahler, A. (2015). New Exports from Emerging Markets: Do Followers Benefit from Pioneers? *Journal of Development Economics*, 114:203–223.
- WCO (2015). Correlation Tables HS 2012–2017. <http://www.wcoomd.org/en/topics/nomenclature/instrument-and-tools/hs-nomenclature-2017-edition/correlation-tables-hs-2012-to-2017.aspx>. accessed on February 20, 2020.

Appendix

A The Harmonized System Codification

The classification of products into Harmonized System (HS) codes is administered by the World Customs Organisation through the HS Convention signed in 1983. Each of the 150+ signatory countries is obliged to align its tariff schedules with the HS code classification. As the HS code system at its most disaggregated level reaches 6 digits, individual countries can expand their tariff schedule to a more detailed classification—at the 8-, 10-, or even 12-digit level—as long as those codes adhere at the 6-digit level to the internationally agreed standard (art.3(3) of the HS Convention).

Amendments to HS codes are proposed by the Harmonized System Committee, consisting of representatives from each contracting country. This Committee shall “propose amendments [to the current HS classification system] as may be considered desirable, having regard, in particular, to the needs of users and to changes in technology or in patterns of international trade” (art.7(1)a of the HS Convention). The proposed changes are then recommended to all contracting parties of the HS Convention and—provided no objections are received within 6 months—enter into force on the 1st of January of the following year (art.16 of the HS Convention).

Changes to the HS code classification may also be due to other reasons than those explicitly listed in the HS Convention. According to the Compilers Manual of the International Merchandise Trade Statistics (2004), some of the HS code changes in 2002 reflected “user needs but many of which were made at the request of other international organizations seeking to be able to identify trade in sensitive goods, including other hazardous materials” (p.32).

The latest HS code vintage to be adopted (HS 2022) will *inter alia* provide new classifications to address environmental issues (better tracking of e-waste), health concerns (separate codes for fentanyl), public security (clearer classification of components for explosives), or to facilitate international medical research (by reclassifying placebo drugs without having to differentiate by their ingredients). However, the 2022 amendments also take into account technological progress through separate HS codes for smartphones, drones, and a reclassification of the glass fiber products.

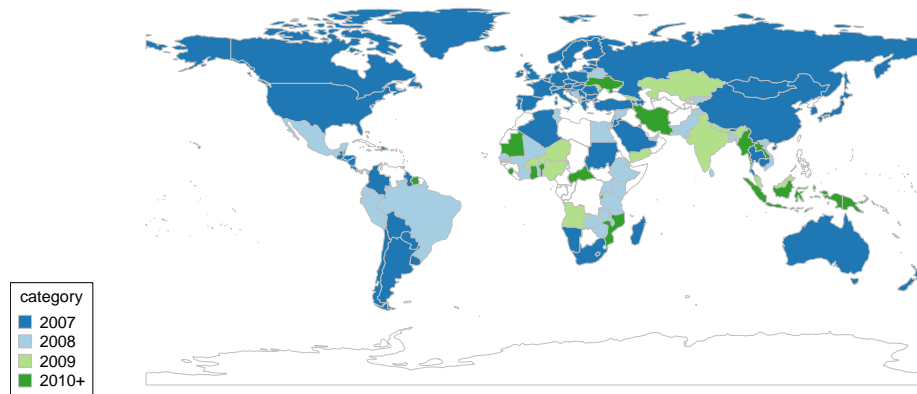
The algorithm in this paper is based on the correlation tables provided by the World Customs Organisation (and published by UN Comtrade²²). Despite these tables being “examined by the Harmonized System Committee, they are not to be regarded as constituting classification decisions taken by that Committee; they constitute a guide published by the Secretariat and whose sole purpose is to facilitate implementation of the 2017 edition of the Harmonized System” (WCO, 2015). The correlation tables are drawn up in

²²The tables can be obtained here: <https://unstats.un.org/unsd/trade/classifications/correspondence-tables.asp>, accessed on February 1, 2022.

a way that is “as comprehensive as possible” (WCO, 2015)—it is therefore not required for a majority of country members of the HS Committee to agree on every link between a given old and new code. The final correspondence tables published by the Secretariat may therefore result in being too broad, i.e., creating too many links. This characteristic further supports our approach of letting trade data “speak for itself” in terms of which HS codes have a quantitative correspondence.

B Additional Figures

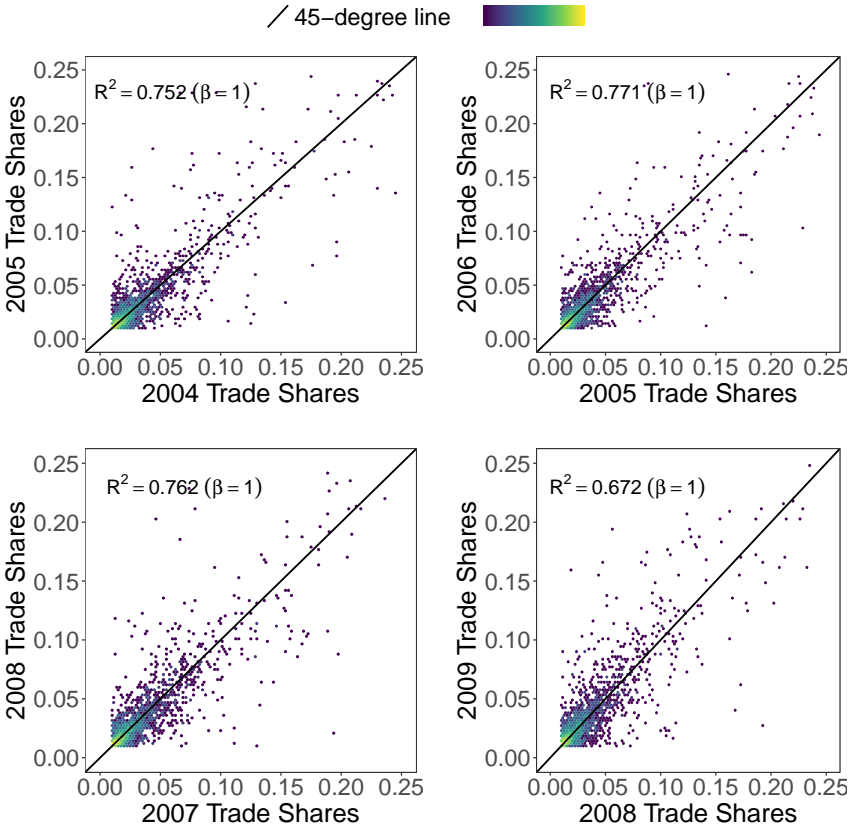
Figure B.1 – Year of HS 2007 Vintage Adoption by Countries Worldwide



Notes: This figure plots for each country the year in which it first adopted the HS 2007 code vintage. Note that some countries never adopted the HS 2007 classification (marked by white filling).

Data source: [UN Comtrade Database](#).

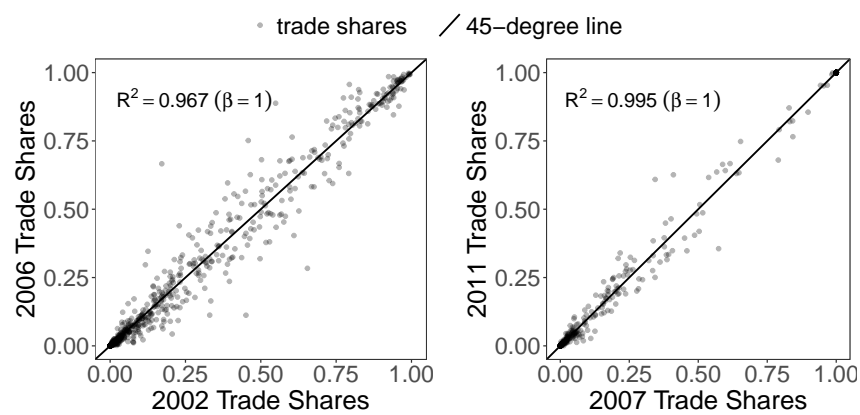
Figure B.2 – Bilateral Product-Level Trade Shares within HS Product Groups over Time



Notes: This figure plots for each bilateral product-level flow its trade share within its specific group across several years. We include only groups which contain at least one 1:n or m:n assignment from HS 2007 to HS 2002. The figure uses Comtrade data provided in the original HS code version and is restricted to 70 importers reporting in the latest available HS code version in a given year. Due to the volume of observations, the data is plotted in 100 bins (per axis), and a lighter color is applied to areas with more observations. Product groups are explained in the main text. The provided R-squared assumes a fit along the 45-degree line. For illustrative purposes, the axes are capped at 0.25, and observations are excluded for which the trade share was below 0.01 in both years.

Data source: [UN Comtrade Database](#).

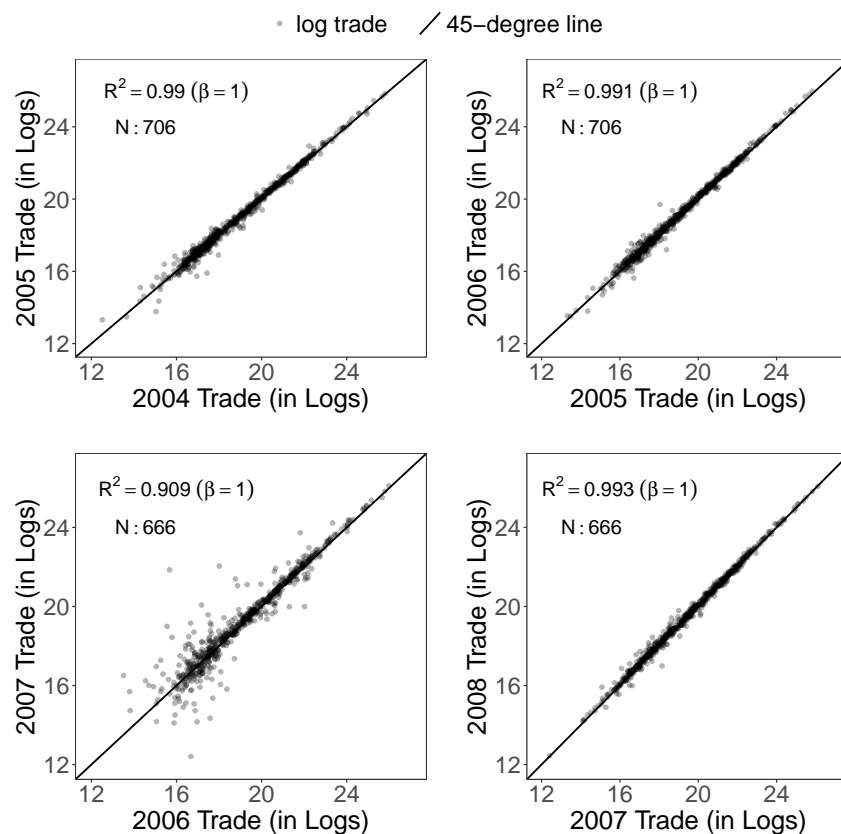
Figure B.3 – Product Trade Shares within HS Product Groups (HS Codes as Reported)



Notes: This figure plots for each product its trade share within its specific group across several years. We include only groups which contain at least one 1:n or m:n assignment from HS 2007 to HS 2002, and trade data for 70 countries which used the latest available HS code classification in each of the years used in the plot. The figure uses Comtrade data provided in the latest available HS vintage. Product groups are explained in the main text. The provided R-squared assumes a fit along the 45-degree line.

Data source: [UN Comtrade Database](#).


Figure B.4 – Log of Product-Level Trade over Time (Our Conversion Weights)



Notes: This figure plots for each product its log of overall trade across several years. We include only products in groups which contain at least one 1:n or m:n assignment from HS 2007 to HS 2002, and trade data for 70 countries which used the latest available HS code classification in each year between 2004 and 2008. The figure uses Comtrade data provided in the latest available HS vintage, and transforms 2007 and 2008 data using our weights back into HS 2002. Product groups and the algorithm to obtain our conversion weights are explained in the main text. The provided R-squared assumes a fit along the 45-degree line.

Data source: [UN Comtrade Database](#).

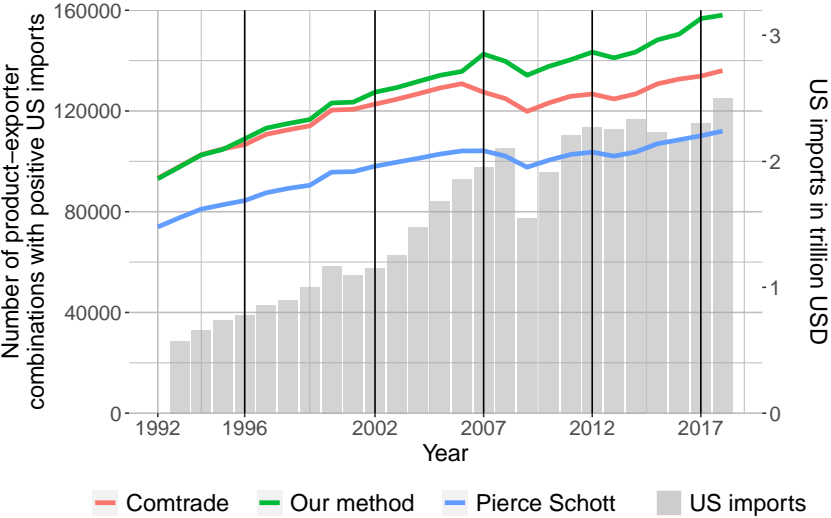
Figure B.5 – Log of Product-Level Trade over Time (Our Conversion Weights – HS 1992)



Notes: This figure plots for each product its log of overall trade across several years. We include only products which had at least one non-1:1 assignment over the period 1995–2017, and trade data for 38 countries which used the latest available HS code classification in each year between 1995 and 2017. The figure uses Comtrade data provided in the latest available HS vintage, and transforms all years after 1995 using our weights back into HS 1992. Product groups and the algorithm to obtain our conversion weights are explained in the main text. The provided R-squared assumes a fit along the 45-degree line.

Data source: [UN Comtrade Database](#).


Figure B.6 – Number of Varieties Imported by the U.S. according to Different Conversion Methods



Notes: This figure plots for each year the number of varieties imported by the United States. A variety is defined as an exporter-product combination (in the case of Pierce Schott: synthetic product codes). The three lines depict the different number of products based on the Comtrade, the [Pierce and Schott \(2012\)](#), and our conversion method; with the first and the last method used to convert all codes back to the HS 1992 vintage. The bars show total U.S. imports in trillion current USD.

Data source: [UN Comtrade Database](#).


Figure B.7 – Conversion Weights for HS 2007 to HS 2002 with and without Autocorrelation Adjustment



Notes: This figure plots every not a priori determined weight for converting HS 2007 to HS 2002, once estimated while adjusting for product-specific trends (cf. Equation (2)) and once without trend adjustments (cf. Equation (1)).

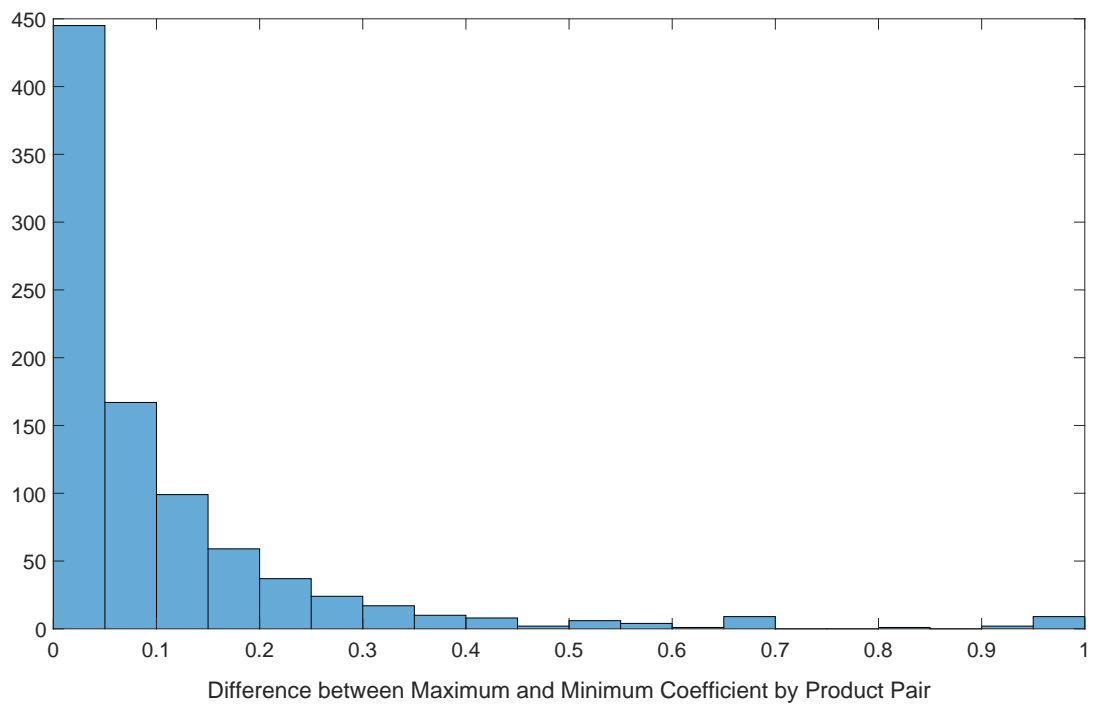
Data source: UN Comtrade Database.

Figure B.8 – Conversion Weights for HS 2007 to HS 2002 using Averaged vs. Single-Year Trade Data



Notes: This figure plots every not a priori determined weight converting HS 2007 to HS 2002 when using trade data for only 2006 and 2007 against using averaged trade for the years 2004–2006 and 2007–2009.
Data source: [UN Comtrade Database](#).

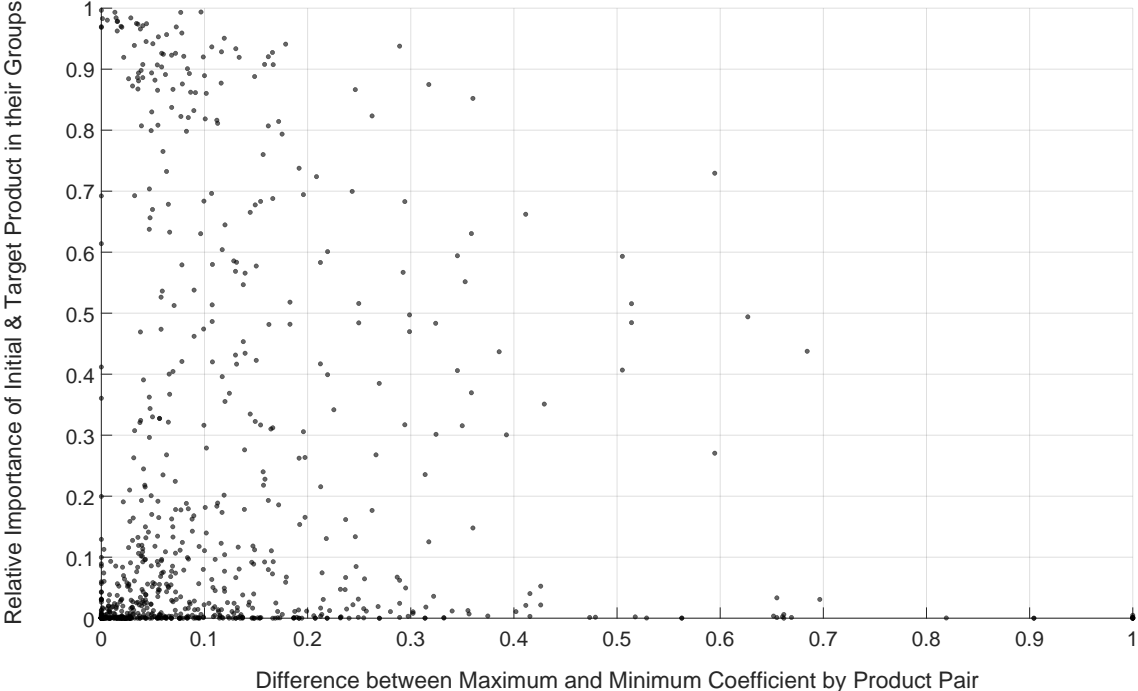
Figure B.9 – Sampling: Maximum Fluctuations of Conversion Weights



Notes: This figure plots the distribution of maximum differences in estimated conversion weights across bins (cf. Section 6). The initial year is 2007 and the target year is 2006 (i.e., converting HS 2007 to HS 2002). The data are randomly split into 20 blocks, and twenty regressions are computed, each leaving out one block. The difference we plot is the difference between the largest and smallest coefficient we get for each product pair. This figure excludes pre-specified weights (i.e., conversions known to be equal to zero or one), as these would, by definition, always yield a difference of zero.

Data source: [UN Comtrade Database](#).

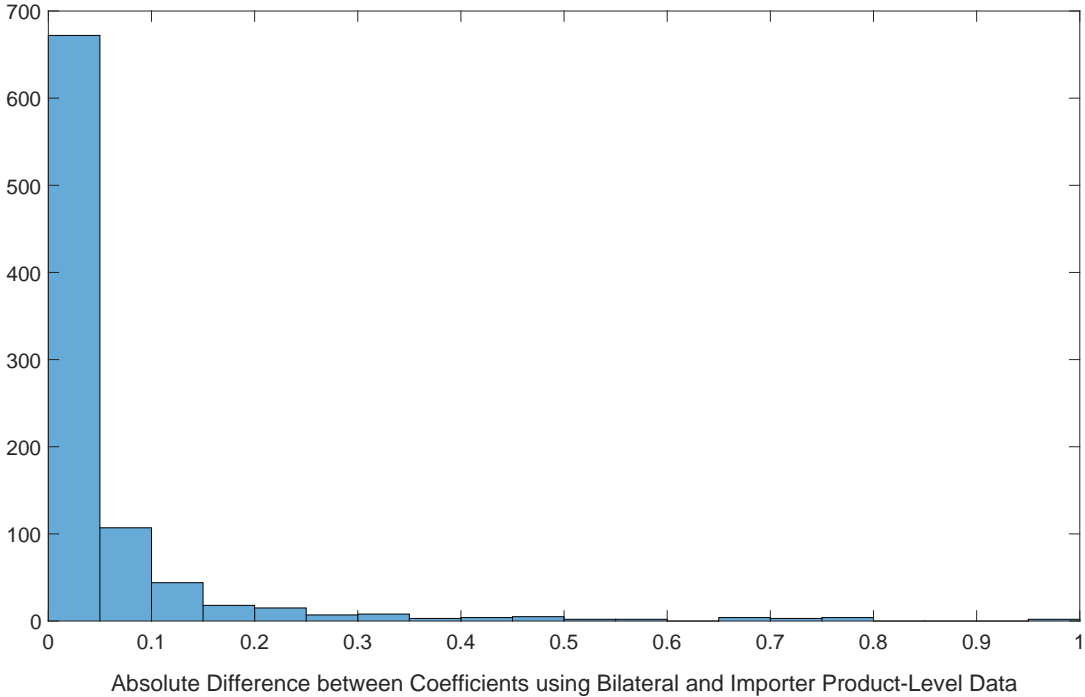
Figure B.10 – Sampling: Maximum Fluctuations of Weights and Relative Importance of Products



Notes: This figure plots the maximum difference in estimated conversion weights across bins against the relative importance of initial and target products in their groups (cf. Section 6). The initial year is 2007 and the target year is 2006 (i.e., converting HS 2007 to HS 2002). The data are randomly split into 20 blocks, and twenty regressions are computed, each leaving out one block. The difference we plot is the difference between the largest and smallest coefficient we get for each product pair. The relative importance is measured as the multiplication of the initial and target product trade shares within their groups. This figure excludes pre-specified weights (i.e., conversions known to be equal to zero or one), as these would, by definition, always yield a difference of zero.

Data source: [UN Comtrade Database](#).

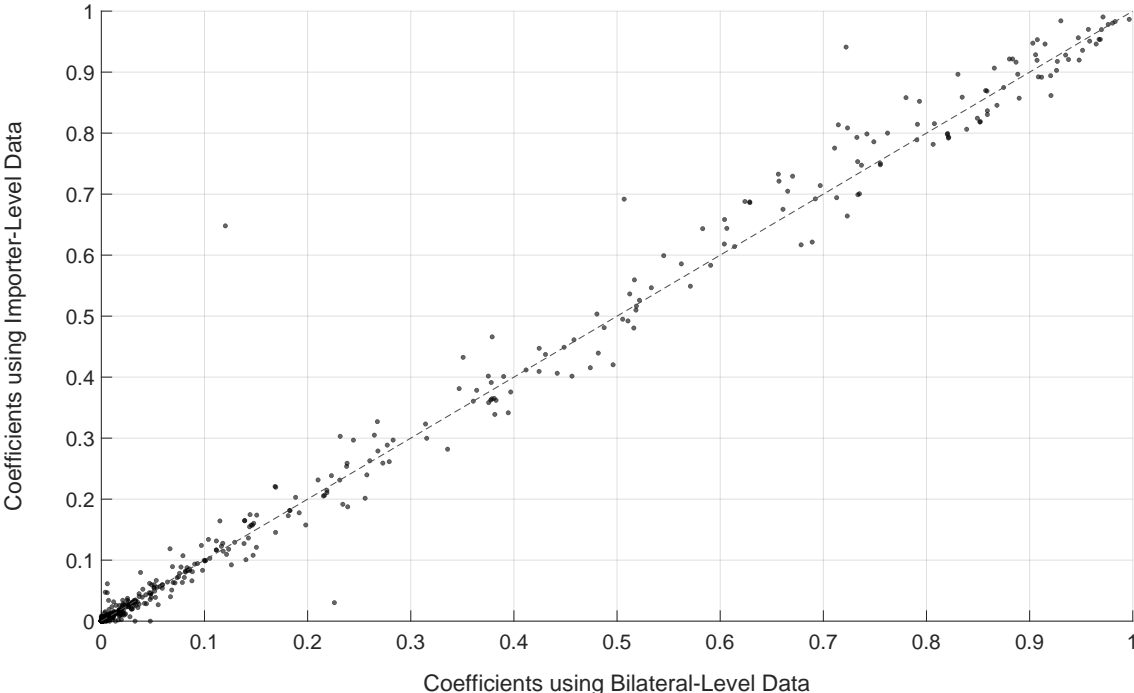
Figure B.11 – Aggregation: Absolute Differences in Coefficients Estimated Using Bilateral vs. Importer Product-Level Data



Notes: This figure plots the distribution of absolute differences between conversion weights estimated using bilateral and importer product-level data, respectively (cf. Section 6). The initial year is 2007 and the target year is 2006 (i.e., converting HS 2007 to HS 2002). This figure excludes pre-specified weights (i.e., conversions known to be equal to zero or one), as these would, by definition, always yield a difference of zero.

Data source: UN Comtrade Database.

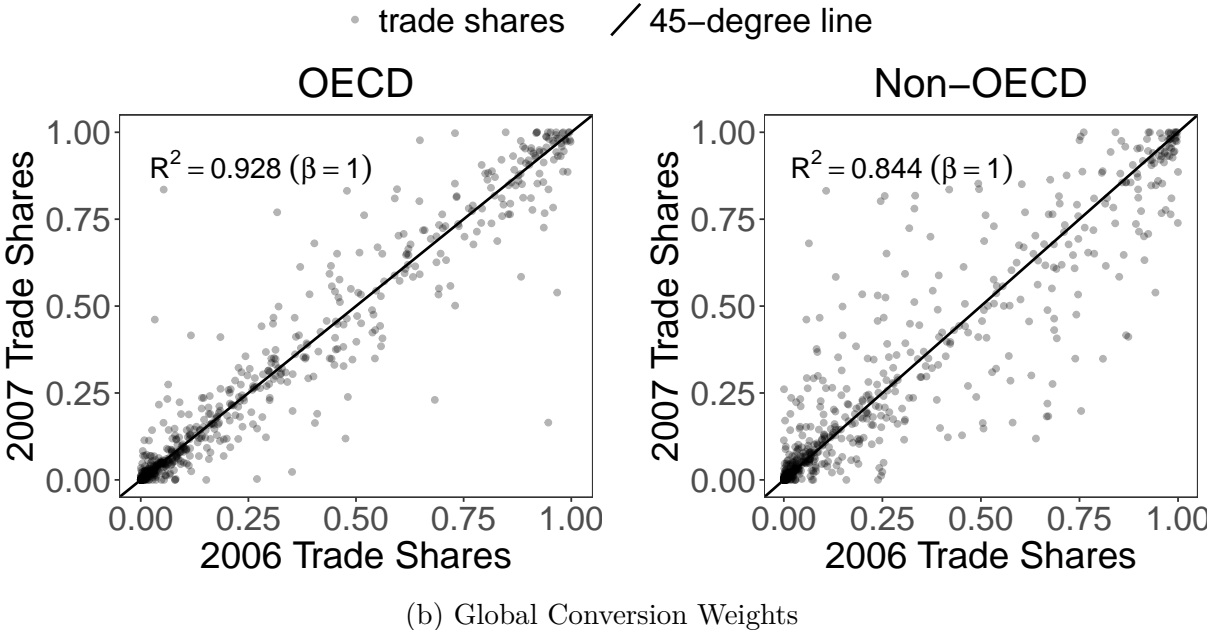
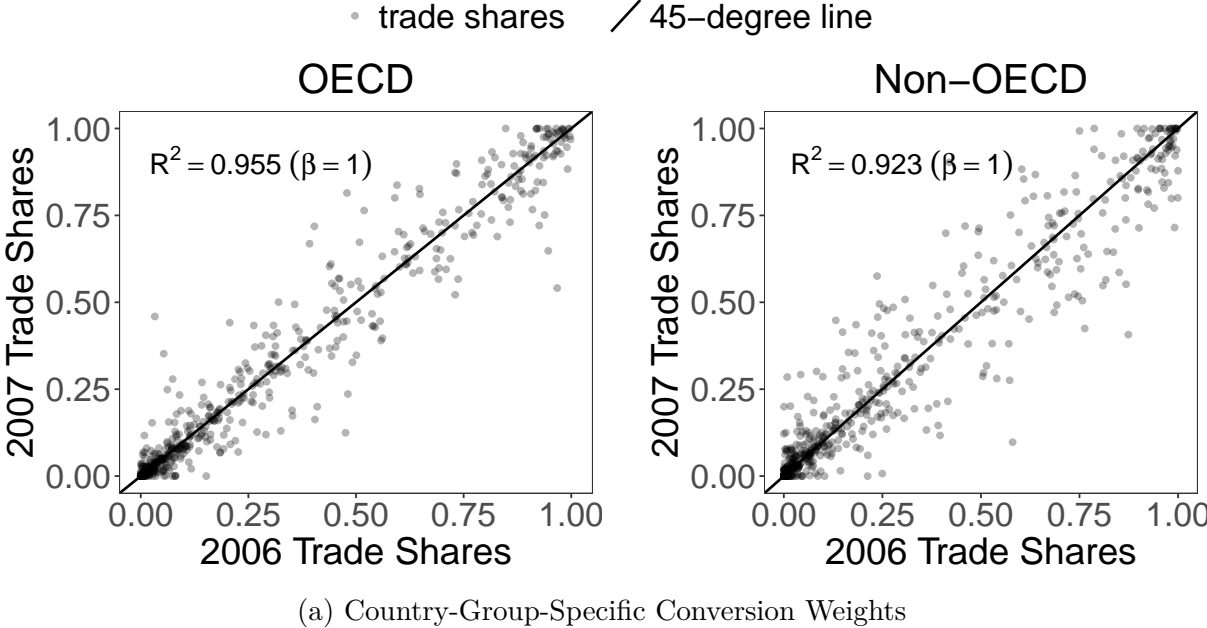
Figure B.12 – Aggregation: Coefficients Estimated Using Bilateral vs. Importer Product-Level Data Scaled by Relative Importance of Products



Notes: This figure plots the conversion weights estimated using bilateral and importer product-level data, respectively, scaled by the relative importance of initial and target products in their groups (cf. Section 6). The initial year is 2007 and the target year is 2006 (i.e., converting HS 2007 to HS 2002). The relative importance is measured as the multiplication of the initial and target product trade shares within their groups. The dashed line depicts the 45-degree line. This figure excludes pre-specified weights (i.e., conversions known to be equal to zero or one), as these would, by definition, always yield a difference of zero.

Data source: UN Comtrade Database.

Figure B.13 – Product Trade Shares within HS Product Groups over Time (Our Conversion Weights – Country-Group-Specific vs. Global Weights)



Notes: This figure plots for each product its trade share within its specific group across several years. We include only product groups which contain at least one 1:n or m:n assignment from HS 2007 to HS 2002, and trade data for 36 OECD (left-hand side) and 41 non-OECD (right-hand side) countries which used the latest available HS code classification in 2006 and 2007. The figure uses Comtrade data provided in the latest available HS vintage, and transforms 2007 data using our weights back into HS 2002. In panel B.13a we use OECD- and non-OECD-specific conversion weights, respectively, while panel B.13b uses the weights we calculated with the entire sample. Product groups and the algorithm to obtain our conversion weights are explained in the main text. The provided R-squared assumes a fit along the 45-degree line.

Data source: UN Comtrade Database.