Missing Tariffs

Feodora A. Teti*

June 11, 2024

Abstract

Many studies use tariff data to measure changes in trade policy. This paper reveals significant errors due to incomplete reporting in one of the standard sources for tariff rates, the World Integrated Trade Solution (WITS). I propose a new interpolation algorithm that takes the misreporting into account, and I use the resulting global tariff dataset to demonstrate the practical relevance of missing tariffs by replicating recent studies that rely on WITS data. I find that cross-country studies leveraging preferential tariff variation such as the estimates of the trade elasticity are particularly sensitive to the mishandling of the tariff data.

Keywords: Tariffs, MFN, Preferences, Trade Elasticity, WITS

JEL-Classification: F13, F14

*ifo Institute, University of Munich, CEPR, and CESifo (e-mail: teti@ifo.de).

I am very thankful to Costas Arkolakis, Treb Allen, Andrew Bernard, Matilde Bombardini, Arnaud Costinot, Carsten Eckel, Gabriel Felbermayr, Lisandra Flach, Alejandro Graziano, Gene Grossman, Dennis Novy, Gianluca Orefice, Monika Schnitzer, Claudia Steinwender, Felix Tintelnot, and Joachim Winter, as well as participants of numerous seminars and conferences for their valuable comments and suggestions. I gratefully acknowledge financial support received from Senatsausschuss Wettbewerb (SAW) under grant no. SAW-2016-ifo-4, the German Research Foundation through CRC TRR 190, and the scholarship from the Bernt Rohrer Foundation. All errors are my own.

1 Introduction

Many studies rely on tariffs to measure changes in trade policy, not only in international economics but also in other fields such as innovation, labor, inequality, and development (cf. Goldberg and Pavcnik 2016; A. Harrison and Rodriguez-Clare 2010; Shu and Steinwender 2019). Tariffs are often used as measures of trade costs because they vary over time and sectors and—in contrast to most nontariff barriers—are quantifiable and easy to compare across countries. They have been leveraged to identify the trade elasticity, which is considered to be *"the single most important structural parameter in gravity models"* (Costinot and Rodríguez-Clare 2014, p. 241), as it determines the gains from trade.[1] Tariffs are also a crucial control variable in many other investigations of research questions in trade such as work that focuses on product quality (Amiti and Khandelwal 2013), the environment (Shapiro 2021), labor markets (Greenland and Lopresti 2022), and nontariff barriers (De Melo and Nicita 2018; Kee, Neagu, et al. 2013b).

This paper shows that one of the standard source of tariff data used in the literature suffers from substantial measurement error and selection bias due to missing tariffs. Many countries fail to report their tariff rates every year, with preferential tariffs more likely than most favored nation (MFN) tariffs to go unreported. The World Bank's World Integrated Trade Solution (WITS), the main provider of tariff data, falsely interpolates the missing preferential tariff rates with MFN rates. This practice leads to artificial spikes in bilateral time series data and hence induces massive measurement error. Additionally, tariff data are available only for importers that report trade flows to UN Comtrade. Therefore, low-income countries are underrepresented in the resulting sample, and tariffs for products with zero trade flows are discarded. Hence, the WITS tariff data are plagued by two issues, false interpolation and positive selection, making them unsuitable for empirical research. To illustrate the practical importance of missing tariffs I replicate papers that were recently published in leading economics journals and that use the data provided by WITS. I find that existing estimates of the trade elasticity are particularly sensitive to the mishandling of the tariff data.

In view of the issues present in the tariff data available through WITS, the question arises: how widely have these data been used in empirical work? In total, I identify 125 papers using WITS data for analysis that have been published in prominent journals since 2010, out of which 21 have been published in leading journals often referred to as *"top 5"*.[2] Some authors successfully address the problem of missing tariffs head on, typically when the focus is on a

[1] The trade elasticity is the elasticity of bilateral imports with respect to variable trade costs.

[2] I consider publications in the *American Economic Review*, *Quarterly Journal of Economics*, *Journal of Political Economy*, *Econometrica*, and *Review of Economic Studies* as the leading journals or *top 5*. Furthermore, I also included publications in the *Review of Economics and Statistics*, *AEJ: Economic Policy*, *AEJ: Microeconomics*,

specific policy context and the flaws in the data series are therefore apparent (i.e., Conconi et al. 2018; Head and Mayer 2019). However, the difficulty of spotting errors in the WITS tariff data varies significantly with the exact context. Especially in cross-country studies, the sheer size and complexity of the data sets—each tariff observation has four dimensions, i.e., importer, exporter, product, and time—generates many possibilities along which error can be embedded in the data, particularly if researchers do not know where to look and what the characteristics of the errors are.

To address the unsatisfactory data situation with respect to missing tariffs, I build a new global tariff dataset (GTD) that accounts explicitly for the issues due to missing tariffs. I proceed in three steps: First, I minimize the number of missing observations by combining five different sources of tariff data. In addition to including information provided by UNCTAD, the WTO, the International Trade Center (ITC), and national authorities, I have collected the phase-in schedules for 149 free-trade agreements (FTAs) that provide flawless information for the respective FTAs. In the second step, I propose a new algorithm to fill in the observations that are still missing. For MFN tariffs, I use the first preceding nonmissing tariff. If no data from earlier years are available, I use the first subsequent observation. For preferential tariffs, the imputation is performed identically, conditional on the RTA's being in force. Furthermore, I account for phased-out tariffs using information on the phase-out regime. The Design of Trade Agreements (DESTA) database (Dür et al. 2014) has information for more than 600 RTAs on the respective phase-out regime. This information is combined with the panel of tariffs to interpolate the preferential tariff rates as closely as possible to the real rates. In the last step, I retrieve the effectively applied tariff using the fully filled-in MFN and preferential tariff rates for all importer–exporter–product observations. I end up with 5.7 billion tariffs covering 197 countries over 30 years at the HS6 level.

In what direction should we expect the bias from the missing tariffs to go? For the example of estimating the tariff elasticity, i.e., regressing trade on tariffs, I will show that there are three competing forces and their interplay will determine the bias. The false interpolation introduces nonclassical measurement error: First, the falsely interpolated tariff rates always overstate the true tariff rates, which will unequivocally attenuate the tariff elasticity. Second, the measurement error is present only for country pairs with a trade agreement. However, country pairs with a trade agreement are very different from those without: for example, they might be geographically closer to each other (cf. Baier and Bergstrand 2004). Hence, geography is a determinant of both trade and tariffs causing endogeneity. This *differential* nature of the measurement error will hence act as a counteracting force to the attenuation bias.

Macroeconomics, *AEJ: Applied Economics*, *Journal of the European Economic Association*, and the *Journals of International Economics*.

The third element of the total bias is attributable to the sample selection. Bilateral observations in which the tariffs for pair–products with zero trade flows are excluded might yield overstated tariff elasticities because of the endogenous relationship between tariff levels and trade: high tariffs correlate with low levels of trade, and in some extreme cases, tariffs might even be prohibitively high. This discussion shows that the direction of the bias will depend on various factors such as the exact sample and the structure of the error from the false interpolation.

I will focus the replication exercise on replicable papers (i.e., all data and code are publicly available), published in "top 5" journals and empirical setting, in which the results are expected to be highly sensitive to the data improvements done in the GTD. Hence, I will replicate three papers that describe different empirical approaches to estimating the trade elasticity in a cross-country setting using variation in MFN and preferential tariffs, i.e., Arkolakis, Ramondo, et al. (2018), C. E. Boehm et al. (2023), and Caliendo and Parro (2015). For all three of the studies I replicate here, the results change significantly. While in Arkolakis, Ramondo, et al. (2018) the attenuation bias dominates and the tariff elasticity increases from -4.28 to -9.64, I find that the estimates reported by Caliendo and Parro (2015) and C. E. Boehm et al. (2023) overstate the absolute tariff elasticity $-\sigma$.

The coefficient reported by Caliendo and Parro (2015) of -4.55 increases to -0.52 when I use the tariffs from the GTD. Almost three-quarters of this increase is due to the false interpolation, with the impact of the differential nature of the measurement error far outweighing the attenuation bias. The triple differentiation proposed by Caliendo and Parro (2015) is an elegant way of parsimoniously controlling for unobservables that might bias the estimates of $-\sigma$. However, the main assumption—namely, that once all symmetric trade costs and country-specific characteristics are accounted for through the triple differentiation, the remaining asymmetric determinants of trade costs are orthogonal to tariff rates—is violated when the researcher uses the data provided by WITS. The false interpolation leads us to wrongly assume that a country pair has MFN relations. The triple differentiation will take out the level effect of the MFN tariff on bilateral trade, but the tariff margin, i.e., the difference between the MFN and preferential tariffs, which equals the measurement error u, will be pushed into the error term of the regression, causing endogeneity.

The missing tariffs lead to overestimation of the absolute value of the tariff elasticity in C. E. Boehm et al. (2023): when I use the GTD, the coefficient increases to -0.61 instead of -2.23 for the long-run and to -0.33 instead of -0.65 for the short-run elasticity. While the time path of the elasticities over different horizons is similar in the sense that the short-run elasticities are much smaller than the long-run elasticities, the stabilization occurs earlier in my replicated results than in C. E. Boehm et al. (2023), namely, after five years instead of seven. Correcting for false interpolation only but keeping the sample selection, increases the

short-run tariff elasticity to -0.45 and the long-run elasticity to -1.48. This increase can be entirely attributed to the differential nature of the measurement error in the reduced form, which leads the authors to overstate the real effects of tariff changes.

Correcting for the false interpolation noticeably changes the coefficients reported by C. E. Boehm et al. (2023), but they become statistically different from each other only when I also correct for the positive selection. This indicates treatment heterogeneity across the different samples: trade between low-income countries reacts less strongly to changes in MFN tariffs than the remaining observations also contained in the WITS sample. A sample split using the GTD confirms this: While I find a long-run elasticity for high-income countries of -1.76, for low-income countries it is -0.08 and insignificant. This finding can be rationalized by heterogeneity across importers in terms of uncertainty over the duration of the change in MFN tariffs, which might hamper the reaction of trade flows. In addition, low-income exporters might be capacity constrained and therefore might have a more muted reaction to tariff changes.

My paper makes three important contributions: it is the first to systematically unpack the inconsistencies in tariff data that are widely used in the field. Second, it introduces a new global tariff database that can be used in many applications, also outside of the trade literature. Last, the paper shows that not addressing the missing tariffs leads to substantial bias in existing estimates one of the most important parameters in international economics: the trade elasticity, which has significant implications for the estimates of the gains from trade in most static trade models.

Related Literature The results of the replication study relate to the extensive literature seeking to estimate the trade elasticity with tariff variation, as surveyed, e.g., in Costinot and Rodríguez-Clare (2014), Head and Mayer (2014), and Hillberry and Hummels (2013). Despite the centrality of the trade elasticity, so far, there is no consensus on its magnitude. Estimates vary between -10 and -1 and depend heavily on the used method and the exact setting and sample (C. E. Boehm et al. 2023; Caliendo and Parro 2015; Fajgelbaum et al. 2020; Fink et al. 2005; Fontangé et al. 2022; Imbs and Mejean 2015; Romalis 2007). My results show that these discrepancies are caused not by the issues in the tariff data but rather by differences in method and sample. Furthermore, my paper shows that different empirical approaches and settings respond differently to the issues arising from the false interpolation and selection present in the WITS data. Last, I find compelling evidence for significant heterogeneity across importers reacting to MFN tariff changes. This results relates to Anson Soderbery (2018) as it highlights heterogeneity in the tariff elasticity across-countries and has important implications on how we think about the trade elasticity.

There have been previous efforts in putting together new global tariff datasets (Barari and Kim 2020; Bouët et al. 2008; Bown and Crowley 2016; Caliendo, Robert C. Feenstra, et al. 2023; Fontangé et al. 2022). However, the GTD is unique in terms of country, time coverage, and level of disaggregation. I also cross-validate all the tariffs with external sources, which is another new feature of the data. In addition, the level of detail of the information used to fill in the missing observations, especially for preferential tariffs, i.e., the detailed tariff schedules and explicit accounting for the phasing-out of tariffs, is novel. Methodologically, the paper draws from the literature on the measurement error to decompose the bias from the missing tariffs into three components (Bound, Brown, Duncan, et al. 1994; Bound, Brown, and Mathiowetz 2001; Bound and Krueger 1991; Imai and Yamamoto 2010; Schennach 2022).

The rest of this paper is structured as follows. Section 2 discusses how WITS mishandles and even exacerbates the issue of missing data. Second 3 introduces the GTD and carefully explains all steps taken to address the issue of missing tariffs. Then, I discuss the direction of the bias from the nonclassical measurement error and sample selection in Section 4, before replicating the three aforementioned papers on the tariff elasticity in Section 5. Section 6 concludes.

2 Missing Tariffs in Standard Sources

The World Bank's World Integrated Trade Solution (WITS) is researchers' standard source for downloading tariff data. In this section, I will show how WITS's data processing introduces error into the tariff data that makes them unsuitable for empirical research. Two issues affect the WITS tariff time series: false interpolation and positive selection.

2.1 False Interpolation

WITS offers worldwide information about tariffs stretching back to 1988, making it the key publicly available source for global panel tariff data.[3] The WITS download tool pools data from UNCTAD's Trains and the WTO, namely, the latter's IDB and CTS. Since 2010, most of the raw data used in Trains have come from the ITC.[4] The data include information for almost 200 countries at the Harmonised System (HS) 6-digit product level, with some of the data dating back to 1988.[5] The information about preferential and MFN tariffs is derived from both Trains/ITC and the IDB, while the CTS is the only provider of bound tariff data. In addition to tariff data, WITS offers information on trade flows and nontariff barriers, which can be used to simulate different trade policies.

To determine the tariffs that countries impose on each other, WITS uses the concept of the effectively applied tariff or min-tariff, defined as the lowest available tariff rate: that is, if a preferential tariff exists, it is used; otherwise, the effectively applied tariff equals the MFN tariff. More formally, the effectively applied tariff $t_{ijkt} = min\{t_{ikt}; t_{ijkt}\}$.[6] The same logic applies when, for example, users interested in sectoral tariffs download aggregated data: WITS first determines the effectively applied tariff t_{ijkt} and then aggregates the tariffs to the requested level.

The min-tariff rule hinges on a crucial assumption: when no preferential tariff is reported for a year, it is assumed that there is no trade agreement between the focal countries. However, countries do not necessarily report all of their tariff schedules, i.e., their MFN and all preferential tariff rates under every trade agreement, for all years. In cases where annual observations for preferential tariffs are missing, the min-tariff rule leads to false interpolation, with the imputing method yielding jumps in the tariffs to the MFN tariff level, which is higher

[3] WITS can be accessed here: https://wits.worldbank.org/.

[4] See the WITS homepage for more information on the data providers: https://wits.worldbank.org/dataproviders.html.

[5] For a few countries, tariffs are even available at the tariff-line level. Researchers can access these when downloading the data for individual countries instead of using the bulk download option.

[6] UNCTAD (2011, p. 95).

than the preferential tariff level. The resulting spikes in the tariff time series are not due to changes in trade policy but instead are an artifact of improper interpolation techniques.

Figure 1 shows the tariff relationship for selected country pairs that signed an RTA during the observation period, graphically illustrating how the missing observations in the reported preferential tariff rates and false imputation by WITS lead to substantial measurement error. The figure shows averages for more than 5,000 HS6-level products, where the red dots are the simple average of the preferential tariffs over all products k for a given year and the blue dots give the simple average of the MFN tariffs. The solid line is the effectively applied tariff, defined as the lower of either the observed preferential or MFN tariff for a given country pair, according to the data downloaded from WITS. Panel (a) displays Mexico's tariffs on trade with the United States. Until NAFTA came into effect on January 1, 1994, Mexico imposed the MFN tariff on US imports; for all years thereafter, the preferential tariffs negotiated in NAFTA were in place.[7] As Figure 1 shows, Mexico did not report its MFN and preferential tariffs for all years: from 1995 onward, the MFN tariff rates are available for every year, while for years prior, the information is available only for 1991. The country's reporting of its preferential tariffs was even more sporadic, as the sparse red dots illustrate. The min-tariff rule, i.e., the practice of defining the effectively applied tariff as the lower of either the MFN or the preferential tariff without accounting for the spotty reporting pattern, leads to massive variation over time in the effectively applied tariff, with spikes in each year for which the preferential tariffs are missing but the MFN tariffs are available—most certainly not the type of variation any researcher would wish to use to conduct empirical analyses.

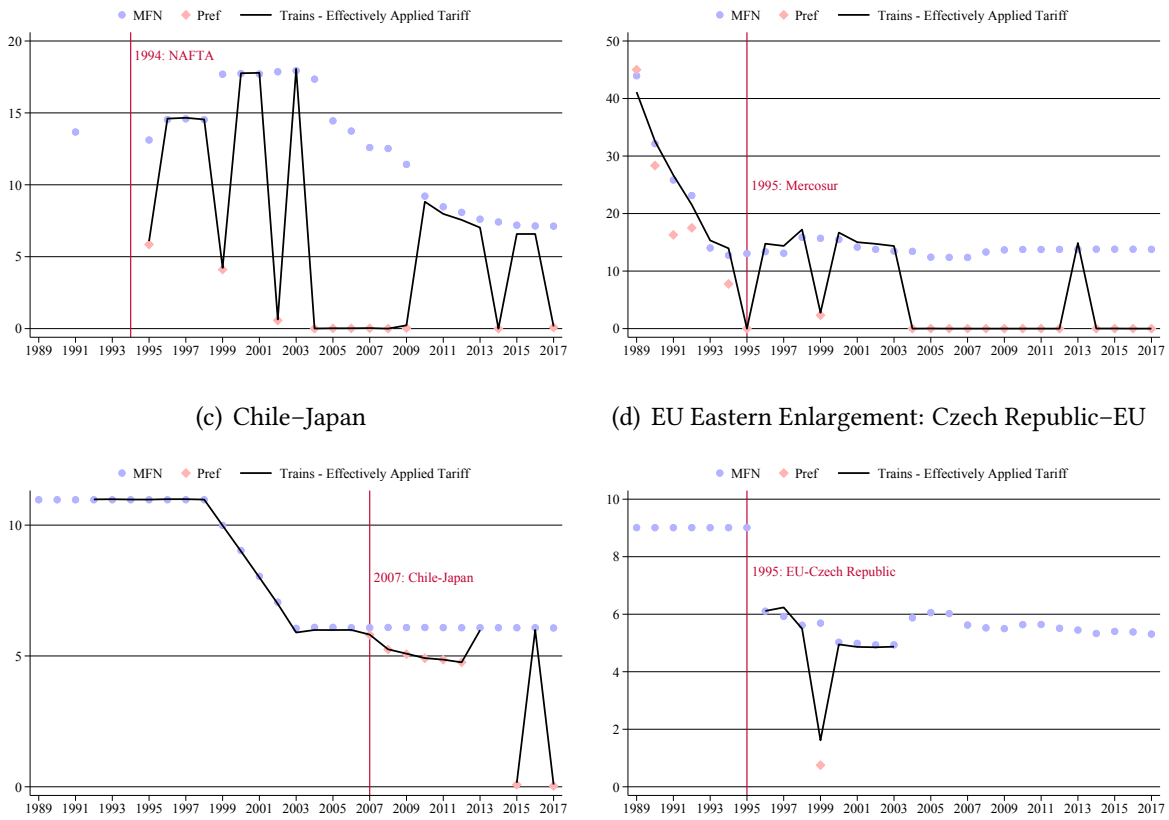
The problem of false imputation due to missing preferential tariff rates is not specific to Mexico—rather, it is prevalent for countries around the world involved in other important RTAs. Trade within Mercosur, Latin America's largest active RTA, is free for almost all products. After a four-year-long phase-in period, the Mercosur customs union entered into force in 1995. Again, the misreporting leads to measurement error in the WITS data (see Figure 1, Panel (b)). Similar issues affect the data on the tariffs imposed by Chile on Japanese imports, as illustrated in Panel (c). In Panel (d), we see an example of a time series affected by the eastern enlargement of the EU, where thirteen new member states joined the bloc, with the largest accession wave taking place in 2004, when ten countries joined. However, most of the tariff cuts had already been realized long before the accession because the new member states had implemented association agreements with the EU many years beforehand. For example, the Europe Agreement with the Czech Republic came into effect in 1995. Tariffs on EU ex-

[7] As Besedes et al. (2020) show, Mexico agreed to immediately cut all its tariff lines vis-à-vis Canada and the United States by 51%, with the remaining tariff lines phased out on different schedules, where the maximum phase-out period lasted 15 years.

Figure 1: Missing Tariffs and False Imputation in WITS: Four Examples

(a) NAFTA: Mexico–United States

(b) Mercosur: Brazil–Argentina



(c) Chile–Japan

(d) EU Eastern Enlargement: Czech Republic–EU

Note: The figure shows the simple average over all HS6-level products k of the preferential, MFN, and effectively applied tariff reported by WITS for selected country pairs. The spikes indicate missing preferential tariffs in the presence of reported MFN tariffs. The years of RTA implementation are from the Economic Integration Agreements database of Scott Baier and Jeffrey Bergstrand (Baier, Bergstrand, and Feng 2014).

ports to the Czech Republic were mostly eliminated under this agreement, a fact that is hard to recognize in the data from WITS, as the preferential tariffs are reported only for 1999.

Can the observed spikes be rationalized by any events related to trade policy? I will examine two types of events that potentially offer plausible explanations because they both lead to temporary tariff hikes: suspensions of trade agreements and antidumping tariffs. Especially in the case of unilateral trade arrangements, sometimes the granted preferences are temporarily revoked.[8] In such cases, the effectively applied tariff would correctly jump to the level of the applied MFN. However, suspensions are rare events and therefore cannot explain the

[8] Albornoz-Crespo et al. (2021) provide the example of the United States' suspension of the General System of Preferences for Argentina because of a dispute over intellectual property.

frequently observed spikes.[9] Furthermore, if suspensions were the sole explanation for the spikes, countries should always report their preferential tariff rates whenever an RTA is in force—this is not the case. In 2017, only 60% of all preferential tariffs are reported for country pairs that have an RTA (see Figure B1 (b)). Last, for the four examples discussed in Figure 1, I find no evidence that the respective agreements were suspended. Antidumping tariffs target only a handful of products. Hence, it is highly unlikely that the presence of a few antidumping tariffs would make the average effectively applied tariff of more than 5,000 products shoot up—let alone to the exact same level as that of the applied MFN tariffs.[10] In short, alternative explanations revolving around other trade policy tools fail to fully explain the observed pattern of spikes.

Last, one might think that countries report their preferential tariff rates conditional on their having positive trade flows. However, this is not the case; instead, an additional analysis in Appendix $A.1$ shows that false interpolation correlates negatively with zero trade, supporting the argument that the false interpolation affects missing preferential tariff observations: countries that trade more have a higher probability of participating in an RTA and therefore also a higher probability of their tariff series being affected by false interpolation, which can explain why a trade value of zero correlates negatively with the probability of false interpolation.

By understanding the logic of WITS's false interpolation, we can characterize the measurement error attributable to the treatment of missing preferential tariffs. First, it will be present only for country pairs with a trade agreement. Second, the falsely interpolated tariff rates will always overstate the true tariff rates. Unsurprisingly, reporting compliance correlates negatively with income (see Appendix A.1).

2.2 Positive Selection

When downloading the effectively applied tariff data from WITS, the user will notice that the number of observations is surprisingly small: in principle, for every year, WITS should provide the effectively applied tariff for approximately 174 million observations since there are 187 importing countries that impose tariffs on 186 exporters for approximately 5,000 products. However, the actual number of observations that can be downloaded is only a fraction of this.

[9] Upon cross-checking all 241 RTAs in force in 2019 using Mario Larch's RTA database (Egger and Mario Larch 2008), I find suspensions for only 4.72% of country pairs. A suspension is defined by a dummy variable for an effective RTA that switches more than once for a country pair between the year of entry into force and 2019. This definition overstates the presence of suspensions because it also includes RTAs in place for only a subset of years. Therefore, the share reported above is likely an upper bound on the true number of suspensions (i.e., the real share is even smaller).

[10] When I look at product-level tariff data, the discrepancies between the applied MFN and effectively applied tariffs downloaded from WITS vanish.

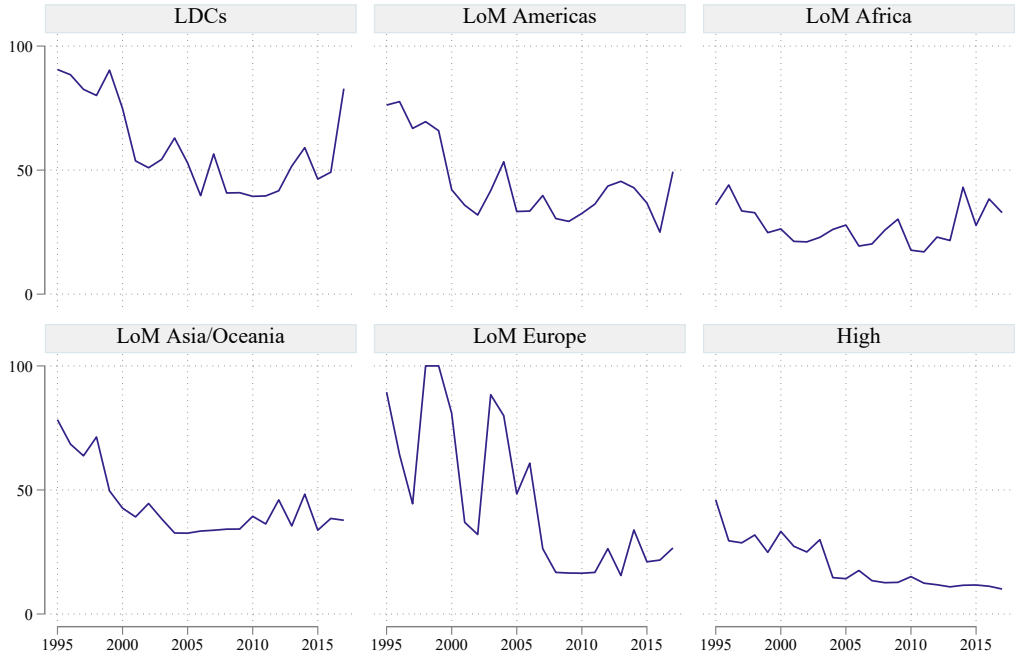
For 2000, WITS reports a mere 2.9 million observations. The reason is that WITS reports tariffs only for importer–exporter–product combinations with positive trade flows. Recall that WITS contains information not only on tariffs but also on trade flows as provided by UN Comtrade. Because the two data types are linked on the platform, WITS discards a great deal of available information on tariffs, leading to even more missing tariff observations.

For empirical work, this is troublesome because of the endogenous relationship between tariff levels and trade: high tariffs correlate with low levels of trade, and in some extreme cases, tariffs might even be prohibitively high. The tariff observations available on WITS will therefore be positively selected. This issue is related to the literature on the correct method for aggregating tariffs and other protectionist measures, which discusses how the common practice on using import weights underestimates the restrictiveness of a country's tariff regime (James E. Anderson and Neary 1994; Kee, Nicita, et al. 2008).

Even more troublesome is that WITS underreports trade flows because UN Comtrade is not complete. Thanks to the brilliant work by CEPII, we know that many countries do not report all of their trade to UN Comtrade or they report trade flows only at a more aggregated level (Gaulier and Zignago 2010). Because countries report both their imports and their exports, CEPII fills in the gaps for those countries that fail to report. Unsurprisingly, the countries with poor compliance are not random but rather are mostly low-income countries. Figure 2 shows the share of missing observations in WITS attributable to underreporting of trade flows in UN Comtrade. The share is determined by the number of observations available through WITS as a fraction of the larger number of observations available through CEPII's Baci trade data. Over all income groups, the share of missing observations attributable to underreporting in UN Comtrade is much higher for the earlier years of the sample, reflecting significant improvement in the coverage of the UN Comtrade data.[11] Furthermore, the share of missing observations correlates negatively with income. While it is relatively low for high-income countries, for least-developed countries (LDCs), it fluctuates for years since 2001 around 50%. In addition, low- and middle-incom (LoM) countries in the Americas are massively underrepresented in the WITS sample.

[11] The year 1995 is the first with available data in Baci.

Figure 2: Share of Missing Observations in WITS Attributable to Underreporting in UN Comtrade



Note: The graph shows the share of missing observations in WITS attributable to underreporting of trade flows in UN Comtrade. To share equals $1 - \frac{N_{WITS}}{N_{Baci}}$. The income groups are as defined by the World Bank, and I distinguish between least-developed countries (LDCs) and low- and middle-income (LoM) countries in the Americas, Africa, Asia/Oceania, and Europe as well as high-income countries.

3 Introducing the New Global Tariffs Database

In this section, I will present how the new global tariffs database (GTD) deals with missing data to alleviate the associated problems. The goal is to offer information on the correct effectively applied tariffs for as many country pairs as possible. The procedure can be broken down into three main steps: In *Step 1*, I combine the data available through WITS with those from other databases containing information on tariffs to minimize the number of missing observations; I end up using five different primary sources for tariffs. In *Step 2*, I develop a new algorithm to impute the data that are still missing. To increase precision for the imputation of preferential tariffs, I use detailed information about phase-out regimes for more than 500 RTAs provided by the DESTA project (Dür et al. 2014). Steps 1 and 2 are conducted separately for MFN and preferential tariffs, as the goal is to fill in the missing observations for both types of tariffs. In *Step 3*, I combine the resulting MFN and preferential tariff data—for which there are no longer any missing observations—and use them to retrieve the effectively applied tariffs t_{ijt} for all possible importer–exporter–product–time combinations to address the positive selection that affects the WITS data.

3.1 *Step 1:* Combining Different Sources to Minimize the Number of Missing Observations

I obtain information from five primary sources: *1)* detailed phase-out schedules for 149 RTAs available through the WTO's RTA database; raw tariff schedules that countries report from *2)* Trains and *3)* the IDB (both accessed via WITS) and *4)* the ITC; and *5)* US and EU tariff schedules from national authorities kindly provided by researchers from the World Bank (Forero-Rojas et al. 2018). Table 1 summarizes the coverage of the respective sources and shows the available years, number of importing countries, and type as well as form of the tariff data that they contain. The number of the rows gives the order of replacement of the missing data: The first row is the base. If information is still missing, I use the information contained in the source described by the second row, and so on. I will next describe in detail the different sources and how they are used.

According to Baier and Bergstrand (2007), "*virtually every FTA is phased-in, typically over 10 years*" (p. 89–90). Countries decide on a product-by-product basis whether the respective tariff is to be phased out. For NAFTA, 51% of all tariff lines were cut immediately; for the remaining tariff lines, the MFN tariffs were phased out over different time horizons (Besedes et al. 2020). This practice makes imputing missing preferential tariffs particularly hard: simply assuming that preferential tariffs equal zero once an RTA is signed will understate the actual

Table 1: Primary Sources for Tariffs

Source	type of tariffs	ad valorem equivalents	covered years	# importing countries
1) Phase-out schedules	preferential	No	1994–2017	91 (149 FTAs)
2) Trains	MFN & preferential	Yes	1988–2017	197
3) IDB	MFN	No	1996–2017	163
4) ITC	MFN & preferential	Yes	2007–2009	197
5) National authorities	MFN & preferential	No	1997–2017	EU & US

Note: The table shows the used primary sources for tariffs and their coverage with respect to available years, number of importing countries, and type as well as form of the tariff data that they contain.

tariffs; i.e., the imputed tariffs will be too low. Hence, dealing with missing observations for preferential tariffs is a major challenge, and one goal of the GTD is to address this issue.

For 149 FTAs, I have gathered information on the exact phase-out schedules for all tariff lines and participating countries; when this information is available, no imputation is necessary in the first place. The schedules can be found through the WTO's RTA database, with especially good coverage of more recent FTAs.[12] The phase-out schedules are the digitized version of the long appendices attached to every FTA, and they define the exact timing of the tariff cuts for the respective tariff lines and parties for each of the covered FTAs. Overall, 384 schedules are available. Unfortunately, the format of the raw files is not identical; therefore, as a first step, the layout of the hundreds of files had to be manually standardized. Then, I aggregated the information to the HS6 product level by taking the average of the corresponding tariff lines.[13] The main advantage of the phase-out schedules is their completeness: they cover all tariffs for all national tariff lines and all years. Hence, whenever this type of information is available, no imputation is necessary. Therefore, for preferential tariffs, the schedules serve as my primary source of information—only when they are not available do I use data from the other sources.

While the WTO's phase-out schedules cover only a subsample of all FTAs currently in force, Trains gives all the preferential tariffs offered through any type of RTA (including FTAs, customs unions, partial scope agreements, and nonreciprocal arrangements), as well. In fact, Trains has the broadest coverage in terms of countries (197 importers and all their partners) and years (starting in 1988). It provides information on both MFN and preferential tariffs. Un-

[12] In 2006, the General Council established a new transparency mechanism for all RTAs that is supposed to help ensure that RTAs fulfill the requirements of Articles XXIV and V. As a result, so-called factual presentations have to be distributed among WTO members. One part of the factual presentation is the tariff schedule with all phasing-in schemes (see https://www.wto.org/english/tratop_e/region_e/trans_mecha_e.htm for details). The data are available at https://rtais.wto.org/.

[13] Only ad valorem tariffs are included.

like the IDB, Trains shows preferential and MFN tariffs separately, which is necessary when the goal is to correct for the troublesome spikes. Instead of specifying the preferential tariffs, the IDB has only the effectively applied tariffs that include the spikes, making it impossible to distinguish between imputed MFN and reported preferential tariffs. Therefore, I use only the MFN tariff information from the IDB and discard the effectively applied tariff altogether. Trains is therefore the core primary source of tariff data in the GTD for most country pairs. Tariffs collected by Trains and the IDB are available through WITS.[14]

The ITC's Market Access Map (MAcMap) is another established source for tariff data.[15] It incorporates bound, applied MFN and preferential tariffs from 2007 onward for 197 countries on the tariff line. The ITC's MAcMap provides raw data that is available to the public. Thus, unless countries report perfectly, problems similar to those in WITS can be expected: i.e., missing observations, especially for preferential tariffs.[16] Since 2010, Trains has supplied WITS with data on tariffs collected by the ITC. Indeed, for the more recent years, the raw data reported through WITS are identical to those that can be downloaded from the ITC website. For the remaining years, i.e., 2007 to 2009, I add the additional data available through the ITC to the raw data that the phase-out schedules and WITS provide. Again, the original data are available for national tariff lines. To match them with WITS and to make them comparable across countries, I aggregate to HS6 product level.

Last, I use data provided by US and EU national authorities. The US International Trade Commission gathers and publishes tariffs for the United States and the European Commission those for the EU.[17] All data cleaning and preparing has been done by researchers from the World Bank (Forero-Rojas et al. 2018), who have kindly provided me with access.[18] Similarly to the phase-out schedules, the original data give tariffs at the lowest level of disaggregation, i.e., the national tariff line, and have to be aggregated to HS6 product level to be comparable across countries and to be matched to the tariff data from WITS.

For preferential tariffs, the phase-out schedules from the WTO's RTA database serve as the primary source of information. Missing observations of preferential tariffs are sequentially filled in with information provided by Trains, the ITC and, last, for the US and EU, with data from the national sources. For MFN tariffs, Trains serves as the primary source; the

[14] The data are available at https://wits.worldbank.org/.

[15] It is available at https://www.macmap.org/.

[16] In addition, the ITC offers a built-in imputation method that is supposed to yield the effectively applied tariff. When the user chooses "minimum rate" instead of "by each trade agreement", the data will follow the min-rule resulting in false interpolation.

[17] The data are available at https://dataweb.usitc.gov/tariff/annual and https://circabc.europa.eu/faces/jsp/extension/wai/navigation/container.jsp.

[18] Only ad valorem tariffs are included.

missing data are sequentially filled in with information provided by the IDB, ITC, and national authorities.

Another complication arises due to non-ad valorem tariffs. Regardless of the type of tariff—bound, MFN or preferential—a tariff can take one of two forms, i.e., ad valorem (for example, 8%) and non-ad valorem (for example, 1.22 USD/kg or 1.22 USD/kg + 8%). In theory, it is possible to convert non-ad valorem tariffs into ad valorem equivalents (AVEs) by dividing the non-ad valorem element of the tariff by the value of the product per unit. In practice, calculating ad valorem equivalents (AVEs) is very complicated; see Bouët et al. (2008) for a more detailed discussion. Trains and the ITC offer AVEs for specific tariffs; however, they use different methods, yielding different AVEs.[19] For the IDB data, the phase-out schedules, and the data provided by national authorities, AVEs are simply not available.[20]

For most empirical analyses, it might be sensible to exclude non-ad valorem tariffs altogether since they vary even if no change in tariffs occurs. To facilitate this, the GTD flags all AVEs based on information provided by Trains and the ITC. For 2017, the available data suggests that less than 2% of country pair–products are subject to non-ad valorem tariffs. However, this most likely understates the true share because Trains/ITC record only the non-ad valorem tariffs for which Trains/ITC calculate AVEs. Whenever countries calculate AVEs themselves and end up reporting only ad valorem tariffs, it is impossible to trace back for which product the reported tariff is a true ad valorem tariff or is instead a non-ad valorem tariff. The EU is a prominent example of a jurisdiction that engages in this practice. Unfortunately, there is no way of systematically identifying all products with non-ad valorem tariffs. For practitioners, one solution might be to exclude extreme values of tariffs, e.g., those that suggest rates high as 150% or higher, as these have a high probability of being AVEs.

3.2 *Step 2:* Imputing the Remaining Missing Observations

Having combined all available raw data on tariffs, I carefully correct for the remaining missing data. I start by imputing MFN tariffs as I need them when dealing with preferential tariffs. Although the missing data are much more pronounced for preferential tariffs, reporting compliance is also imperfect for MFN tariffs t_{ikt} of importer i for product k in year t. Rather than replacing missing MFN tariffs by linearly interpolating between available observations, I set them equal to the nearest preceding observation. This procedure accounts for anecdotal ev-

[19] From 2010 onward, the ITC has delivered tariffs to Trains. The only differences between the Trains and ITC data for the years 2010 to 2017 are observable for the ad valorem equivalents of non-ad valorem tariffs, as Trains and the ITC do not use the same method for computing the AVEs.

[20] The IDB ignores specific tariffs altogether, resulting, yet again, in missing data; a specific tariff of 1.22 USD is recorded as missing and a mixed tariff of 1.22 USD + 8% is recorded as equal to 8%.

idence that countries are more likely to update schedules after a significant tariff change. If there is no preceding observation, the missing MFN tariffs are set equal to the nearest subsequent observation. Caliendo, Robert C. Feenstra, et al. (2023) propose the same method.

To use all available information over time, it is crucial to convert all six-digit products into the first available nomenclature, HS88/92. Otherwise, I can fill only missing tariffs within the same nomenclature or for those product codes that do not change over the 30 years. In addition, the raw data do not allow me to readily compare tariffs within the same product across importers because many countries do not use the correct nomenclature: in theory, every country should use HS88/1992 for the years 1988 to 1995, HS1996 for 1996 to 2001, and so on. If this were the case, tariffs for product k could be compared across importers i at any time t. However, many countries update to the new nomenclature only with a significant lag—independently of which original source data is used. For example, in WITS, Guyana reports its MFN tariffs for 2017 using the outdated HS2002. Using one consistent nomenclature deals with this issue.[21]

Once all the MFN tariffs are imputed, I turn to the preferential tariffs. The imputation method accounts for the presence of RTAs: preferential tariffs must not be filled in for years when no RTA is in place. This sounds trivial but is crucial, as otherwise the imputation will do more harm than good. To account for any sort of RTA being in effect, I combine Mario Larch's RTA database (Egger and Mario Larch 2008), the database on economic integration agreements put together by Scott Baier and Jeffrey Bergstrand (Baier, Bergstrand, and Feng 2014), and information provided by DESTA (Dür et al. 2014). Only when for a country pair ij does at least one of the three mentioned RTA databases confirm that there is in fact an RTA in force do I impute the preferential tariff.

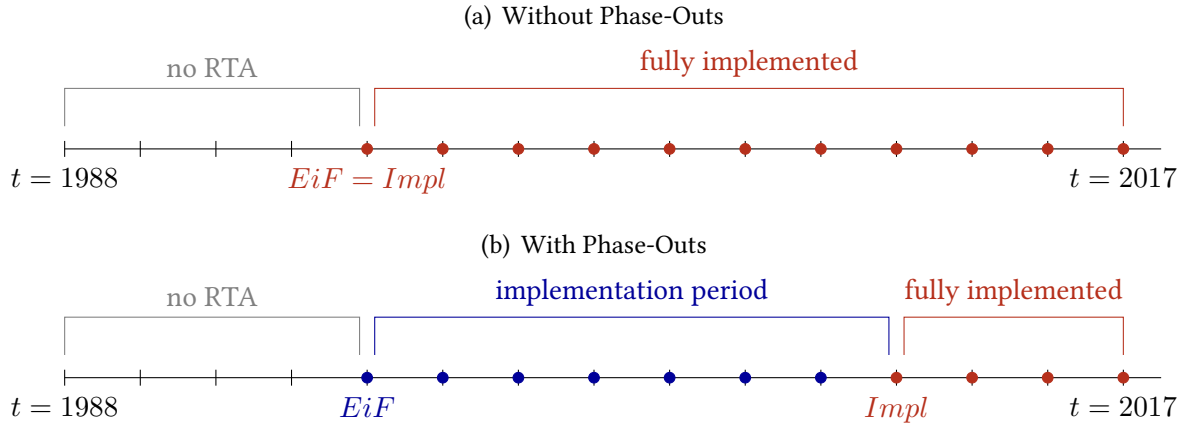
The imputation method should account for tariff cuts being phased in, instead of simply assuming that all tariffs equal zero once the RTA is in place. For 149 FTAs, I do not have to make any assumptions because, fortunately, the WTO's RTA database provides full information on the phase-out schedules. For the remaining ones, I distinguish between RTAs with and without phase-outs, where for the latter all tariffs are cut immediately. The DESTA database (Dür et al. 2014) has gathered phase-out regime information for more than 600 RTAs and tells us whether there is at least one product for which tariffs are being phased out; DESTA specifies the year when the last tariff cuts must be implemented. Hence, I know for all RTAs in the sample the year of entry into force EiF_{ij}, the year when all tariff cuts must be fully implemented $Impl_{ij}$, and the effectively applied tariff in $EiF - 1$, as it equals the MFN tariff one year prior

[21] Concordance tables are available through WITS.

to the entry into force ($t_{ijk,EiF-1} = t_{ik,EiF-1}$). With these three pieces of information, I can now start filling in the missing preferential tariffs.

The imputation algorithm splits the available data into segments and distinguishes *(i)* the time before the RTA enters into force (marked in gray in Figure 3), *(ii)* the implementation period (in blue), and *(iii)* the years when the RTA is fully in force (in red). The time is displayed on the x-axis of the schematic figure. When no phase-out is stipulated in the RTA, the year of entry into force is identical to the year when full implementation is reached ($EiF_{ij} = Impl_{ij}$; see Panel (a)).

Figure 3: Imputation Algorithm for Preferential Tariffs



I start by filling in the observations for the red segment, i.e., those for years when the RTA is fully implemented. For this period, any tariff that can be observed equals the final one, typically zero. Using the year of full implementation as the start date and 2017 as the end date, I sequentially fill in the missing observations for the red segment with the first subsequent nonmissing observation. If no subsequent observation is available, the missing preferential tariffs are set to the nearest preceding observation. For RTAs that do not stipulate any phase-outs, this is all that has to be done. Note that, in these cases, the algorithm delivers correct imputation even if the importing country reports only once in this period.

For RTAs with phase-outs, I now turn to the implementation period, i.e., the blue segment. I model the missing phase-out tariffs by linearly interpolating between the available observations in the implementation period. If the preferential tariff for the year of the entry into force is missing, I use the MFN tariff that was in effect in the year prior to the entry into force $t_{ik,EiF-1}$. The preferential tariff at $t = Impl_{ij}$ is the end point for the linear interpolation. I decided to linearly interpolate because most phase-out schedules follow some sort of linear reduction: for example, tariffs might be gradually reduced in five or 10 equal cuts to zero.[22]

[22] See Besedes et al. (2020) for details on NAFTA.

For the last segment (gray in Figure 3), i.e., the time before the RTA is in place, no preferential tariffs are in place, and the effectively applied tariff equals the MFN tariff.

In some cases, multiple preferential tariff schemes are available; for example, least-developed countries qualify for the Everything but Arms arrangement and the General System of Preferences of the EU. Whenever a country is eligible for multiple preferential tariffs, I assume that the lowest tariff applies.

3.3 *Step 3:* Retrieving the Effectively Applied Tariff

At this point, neither the MFN nor the preferential tariffs have missing observations anymore. Next, I set up a matrix that consists of every country pair–product–year combination $ijkt$, hence, taking care of the selection present in the WITS data. For every year t, I combine 197 importers i with 196 exporters j for on average 4,900 products k, resulting in a total of 5.7 billion observations. Most countries report tariffs for all 5,018 products of the HS88/92 nomenclature, but some low-income countries and countries that have not yet joined the WTO do not report tariffs for the whole product range. Figure 4 shows how the imputation method changes the time series for the examples given in Section 2. For all country pairs, the troublesome spikes have vanished.

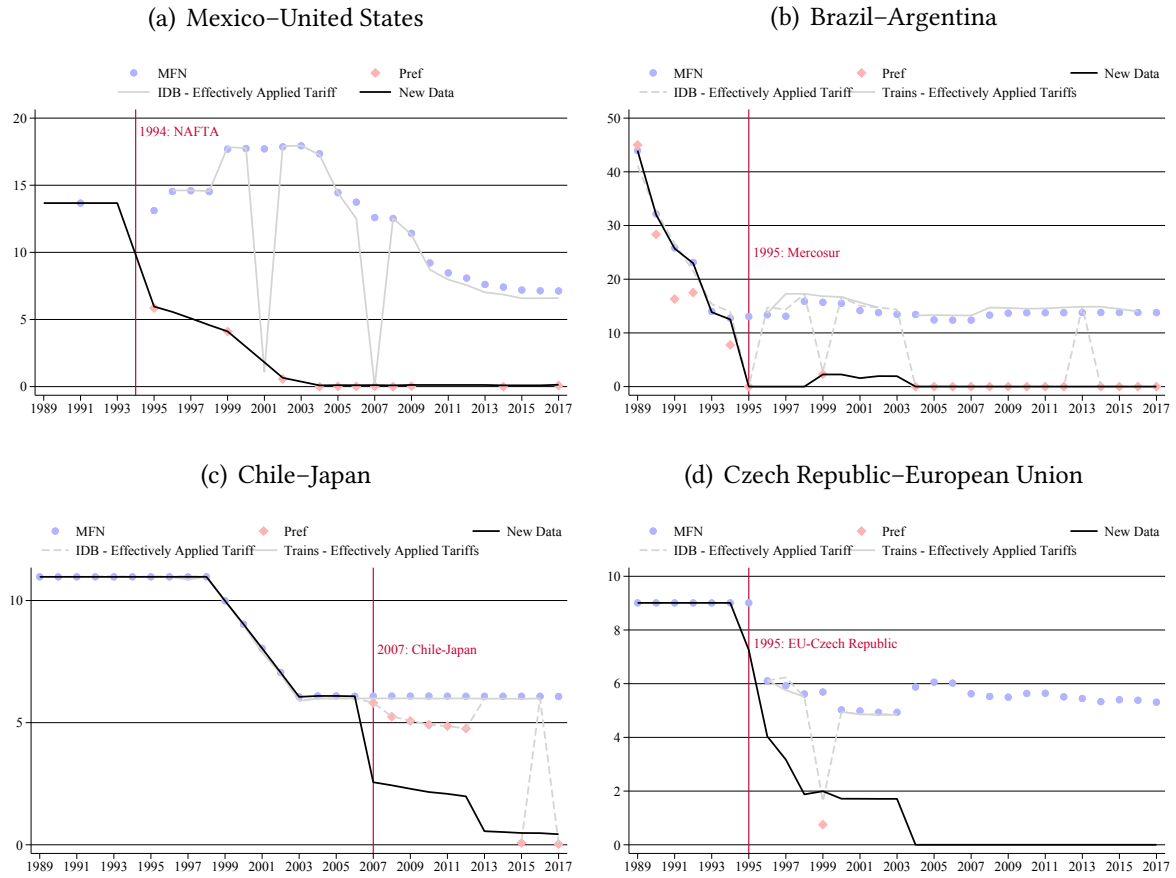
3.4 Comparison with Other Sources

While there are a many databases available that inform about tariffs for specific countries, years or products, such as, for example, the Agricultural Market Access Database Bouët et al. (see 2008) for a summary of alternative databases), very few databases provide information comparable to the data presented in this paper with respect to country and time coverage and level of disaggregation (HS6).

CEPII's *MAcMap-HS6* is a collaborative project between CEPII and the ITC. Its data are significantly more reliable than those available through WITS. First, they do not suffer from selection bias; i.e., the database gives tariffs for all importer–exporter–product combinations. Second, the ITC addresses the missing preferential tariffs. This processed version of the ITC data is not, however, publicly available as of January 2024.[23] Unfortunately, to the best of my knowledge, there is no transparent documentation available that describes in detail what exact steps have been performed by the ITC to fill in the missing preferential tariffs. The description that follows is based on information provided most kindly by Gianluca Orefice,

[23] The raw and unprocessed ITC data can be accessed here: https://www.macmap.org/.

Figure 4: Imputed Effectively Applied Tariff for Selected Country Pairs



Note: The graph shows the simple average of the imputed effectively applied tariff for selected country pairs.

who used the data in his own research, and on my own understanding from checking the available data.

The steps performed by the ITC to fill in the missing preferential tariffs are similar to those I follow to produce the GTD: it first tries to minimize the number of missing observations by going back to the text of the signed trade agreements to extract information about the phasing-out of tariffs. When tariffs are not reported in t, the ITC fills them in with the tariffs reported in $t-1$ and $t-2$; if the reporting country does not report in those three years, the ITC chooses the subsequent year $t+1$ instead.

However, my understanding is that the infilling algorithm relies on countries' reporting tariffs in one of the years around the missing observation. Therefore, the data will fail to deliver correctly imputed tariffs in cases where back-to-back observations for many years are missing. For example, the trade agreement between Canada and Chile entered into force in 1997, but Chile did not report tariffs until 2006. In contrast, the GTD uses the data that

Chile reported in 2006 for all years from the entry into force of the Canada–Chile agreement. The ITC data fail to retrieve all preferential tariffs for this bilateral relationship, while the GTD also reports preferential tariffs for the years 1999–2006. This is possible because I link every preferential tariff with a corresponding year of entry into force, which, according to my understanding of the available documentation, the ITC does not do. More generally, my algorithm uses all available information over all years, while the ITC's infilling procedure uses less information. Furthermore, since the ITC performs the filling-in for each country in the reported nomenclature, for products for which product codes break because of changes in the nomenclature, interpolation is not possible.

I also encountered some irregularities that I was not able to fully understand with the available documentation. For example, CEPII's MAcMap reports for all ten new member states that joined the EU in 2004 zero tariffs on all products in 2001. This is wrong, as there were exemptions and phase-outs.[24] For example, Poland did not fully eliminate tariffs until 2003. The GTD accounts for this because in 2001 Poland actually reported information about its preferential tariffs on imports from the EU. Thus, no imputation was necessary. Since the data are in principle available, I do not fully understand why CEPII's MAcMap reports free trade already in 2001.

CEPII then processes the tariff data further and improves them, especially with respect to non-ad valorem tariffs. CEPII first converts all HS6 products to one nomenclature to make intertemporal and cross-country comparisons possible. Second, there is a special focus on the calculation of AVEs. Without doubt, CEPII's *MAcMap-HS6* is the best source for non-ad valorem tariffs and, in particular, the AVEs of tariff rate quotas. Bouët et al. (2008) describe the exact methods used to convert all five forms of non-ad valorem tariffs to AVEs. The data are publicly available through CEPII's website.[25] Fontangé et al. (2022) are a prominent example of authors using these tariff data in the context of estimating trade elasticities.

Another advantage of the GTD over CEPII's *MAcMap-HS6* is the coverage of the former in terms of available years: *MAcMap-HS6* covers only the six years 2001, 2004, 2007, 2010, 2013, and 2016.[26]

Caliendo, Robert C. Feenstra, et al. (2023) have constructed a database similar to the one presented here. However, their dataset differs with respect to the covered years (1984–2011) and degree of precision of the preferential tariffs. They add data from three other sources to the tariffs provided by WITS: manually collected tariff schedules published by the International

[24] Official Journal of the European Communities No L114/27.

[25] The data can be downloaded at http://www.cepii.fr/CEPII/en/bdd_modele/presentation.asp?id=12.

[26] The available years correspond to those with releases of the Global Trade Analysis Project (GTAP) database, for which CEPII provides information on tariffs.

Customs Tariffs Bureau, US tariff schedules from the US International Trade Commission, and US tariff schedules derived from detailed US tariff revenue and trade data provided by the Center for International Data at UC Davis. The imputation algorithm used in the two databases is very similar, most likely resulting in very similar MFN tariffs. To account for the phasing-in of preferential tariffs, Caliendo, Robert C. Feenstra, et al. (2023) include information on approximately 100 FTAs and their phasing-in regimes, i.e., whether most tariff lines are cut immediately or phasing-in is common. This information is coarser than that in the GTD, where I implement a considerable improvement by including detailed phasing-in schedules at the tariff-line level for 149 FTAs. For the agreements for which this information is not available, similarly to Caliendo, Robert C. Feenstra, et al. (2023), I use the information on phasing-in regimes to construct the filled preferential tariffs. In contrast to CEPII's MAcMap, the data developed by Caliendo, Robert C. Feenstra, et al. (2023) perform the filling conditional on the entry into force of agreements, i.e., preferential tariffs are filled for the entire time period. In terms of aggregation, the data are available only at the Standard International Trade Classification (SITC) level and hence are more aggregated than the data in the GTD. Last, as of January 2024, the publicly available data provide tariffs only for country pair–industry combinations with positive trade flows.[27]

To the best of my knowledge, the database presented here is the first to simultaneously deal with the missing tariffs, account explicitly for the phasing-in schedules of RTAs, and cross-validate the information to minimize error in the original data. The main contribution of the new tariff database is twofold: First, the coverage in terms of number of countries, level of disaggregation, and years available is unique. Second, the level of precision of the preferential tariffs is higher than that in other databases of this kind.

[27] The data can be downloaded here: https://rcfeenstra.github.io/CFRT/.

4 Understanding the Bias

This section contains a general discussion of the bias in estimates of the effect of tariffs on trade due to the missing tariff observations. This bias consists of three components: i) the bias due to the nonclassical nature of the measurement error in the falsely interpolated tariffs, ii) the bias due to the differential nature of the measurement error, and iii) the positive selection in the reported trade flows. The net effect and in particular its direction depend on the interplay of these three competing forces. I will discuss the two issues of false interpolation and positive selection separately.

4.1 Bias Due to Measurement Error

To understand the bias attributable to the false interpolation, I closely follow the approach for decomposing bias due to measurement error put forward in the applied labor literature (Bound, Brown, Duncan, et al. 1994; Bound, Brown, and Mathiowetz 2001; Bound and Krueger 1991). Suppose that the true relation between the outcome variable y_{ijk} and ad valorem tariffs $\tau_{ij} = ln(1 + t_{ij})$ that importing country i imposes on country j is given by

$$y_{ij} = -\sigma_s \, \tau_{ij} + \epsilon_{ij} \tag{1}$$

where ϵ_{ij} is uncorrelated with τ_{ij}. The index s in the coefficient $-\sigma_s$ indicates that the resulting tariff elasticity is estimated on the same sample available in tariff data downloaded from WITS. Hence, it is the resulting elasticity when the messy tariff data are purged of the falsely interpolated data but when any issues arising from sample selection are ignored. For ease of notation, I will omit the product dimension k and the time index t, but all the results can be extended easily. As some of the preferential tariffs t^*_{ij} are falsely interpolated with MFN tariffs t_i, I do not observe τ_{ij} but instead

$$\tilde{\tau}_{ij} = \tau_{ij} + u_{ij}, \tag{2}$$

with $u_{ij} = ln(1 + t_i)$ for share ω_{ij} of country pairs with an RTA.

The least squares estimator of $-\tilde{\sigma}_s$ is

$$-\tilde{\sigma}_s = \frac{cov(y_{ij}, \tilde{\tau}_{ij})}{var(\tilde{\tau}_{ij})}$$

and can be rewritten as

$$
\begin{aligned}
-\tilde{\sigma}_s &= -\sigma_s \left(\frac{var(\tau_{ij}) + cov(\tau_{ij}, u_{ij})}{var(\tau_{ij}) + var(u_{ij}) + 2cov(\tau_{ij}, u_{ij})} \right) + \frac{cov(\epsilon_{ij}, u_{ij})}{var(\tau_{ij}) + var(u_{ij}) + 2cov(\tau_{ij}, u_{ij})} \\
&= -\sigma_s \underbrace{\left(\frac{cov(\tau_{ij}, \tilde{\tau}_{ij})}{var(\tilde{\tau}_{ij})} \right)}_{b_{\tau,\tilde{\tau}}} + \frac{cov(\epsilon_{ij}, u_{ij})}{var(\tilde{\tau}_{ij})} = -\sigma_s b_{\tau,\tilde{\tau}} + \Omega
\end{aligned}
\tag{3}
$$

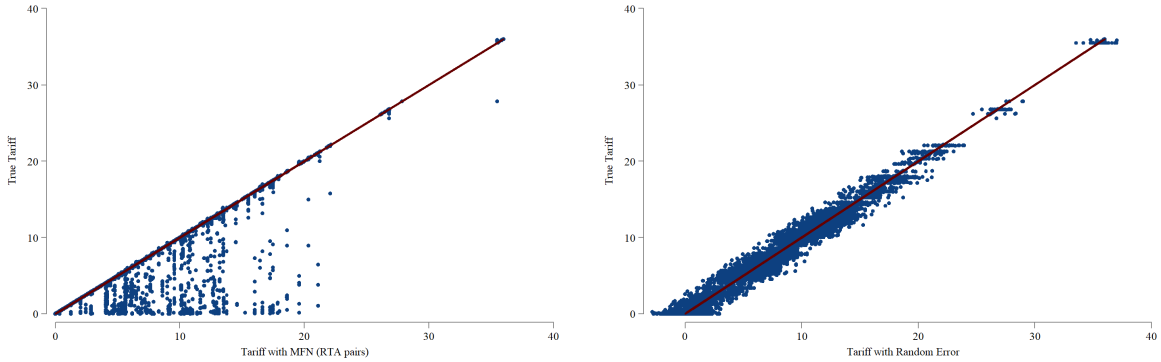
If we could conveniently assume that the measurement error u_{ij} were uncorrelated with τ_{ij} and ϵ_{ij}, the biased $\tilde{\sigma}_s$ would be unambiguously smaller than σ_s. In this case, the bias would be entirely determined by the first term $-\sigma_s(.)$ because the covariance terms $cov(\tau_{ij}, u_{ij})$ and $cov(\epsilon_{ij}, u_{ij})$ would be equal to zero by assumption. It is easy to see that the term in parentheses will always be smaller than one and, in fact, bounded by zero and one—the well-known result from classical measurement error. However, in the context of missing tariff data and the resulting false imputation, these two assumptions do not hold.

Recall that the WITS imputation method replaces missing preferential tariff observations with the higher MFN tariff rates; therefore, the corrupted tariffs will always be larger than the true tariffs. Figure 5 Panel (a) shows this relationship for the year 2001 by plotting the true tariffs τ_{ij} against the tariffs affected by measurement error $\tilde{\tau}_{ij}$. If the error u_{ij} were indeed random, the corrupted data would align around the 45-degree line, introducing solely noise to the explanatory variable but no systematic pattern (see Panel (b)). Note that the false interpolation is much more frequent for low tariffs because preferential tariffs are in most cases zero or very low. Therefore, the error u_{ij} correlates negatively with the true tariff τ_{ij}; i.e., positive values of u_{ij} are associated with smaller true tariffs, making the measurement error nonclassical. With $cov(\tau_{ij}, u_{ij}) < 0$, it is a priori unclear in which direction the false interpolation biases estimates—even a flip in sign is possible if the error and the true tariffs are sufficiently negatively correlated.

The term in parentheses can be compactly rewritten as the regression coefficient $b_{\tau,\tilde{\tau}}$ obtained from regressing τ on $\tilde{\tau}$. This is a general result and can be extended to other empirical approaches that use ordinary least squares (OLS) to estimate the tariff elasticity, such as those accounting for multilateral resistance terms, other control variables or time-differencing, which are common approaches to addressing endogeneity concerns in the literature. While the size of $b_{\tau,\tilde{\tau}}$ depends on the identification strategy and how it changes the error structure, it can always be recovered by regressing the true on the falsely imputed tariffs. As $\tilde{\tau}_{ij}$ will always be smaller than τ_{ij} and can never become negative, $b_{\tau,\tilde{\tau}}$ will be always smaller than one. Hence, the first element of the bias attributable to false interpolation will always lead to attenuated results and push the coefficient toward zero.

Figure 5: Measurement Error Due to False Interpolation

(a) Nonclassical Measurement Error: $\tau \leq \tilde{\tau}$ (b) Random Measurement Error



Note: The figure shows the measurement error due to false interpolation. Panel (a) plots the true tariffs τ against the corrupted ones $\tilde{\tau}$ for the year 2001; Panel (b) shows how the two variables would be associated with each other if the measurement error were random with a mean of zero and variance of 1.

The second part of the bias, denoted as Ω in Equation (3), comes from the fact that the false interpolation occurs only for country pairs with an RTA—observations of tariffs with error are not random neither across country pairs nor over time. Of course, pairs with a trade agreement systematically differ from those without: for example, pairs with a trade agreement might be geographically closer to each other. Therefore, geography is a determinant not only of trade but also of tariffs—both the true tariffs τ_{ij} and $\tilde{\tau}_{ij}$. The error term ϵ_{ij} in Equation (1) contains geography and any other residual variables that might simultaneously determine tariffs and trade, violating the assumption $cov(\epsilon_{ij}, u_{ij}) = 0$. The measurement literature refers to this phenomenon as *differential measurement error* (e.g., Bound, Brown, and Mathiowetz 2001; Imai and Yamamoto 2010). As many of the variables contained in the residual ϵ_{ij}, such as bilateral distance, income levels and consumer preferences, correlate with trade and tariffs in opposite directions, using the mismeasured tariffs as an explanatory variable will most likely overstate the true effect of tariffs on trade.

If $cov(\epsilon_{ij}, u_{ij}) = 0$ were true, the error would not systematically differ across country pairs, and hence only nonclassical measurement error would be present, always attenuating the estimates. However, because the error is not just nonclassical but also differential, we also have the competing element stemming from the differential nature of the error that pushes the coefficient of interest down. Note that the direction of this bias might be different if we were interested in a different outcome since the sign of $cov(\epsilon_{ij}, u_{ij})$ might switch. The total

bias attributable to false interpolation $\Delta\ \sigma^M$ can be summarized as

$$\Delta\ \sigma^M = plim\ -\tilde{\sigma}_s - (-\sigma_s) = \underbrace{-\sigma_s(b_{\tau,\tilde{\tau}} - 1)}_{\textit{nonclassical ME}} + \underbrace{\Omega}_{\textit{differential ME}},$$

and will be the absolute difference in coefficients when we estimate the tariff elasticity with the WITS data and the new data. It is straightforward to decompose the total bias into two distinct elements. The first is the part that is attributable to the messy data but that does not take into account that country pairs with corrupted tariffs might be very different from those without, which I will refer to as *nonclassical ME*. Its magnitude depends on the signal-to-noise ratio measured by $b_{\tau,\tilde{\tau}}$ and the unbiased $-\sigma_s$ and and is therefore easy to determine. The second part of the bias comes from the differential nature of the measurement error, i.e., $cov(\epsilon_{ij}, u_{ij}) \neq 0$, and we can quantify its magnitude by subtracting the observable nonclassical measurement error from the total bias.

4.2 Bias Due to Selection

Even if false interpolation were no longer an issue—for example, because clean data are available and u_{ij} is always zero—estimating $y_{ij} = -\sigma\tau_{ij} + \epsilon_{ij}$ with the available data through WITS would lead to biased estimates because of positive selection. WITS provides tariffs only for importing countries that report positive trade flows to UN Comtrade. As is the case for tariffs, trade reporting patterns in the UN Comtrade data are not random; rather, low-income countries have a significantly lower probability of reporting: for LDCs, the probability of reporting trade is 15 percentage points lower than that of the remaining countries for 1995 to 2017 and is a staggering 34 percentage points lower for the first five years (1995 to 2000) of the sample period.

This systematic pattern of tariff missingness creates selection bias. To see this, let $s_i = 1$ if all of importing country i's bilateral links are included in the WITS sample and equal zero otherwise. s_i can be rewritten as a function of income gdp_i and other unobservables e_i; hence, the selection equation is

$$s_i = \gamma\ gdp_i + e_i.$$

Estimating Equation (1) for the available sample will give a biased estimate for the elasticity $-\sigma_s$. In addition to the true value of $-\sigma$, the biased coefficient contains the selection bias $\pi = \frac{cov(\epsilon_{ij},\tau_{ij}|s_i=1)}{var(\tau_{ij}|s_i=1)}$, as country i's income is part of the error term ϵ_{ij}. If the selection of the sample were indeed random, π would equal zero, and the elasticity would be estimated without error; i.e., $-\sigma_s = -\sigma$. Unfortunately, this is not the case, and the covariance term $cov(\epsilon_{ij}, \tau_{ij}|s_i = 1) \neq 0$, as income is also part of the error term ϵ_{ij}.

$$-\sigma_s = -\sigma + \underbrace{\frac{cov(\epsilon_{ij}, \tau_{ij}|s_i = 1)}{var(\tau_{ij}|s_i = 1)}}_{\text{selection}} = -\sigma + \pi.$$

While the conditional covariance and variance are not directly observable, we can back out the selection bias by comparing $-\sigma_s$, which is the resulting coefficient when we estimate the tariff elasticity only for the WITS sample, and the true $-\sigma$. The selection most likely pushes down the tariff elasticity ($-\hat{\sigma} < -\sigma$) because low-income countries are less active in international trade and, at the same time, have higher tariffs on average.

4.3 Decomposition

The total bias $\Delta\ \sigma$ attributable to missing tariff observations consists of three competing elements: 1) the bias due to the nonclassical nature of the measurement error in the falsely interpolated tariffs, 2) the bias due to the differential nature of the measurement error, and 3) the bias due to positive selection. By comparing the four coefficients obtained from the data, $-\sigma$, $-\tilde{\sigma}_s$, $-\sigma_s$, and $b_{\tau,\tilde{\tau}}$, we can quantify each of the elements.

The total bias $\Delta\ \sigma$ equals the difference between the coefficient resulting from our estimating the tariff elasticity with the data downloaded from WITS $-\tilde{\sigma}_s$, i.e., the falsely interpolated tariffs and the selected sample, and the coefficient using the new global tariffs database (GTD) that corrects for both issues $-\sigma$. The total bias is the sum of the biases attributable to false interpolation $\Delta\ \sigma^M$ and selection $\Delta\ \sigma^S$. The bias attributable to measurement error $\Delta\ \sigma^M$ equals $-\tilde{\sigma}_s - (-\sigma_s)$, and we can further decompose it by regressing the corrupted tariffs $\tilde{\tau}_{ij}$ and the correct ones τ_{ij}, which will give $b_{\tau,\tilde{\tau}}$, the element of the bias attributable to the nonclassical nature of the measurement error; Ω is then a residual. Last, the selection bias equals $\Delta\ \sigma - \Delta\ \sigma^M$.

$$\Delta\ \sigma = \Delta\ \sigma^M + \Delta\ \sigma^S = \underbrace{\overbrace{-\sigma_s(b_{\tau,\tilde{\tau}} - 1)}^{\textit{nonclassical ME } \lambda} + \overbrace{\Omega}^{\textit{differential ME}}}_{\text{false interpolation}} + \underbrace{\pi}_{\text{selection}} \tag{4}$$

Note that the decomposition is a general result and extends to various empirical approaches that use ordinary least squares to estimate the tariff elasticity, such as those accounting for multilateral resistance or time-differencing, which are common approaches to overcoming endogeneity issues in the empirical tariff elasticity literature.

5 Missing Tariffs in Practice: A Replication Study

In view of all of the issues present in the tariff data available through WITS, the question arises: how widely have these data been used in empirical work? In total, I identify 125 papers using WITS data for analysis that have been published in prominent journals since 2010, out of which 21 have been published in leading journals often referred to as "top 5".[28] Unsurprisingly, the majority of the remaining papers were published in the *Journal of International Economics*, the top filed journal in international economics (76 papers); 13 have been published in the *AEJ: Economic Policy*, *AEJ: Microeconomics*, *Macroeconomics*, or*AEJ: Applied Economics*, 14 papers in the *Review of Economics and Statistics*, and one paper in the *Journal of the European Economic Association*.

I used Google Scholar to compile a list of all relevant papers. As researchers have many ways of referring to the WITS tariff data, it was not possible to perform a Google Scholar search that contained all possible combinations. Therefore, I first searched for all published papers in the considered thirteen journals that contain the word "tariffs" anywhere in the full text and downloaded the corresponding paper to build a corpus of papers that potentially use the WITS data. Then, in a second step, I searched within all papers that contain the word "tariffs" for at least one of the following keywords: "WITS", "world integrated trade solution", "W.I.T.S.", "UNCTAD-Trains", "Trains", "T.R.A.I.N.S.", "Trade Analysis and Information System", and "UNCTAD". Last, I checked all the remaining papers one by one to eliminate any remaining false positives; for example, Shapiro (2021) uses WITS to extract information on nontariff barriers (see Appendix B.1 for more details on the compilation).

For the replication study, I will focus on the most influential contributions, i.e., the 21 papers that have been published in leading journals. Table 2 summarizes all of them. Four papers use tariff variation to identify the trade elasticity, six papers are interested in the effects of bilateral or multilateral trade policy, seven papers study firm-level outcomes, and four papers use tariffs to simulate shocks in quantitative models.

Some authors address the problem of missing tariffs head on, typically when the focus is on a specific policy context and the flaws in the data series are therefore apparent. For example, when analyzing the effects of rules of origin under NAFTA, Conconi et al. (2018) replace the missing preferential tariff observations for Mexican imports from the United States with data from the subsequent year, and Head and Mayer (2019) fill in missing preferential tariffs

[28] I consider publications in the *American Economic Review*, *Quarterly Journal of Economics*, *Journal of Political Economy*, *Econometrica*, and *Review of Economic Studies* as the leading journals. Furthermore, I also included publications in the *Review of Economics and Statistics*, *AEJ: Economic Policy*, *AEJ: Microeconomics*, *Macroeconomics*, *AEJ: Applied Economics*, *Journal of the European Economic Association*, and the *Journals of International Economics*.

for cars by going through the relevant trade agreements. However, the difficulty of spotting errors in the WITS tariff data varies significantly with the exact context. Especially in cross-country studies, the sheer size of the data sets—each tariff observation has four dimensions, i.e., importer, exporter, product, and time—generates many possibilities along which error can be embedded in the data, particularly if researchers do not know where to look and what the characteristics of the error are.

Unfortunately, the authors' use of confidential microdata makes replication infeasible for almost half of the papers using WITS data; unsurprisingly, none of the firm-level studies can be rerun.[29] As my goal is to determine the bias attributable to missing tariff data in the regression outcomes and to decompose the error due to this missingness, I will exclude the four papers that use tariffs as a shock in their respective quantitative model. Of the remaining papers, three rely on MFN tariff variation, for which the errors in the data are negligible: false interpolation is not an issue for these observations, and selection is also much less pronounced because WITS always reports MFN tariffs t_{ikt} as long as country i imports product k in year t from at least one partner. The two papers focusing on a single country are also excluded as the error is expected to be negligible in both. Instead, the studies estimating the trade elasticity will most likely suffer from substantial biases: large discrepancies in the relevant data can be expected given that these studies use MFN and preferential tariffs in a cross-country setting. Hence, my replication study will focus on these papers, i.e., Arkolakis, Ramondo, et al. (2018), C. E. Boehm et al. (2023), and Caliendo and Parro (2015).

[29] For Rodney D. Ludema and Anna Maria Mayda (2013), I cannot find the replication data on the journal website.

Table 2: Papers Using WITS Data and Published in Leading Journals

Paper	data availab.	Tariff Types	Aggregation	Countries	Time Period
Trade Elasticity					
Arkolakis, Ramondo, et al. (2018)	partly	MFN and pref.	pairs	62 countries	1999
Brancaccio et al. (2020)	no	MFN and pref.	pairs (bulk-commodities)	worldwide	2010-2016
C. E. Boehm et al. (2023)	yes	MFN and pref.	HS6-level	worldwide	1995-2018
Caliendo and Parro (2015)	yes	MFN and pref.	WIOD sectors	16 countries	1993
Effects of Bilateral Trade Policy					
Conconi et al. (2018)	yes	MFN and NAFTA-pref.	HS6-product	Mexico	1991 and 2003
Rafael Dix-Carneiro and Brian K. Kovak (2017)	no	MFN and pref.	44 aggregated industries	Brazil	1998-2010
Handley and Nuno Limão (2017)	yes	MFN and Column 2	HS6-level	United States	2000 and 2005
Multilateral Trade Policy					
Bagwell and Staiger (2011)	yes	Applied and bound MFN	HS6-level	16 recent-WTO countries	1995-2002
Rodney D. Ludema and Anna Maria Mayda (2013)	no	Applied and bound MFN	4-digit ISIC	36 countries	1993-2000
Nicita et al. (2018)	yes	Applied and bound MFN	HS6-digit	100 importer	1988-2007
Firm-Level Studies					
Alfaro et al. (2016)	no	MFN and pref.	4-digit (SIC)	United States (imposed and faced)	2000
J. Boehm et al. (2022)	no	MFN and pref.	3-digit industry-level	India	2000-2001 and 2009-2010
Loren Brandt et al. (2019)	no	MFN and pref.	4-digit national system	China	1994-2007
Bustos (2011)	no	MFN	4-digit ISIC industries	Brazil and Argentina	1991-1996
Head and Mayer (2019)	no	MFN and pref.	Cars and parts	worldwide	2000-2016
Keller and Yeaple (2013)	no	MFN and pref.	48 industries	faced by US exporters	1994, 1999, and 2004
Ahmad Lashkaripour and Volodymyr Lugovskyy (2023)	no	MFN and pref.	WIOD sectors	worldwide	2014
Quantitative Models					
Bagwell, Staiger, and Yurukoglu (2021)	yes	MFN	49 industries	worldwide	1990 and 2000
Burstein and Vogel (2017)	yes	Applied and bound MFN	2-digit level	60 countries	2007
Caliendo, Parro, et al. (2021)	yes	MFN and pref.	WIOD sectors	EU15 and NMS	2000-2014
Lind and Ramondo (2023)	yes	MFN and pref.	2-digit SITC industries	WIOD countries	2000-2014

Note: The table contains all the papers featuring WITS data and published in leading journals (*American Economic Review*, *Quarterly Journal of Economics*, *Journal of Political Economy*, *Econometrica*, and *Review of Economic Studies*) since 2010; details can be found in the main text. Ten of the papers are excluded from the replication study because they use confidential microdata. When MFN tariffs are not further specified, authors use the applied MFN tariff.

5.1 Arkolakis et al. (2018)

Arkolakis, Ramondo, et al. (2018, henceforth ARRY) develop a new model of trade and multinational production that they use to quantify the welfare implications of shocks driving increased specialization in innovation and production. The model delivers a novel testable implication—namely, that trade flows restricted to the parents and affiliates of firms from a given country are more sensitive to trade costs than overall, or unrestricted, trade flows. Using high-quality data from the Bureau of Economic Analysis (BEA) on the sales of US firms and their foreign affiliates, the authors take their theory to the data and estimate restricted and standard gravity equations. The two estimated elasticities then serve as key targets in the calibration of their model. ARRY rely on cross-sectional variation in tariffs to estimate the restricted and unrestricted trade elasticities. The gravity regressions use unbalanced pair-level data for 62 countries in 1999 and account for the important multilateral resistance terms with importer and exporter fixed effects (James E. Anderson and Van Wincoop 2003). In addition, a vector of standard gravity controls $\boldsymbol{Z_{ij}}$ (bilateral distance, contiguity, same language, colonial history, and domestic trade) is included.

For the restricted gravity estimates, the data on firm-level trade flows are not available for reasons of confidentiality. Therefore, I will analyze only how sensitive the results of the unrestricted gravity equation are to the issues in the WITS data caused by missing tariffs. The regression equation for the unrestricted gravity model is

$$lnX_{ij} = \sigma\tau_{ij} + \sum_k \delta^k[1|d_{ij} \in d_k] + \gamma \boldsymbol{Z_{ij}} + \mu_i + \mu_j + \epsilon_{ij}. \tag{5}$$

To compile information about bilateral tariffs, the authors downloaded information about MFN tariffs form WITS and combined it with a dummy variable that indicates whether country i and j have a trade agreement. ARRY assume that whenever there is a trade agreement, the applied tariff equals zero, and the MFN tariff otherwise. While this interpolation method seems very sensible, I will show that, unfortunately, it produces erroneous data as well.

ARRY download MFN tariffs for the relevant pairs at the HS6-level. While MFN tariffs do not suffer from any false interpolation, the issue of positive selection is very much present: for every pair-product combination only MFN tariffs with positive trade are reported. ARRY fill-in the missing observations at the HS6 level for every importing country with the HS6-average across all partners. For reporting countries that do not import all 5,000 HS6-products, this procedure will not necessarily produce the same simple average MFN tariff at the pair level as when using the GTD. For example, if countries do not trade products with particularly high tariffs, the pair-level average using the GTD will be higher than the one that ARRY's method produces. Hence, while the MFN tariffs at the HS6-level might not contain any measurement

error, the suggested handling of the selection bias most likely induces additional measurement error for the pair-level averages, i.e. $\tilde{\tau}_i = \tau_i + u_i$ where u_i can be positive or negative.

Furthermore, interpolating preferential tariffs always with zero will also add measurement error. In contrast to the false interpolation method used by WITS, it is unclear whether this practice yields corrupted tariff data $\tilde{\tau}_{ij}$ with values smaller or larger than the true values τ_{ij}: On the one hand, the fact that trade agreements usually include phase-outs and exemptions implies that trade agreements do not eliminate all tariffs. Hence, for the country pairs that ARRY identify as having a trade agreement, $\tilde{\tau}_{ij} \leq \tau_{ij}$. On the other hand, not all trade agreements are coded in the first place. For example, the association agreements between the EU and the Eastern European countries that acceded in 2004 were already in place in 1999, but ARRY assume that the United Kingdom—an EU member state during their sample period—imposed MFN tariffs on them, leading to $\tilde{\tau}_{ij} \geq \tau_{ij}$.

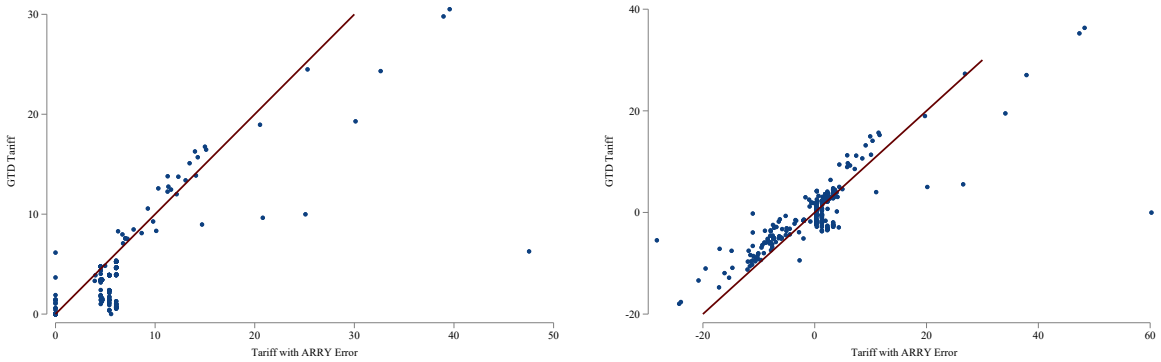
These two issues together, lead to a different error structure than when using the falsely interpolated data available through WITS. Figure 6 compares the pair-level GTD tariffs with the ones used in ARRY. It becomes immediately apparent, that the error in the data is much closer to the case classic ME that is depicted in Figure 5 (Panel (b)), as the corrupted data aligns much closer around the 45-degree line than in the case of non-classical ME. The importer and exporter fixed effects exacerbate the problem of data noisiness, a well-known result. To see why this happens, it is helpful to think about how fixed effects transform the data. When fixed effects are included, the group means will be subtracted from all observations. However, the group mean $\bar{\tau}$ will be measured with error as well. Therefore, when noisy means $\bar{\tilde{\tau}}$ are subtracted from every observation, even formerly clean data become corrupted. Panel (b) of Figure 6 illustrates this point.

Column (1) in Table 3 shows the results of my rerunning Equation (5) with the available data and code provided by the authors for replication.[30] I successfully reproduce the estimates from the paper, which give a coefficient of -4.3. The standard errors differ slightly compared to the paper because the authors run a seemingly unrelated regression using both, restricted and unrestricted data, which are not available to me. To determine the bias, I use the newly constructed tariff data at the HS6 level, aggregate by averaging over all products within a pair, and use those tariffs to estimate Equation (5). This exercise gives σ, i.e., the true elasticity that is plagued neither by false interpolation nor by positive selection. To decompose the total bias into its distinct elements, I will first determine how much of the bias is due to false interpolation. Instead of using all pair–product ijk observations from the GTD, I will keep only those observations that are also available in WITS at the HS6 level to keep the sample

[30] The authors kindly provided me with additional information about the relevant sample that is necessary to reproduce the estimates from the paper and is not available in the replication files.

Figure 6: Measurement Error in ARRY

(a) Tariffs (in %) (b) Demeaned Tariffs ($\tau - \dot{\tau}$)



Note: The graph compares the simple average of pair-level tariffs when using the GTD tariffs with the ones used in ARRY. Panel (a) shows tariffs in %, Panel (b) gives the demeaned tariffs $\tau_{ij} - \dot{\tau}_{ij}$ with $\dot{\tau}_{ij} = \tau_{ij} - \bar{\tau}_i - \bar{\tau}_j + \bar{\tau}_{ij}$).

constant and then average over products. This simple average of tariffs for the WITS sample will then be used to estimate σ_s.

My results suggest that the ARRY data underestimate the true σ by 5.36 with a coefficient of -9.64. The bias is almost entirely attributable to the false interpolation, as the coefficients for σ and σ_s are very similar (-9.64 and -8.97, respectively). This result is not surprising as ARRY make an effort to fill up observations that are missing due to WITS only reporting tariffs for pair-products with positive UN Comtrade flows and at the same time introduce more noise to the data. Interestingly, using the corrected version of the tariff data, produces estimates that are in the same ballpark as an alternative empirical approach that ARRY propose in their paper that does not rely on tariff variation and yields a coefficient of -9.7.

Both the bias due to the nonclassical measurement error and that due to the differential nature of the error attenuate the estimates. The positive sign of Ω might seem at first glance counterintuitive because, in general, one would expect Ω to lead to an overstated, i.e., more negative, elasticity because of the differential nature of the measurement error. However, in contrast to the WITS data, in the ARRY data, the true tariff rate could also be higher than the corrupted rate when missing preferential tariffs are mistakenly set to zero. Hence, the error u_{ij} becomes negative for country pairs with trade agreements, while it remains zero for countries with MFN relations. With this pattern of the error in mind, it is easy to see that the sign of the usually negative covariance between ϵ_{ij} and u_{ij} now might flip, leading to the positive Ω.

Table 3: Results: Replication Study

Paper	Coefficients				Decomposition				
	$-\tilde{\sigma}_s$	$-\sigma$	$-\sigma_s$	$b_{\tau,\tilde{\tau}}$	$\Delta\ \sigma$	$\Delta\ \sigma^S$	$\Delta\ \sigma^M$	λ	Ω
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Arkolakis, Ramondo, et al. (2018)	-4.28*	-9.64***	-8.97***	0.67***	5.36	0.66	4.69	2.97	1.72
	(2.36)	(3.17)	(3.00)	(0.03)		(12%)	(88%)	(63%)	(37%)
Caliendo and Parro (2015)	-4.55***	-0.52*	-1.56***	0.39***	-4.03	-1.04	-2.98	0.96	-3.94
	(0.35)	(0.29)	(0.31)	(0.02)		(26%)	(74%)	(-32%)	(132%)
C. E. Boehm et al. (2023)									
— *short run*	-0.65***	-0.33***	-0.45***	-	-0.32	-0.12	-0.20	-	-
	(0.12)	(0.06)	(0.13)			(37%)	(62%)		
— *long run*	-2.23***	-0.61***	-1.48***	-	-1.62	-0.87	-0.75	-	-
	(0.35)	(0.14)	(0.357)			(54%)	(46%)		

Note: The table shows the results when replicating Arkolakis, Ramondo, et al. (2018, Table 1, column 1), Caliendo and Parro (2015, Table 1, aggregate elasticity), and C. E. Boehm et al. (2023, Figure 2, years 1 and 10).

5.2 Caliendo and Parro (2015)

In their seminal work, Caliendo and Parro (2015, henceforth CP) include sectoral linkages, trade in intermediate goods, and sectoral heterogeneity in production into a Ricardian model to quantify the trade and welfare effects of tariff changes. In their quantification exercise, the trade elasticity is one of the central parameters. Therefore, CP propose a new method for estimating the trade elasticity that has minimal data requirements and is consistent with any trade model that delivers a gravity equation.

For the estimation of $-\sigma$, consider three countries indexed by n, i, and h. Take the cross-product of goods shipped in one direction between the three countries, from n to i, from i to h, and from h to n, and then the cross-product of the same goods shipped in the other direction, from n to h, from h to i, and from i to n. Employing the gravity equation and taking the ratio, it follows that

$$\frac{X_{ni}X_{ih}X_{hn}}{X_{nh}X_{hi}X_{in}} = \left(\frac{\kappa_{ni}\kappa_{ih}\kappa_{hn}}{\kappa_{nh}\kappa_{hi}\kappa_{in}}\right)^{-\sigma}. \tag{6}$$

All the terms involving prices and other parameters cancel out, and bilateral trade is a function only of trade costs κ, which consist of tariffs $\tau_{ni} = ln(1 + t_{ni})$ and nontariff trade costs d_{ni}. The nontariff trade costs can be modeled quite generally as a linear function of bilateral symmetric trade costs $\mu_{ni} = \mu_{in}$ and importer and exporter characteristics ν_i and ν_n. All remaining asymmetric trade costs incurred in exporting from n to i are summarized in ϵ_{ni}. Hence, κ_{ni} can be rewritten as

$$ln(\kappa_{ni}) = ln(\tau_{ni}) + ln(d_{ni}) = ln(\tau_{ni}) + \mu_{ni} + \nu_n + \nu_i + \epsilon_{ni}.$$

The main regression equation follows after we take logs and plug the nontariff trade costs into (6).

$$ln\left(\frac{X_{ni}X_{ih}X_{hn}}{X_{nh}X_{hi}X_{in}}\right) = -\sigma\ ln\left(\frac{\tau_{ni}\tau_{ih}\tau_{hn}}{\tau_{nh}\tau_{hi}\tau_{in}}\right) + \left(\frac{\epsilon_{ni}\epsilon_{ih}\epsilon_{hn}}{\epsilon_{nh}\epsilon_{hi}\epsilon_{in}}\right). \tag{7}$$

CP run their regressions at the World Input–Output Database (WIOD) sector level. They download tariffs for International Standard Industrial Classification (ISIC) industries, fill in potentially missing tariffs with the values from preceding and subsequent years, and then aggregate to the WIOD sector level. To obtain the tariffs that keep the WITS sample constant, I need to mimic exactly what WITS does when the researcher downloads data—i.e., use the positively selected observations at the HS6 level and aggregate to the ISIC industry level. I do this for the GTD for the the years 1989–1995 and then run the infilling algorithm proposed by CP and aggregate to the WIOD sector level. Using these tariffs will give σ_s. When the goal is to determine the true σ that does not suffer from sample selection, I simply use the GTD for the year 1993 and aggregate to WIOD sectors.

The results of the replication show that CP drastically overstate the elasticity: instead of -4.55, my use of the GTD gives $-\sigma = -0.52$. While the bias due to positive selection is much more important here than it was for ARRY, the majority of the bias is still due to the false interpolation. To understand why, it is helpful to revisit how the falsely interpolated tariffs bias the estimates. The triple differentiation proposed by CP is an elegant way of parsimoniously controlling for unobservables that might bias the estimates of $-\sigma$. The main assumption is that once all symmetric trade costs and country-specific characteristics are accounted for, the remaining asymmetric determinants of trade costs ϵ_{ni} are orthogonal to tariffs τ_{ni}. When Equation (7) is estimated with the falsely interpolated tariffs $\tilde{\tau}_{ni} = \tau_{ni} + u_{ni}$, this assumption is violated. The false interpolation leads us to falsely assume that countries n and i have MFN relations. The triple differentiation will take out the level effect of the MFN tariff on bilateral trade, but the tariff margin, i.e., the difference between the MFN and preferential tariffs, which equals u_{ni}, will be pushed into the error term ϵ_{ni}. As established before, the error u_{ni} is not random but instead differential and nonclassical, violating the identifying assumption, i.e., $cov(\tau_{ni}, \epsilon_{ni}) \neq 0$. As the decomposition shows, most of the bias attributable to the false interpolation is associated with Ω, i.e., the differential nature of the error.

My using the corrected version of the tariffs fails to give meaningful estimates for the tariff elasticity $-\sigma$, which needs to be smaller than -1 to be theory consistent; when I estimate the elasticities sector by sector, some of the coefficients even become positive (see Appendix A2).

Ossa (2016) reports similar results upon trying to use the CP method in a different sample than in the original article.

The lack of variation in tariffs could serve as a potential explanation for this finding. Recall that the triple differentiation approach cancels out not only all MFN barriers but also all symmetric country-pair costs—i.e., all identical preferential tariffs between signatory parties are canceled out. Why is this important? Article XXIV of the GATT/WTO says that RTAs are allowed only if they eliminate tariffs on "substantially all trade". Therefore, the vast majority of tariffs for members of the same RTA are in fact zero, i.e., symmetric between country n and i, and, as such, are canceled out by the CP method. The CP sample that is used to estimate the elasticities contains only 16 countries. In 1993, eight out of 240 country pairs had a bilateral trade agreement, and for them, 15% of all the tariffs at the HS6 level were asymmetric; aggregating to the industry level will most likely wipe out even more of the remaining variation because many of the asymmetric tariffs are in the same industry. The remaining asymmetric tariffs correspond to products for which the signatory parties have negotiated either phase-out periods or asymmetric exemptions. Even if we were willing to assume that the asymmetry in these tariffs is exogenous—quite a strong assumption given that the remaining tariffs fall on products that enjoy extra protection in one of the countries—little variation would remain, especially after we aggregate to WIOD sectors.

In addition to the bilateral RTAs notified under Article XXIV, trade agreements of another type were in place in 1993, namely, agreements notified to the WTO under the enabling clause. The purpose of such agreements is to help low-income countries grow through trade, and therefore, their terms may deviate from the "substantially all trade" criterion. This leaves in principle a great deal of variation for identification. Some 72% of country pair–HS6 product observations in the CP sample with an agreement under the enabling clause correspond to tariffs set asymmetrically, and 42% of the country pairs have such agreements. However, these agreements are notorious for their low utilization—i.e., exporters do not use the available preferences but instead keep paying the MFN tariffs. Costly compliance with rules of origin might explain this underutilization. Unfortunately, even with perfectly clean tariff data, the underutilization is not observable and would push the CP estimates toward zero.

5.3 Boehm et al. (2023)

Next, I will discuss how missing tariffs bias the estimates in C. E. Boehm et al. (2023) (henceforth BLP), who determine the trade elasticity for different time horizons using tariff variation. To account for omitted variables that vary by country pair–product, they time-difference the data. Changes in trade policy might still be endogenous as, for example, policymakers might lower tariffs in response to a large increase in trade. To address these kind of endogeneity is-

sues, BLP propose a novel and very innovative identification strategy: they compare imports from countries with MFN relations to imports from countries unaffected by changes in MFN tariffs, i.e., countries with a trade agreement. Crucially, the estimates are based only on the response of minor exporters, for which changes in MFN tariffs are plausibly exogenous. The exclusion of large trade partners is important for identification because changes in their trade policy might be endogenous—for example, because of lobbying activity by multinational enterprises (Blanchard and Matschke 2015). Furthermore, BLP account for autocorrelation in tariff changes using local projections (Jordà 2005). The exogenous year-on-year changes in MFN tariffs are used as an instrument for future tariff changes to identify responses of trade at different horizons.

The instrument z_{ijkt} equals the change in MFN tariffs between t and $t-1$ if country pair–products have MFN relations in t and $t-1$ and equals zero otherwise (see Equation (8)).

$$\begin{aligned} z_{ijkt} = &\left(ln\tau_{ikt}^{MFN} - ln\tau_{ikt-1}^{MFN}\right) \\ &\times \mathbf{1}\left(\tau_{ijkt} = \tau_{ikt}^{MFN}\right) \times \mathbf{1}\left(\tau_{ijkt-1} = \tau_{ikt-1}^{MFN}\right). \end{aligned} \tag{8}$$

Unfortunately, the instrument is corrupted because of the false interpolation done by WITS: for importers i that fail to report preferential tariffs with partner j in two consecutive years, BLP falsely assume that there is no trade agreement, and the instrument will be equal to the change in MFN tariffs instead of zero, i.e., $\tilde{z}_{ijkt} = z_{ijkt} + u_{ijkt}$. Note that the error u_{ijkt} could be smaller or greater than the true instrument z_{ijkt} because it is a priori unclear whether MFN tariffs increase or decrease in the respective year that the country fails to report its tariffs.

The instrument is used to estimate the following first stage, with major partners excluded:

$$\Delta_h ln\tau_{ijkt} = \beta^h z_{ijkt} + \delta_{ikt}^h + \delta_{jkt}^h + \delta_{ijk}^h + u_{ijkt}^h. \tag{9}$$

$\Delta_h ln\tau_{ijkt}$ denotes the time difference in $ln\tau_{ijkt}$ between periods $t-1$ and $t+h$. Since the dependent variable of the first stage is also constructed from WITS tariff data, it suffers from false interpolation as well, i.e., $\Delta_h ln\tilde{\tau}_{ijkt} = \Delta_h ln\tau_{ijkt} + \Delta_h u_{ijkt}$. For pairs with preferential trade relations that fail to report their preferential tariffs, BLP falsely assume that the instrumental variable equals the change in MFN tariffs instead of being zero. Hence, the reported tariff in $t-1$, the base year for determining the horizon h change in future tariff changes, will be a false MFN tariff, too. Therefore, the measurement error in the dependent variable $\Delta_h u_{ijkt}$ is negative whenever countries correctly report their preferential tariffs in $t+h$ since tariffs mechanically drop from the high MFN tariff levels to the low preferential levels in $t+1$. As reporting compliance rises over time while MFN tariffs decrease over time, the error in the future tariff changes will be larger for longer horizons.

The second stage is then

$$\Delta_h lnX_{ijkt} = -\sigma^h \Delta_h ln\tau_{ijkt} + \delta^h_{ikt} + \delta^h_{jkt} + \delta^h_{ijk} + u^h_{ijkt}. \tag{10}$$

Table 3 gives the result of the replication analysis when I rerun the code made available by the authors and use the data provided by WITS and the GTD. First, note that the replicated coefficients are slightly different from those that BLP report in their paper because of minor changes to the sample that were necessary to allow comparison of the BLP results with those from the GTD. Recall that the GTD stops in 2017 and has data available only for HS6 products in the 1988/92 nomenclature. Trimming the sample accordingly gives a somewhat lower number of observations.[31] As Figure A3 shows, the results of BLP are very robust to this change in sample.

The missing tariffs lead to overestimation of the absolute value of the tariff elasticity: my use of the GTD gives a coefficient of -0.61 instead of -2.23 for the long-run and -0.33 instead of -0.65 for the short-run elasticity. Figure 7 summarizes the results for the other time horizons. While the time path of the elasticities over different horizons is similar in the sense that the short-run elasticities are much smaller than the long-run elasticities, the stabilization occurs earlier than in BLP, namely, after five years instead of seven. Furthermore, both corrections—false interpolation as well as selection—lead to lower absolute tariff elasticities; however, including the full sample instead of the selected WITS sample yields the largest change in coefficients for all time horizons, i.e., $-\tilde{\sigma}_s < -\sigma_s < -\sigma$.

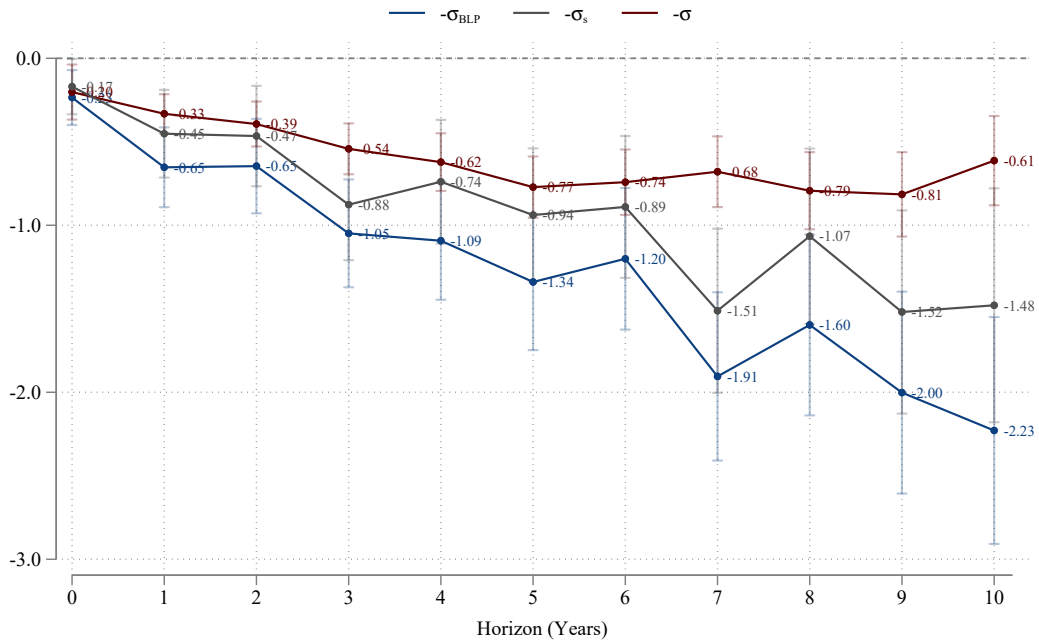
To understand what drives the larger absolute coefficients when I use the WITS data, it is helpful to formalize the IV estimator and split the analysis by the type of problem addressed, that is, false interpolation or positive selection. As Equation (11) shows, $-\tilde{\sigma}_s$ is defined as the ratio of the reduced form $cov(\Delta_h lnX, \tilde{z})$ and the first stage $cov(\Delta_h ln\tilde{\tau}, \tilde{z})$; all indices are suppressed for ease of notation.

$$-\tilde{\sigma}_s = \frac{cov(\Delta_h lnX, \tilde{z})}{cov(\Delta_h ln\tilde{\tau}, \tilde{z})} = \frac{cov(\Delta_h lnX, z) + cov(\Delta_h lnX, u)}{cov(\Delta_h ln\tau, z) + cov(\Delta_h ln\tau, u) + cov(\Delta_h u, z) + cov(\Delta_h u, u)}. \tag{11}$$

As Table A3 shows, the first stage is lower when the tariff data have been cleansed of false interpolation. Following the initial impulse, 65% of the change remains after one year and 62% after ten years, whereas the WITS data suggest these numbers to be 82% and 71%, respectively. Put differently, using the GTD tariffs gives a higher degree of mean reversion.

[31] The number of observations drops from 26.2 million to 19.7 million for the short-run analysis and from 8.3 million to 6.5 million for the long-run analysis.

Figure 7: Replication of Boehm et al. (2023)



Note: The graph compares the tariff elasticity over different horizons when I use the WITS data and the GTD data. The blue line shows the results by BLP using the WITS data, the gray line accounts only for false interpolation, while the red line also corrects for the positive selection in the WITS data.

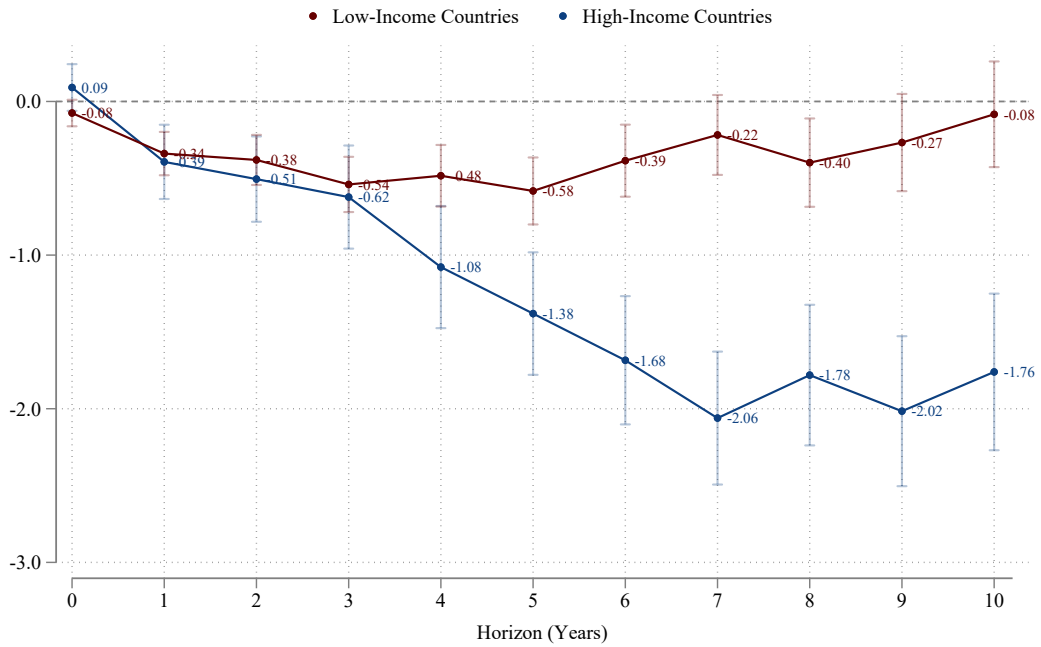
These differences in the coefficients are driven by a mechanically high level of persistence for the observations with the wrongly classified instrument z. Recall that for observations for which the instrument is wrong $z = \tilde{z}$, the h-horizon change in tariffs is mechanically negative and most likely relatively large as tariffs jump from the wrongly coded MFN tariff in $t-1$ to the true (low) preferential tariff in $t+h$. At the same time, most of the year-on-year changes for country pairs with MFN relations correspond to reductions and equal on average three percentage points (excluding observations with zero changes). This pattern reflects that MFN tariffs have been reduced worldwide since 1995. Hence, for the observations with the corrupted instrument, the falsely imputed WITS data wrongly indicate strong persistence overstating the coefficient of the first stage.

All else equal, the larger first-stage estimate attenuates $-\tilde{\sigma}_s$; consequently, $-\tilde{\sigma}_s < -\sigma_s$ must result from the lower reduced-form estimate when I correct for the false interpolation. Equation (11) shows that a smaller reduced-form estimate is possible only if the covariance between the horizon-h change in trade and the error in the instrument is negative. To recover the sign of $cov(\Delta_h lnX, u)$, I regress $\Delta_h lnX$ on u, including all controls from the main specification. The resulting coefficients are negative and larger in absolute terms for $h = 10$ than for $h = 1$ (-0.79 and -0.31, respectively c.f. Table A4).

Correcting for the false interpolation noticeably changes the coefficients reported by BLP, but they become statistically different from each other only when I also correct for the positive selection. This indicates treatment heterogeneity across the different samples: trade between low-income countries reacts differently from that of the remaining observations also contained in the WITS sample. A sample split using the GTD confirms this. As Figure 8 shows, the elasticity for high-income and low-income importers is relatively similar up until $h = 3$, when they start to deviate. While for the high-income countries it decreases until $h = 7$ and then fluctuates around -2, for low-income countries the coefficients remain over all horizons small and become even insignificant in the long-run. In fact, the results for the high-income countries look very similar to the ones reported by BLP stressing again the importance of the selection bias in this setting. This results relates to Anson Soderbery (2018) as it highlights heterogeneity in the tariff elasticity across-countries and has important implications on how we think about the trade elasticity.

There are at least two potential explanations for this finding: First, the change in MFN tariffs of low-income countries might be perceived by exporters as temporary. The resulting higher level of uncertainty might hamper the reaction of trade flows (cf. Handley and Nuno Limão 2015; Handley and Nuno Limão 2017). Second, low-income exporters might be capacity constrained and therefore might barely react to tariffs changes. Most pairs have MFN instead of preferential relations. Therefore, a change in MFN tariffs of importer i will alter trade

Figure 8: Heterogeneous Effects by Income: Boehm et al. (2023)



Note: The graph compares the tariff elasticity over different horizons using the GTD for different income-levels of the importing countries. Low-income countries are depicted in red, high-income countries in blue.

relations with most partners that might be, on top of that, more competitive than the exporters from the low-income countries.

6 Conclusion

This paper shows that a standard source of tariff data used in the literature suffers from substantial measurement error and selection bias due to missing tariffs. Many countries fail to report their tariffs every year, with reporting compliance systematically lower for preferential than for MFN tariffs. The World Bank's WITS, the main provider of tariff data, falsely interpolates the missing preferential tariffs with MFN tariffs. This practice leads to artificial spikes in bilateral time series data and hence induces substantial measurement error. Additionally, WITS presents tariffs only for importers that report trade to UN Comtrade. Therefore, low-income countries are underrepresented in the resulting sample, and tariffs for products with zero trade flows are discarded. I propose a new interpolation algorithm that takes this misreporting into account and combine five different sources of tariff data to introduce a new global tariff dataset at the six-digit product level covering 197 countries and 30 years. I illustrate the practical importance of the tariff missingness problem by replicating papers that were recently published in leading economics journals and that use the data provided by WITS. Estimates of the trade elasticity are particularly sensitive to the mishandling of the tariff data.

Almost twenty years ago, James E. Anderson and Van Wincoop (2004) wrote that *"the grossly incomplete and inaccurate information on policy barriers available to researchers is a scandal and a puzzle"* (p. 693), and the situation is not substantially improved today. Why is collecting data on tariffs so challenging? First, the sheer number of observations—197 importers, 196 exporters, and approximately 5,000 products, amounting to 200 million observations per year—makes handling tariff data very difficult, confusing, and therefore error prone. Second, the large-scale trade integration through RTAs in recent decades and the resulting multiplicity of different tariff schemes increase complexity. Last, data-collecting institutions rely on countries to report their tariffs but have very few tools to enforce reporting compliance (Hoda 2001).

To resolve the issue of misreporting, international institutions could make access to WITS conditional on reporting on all tariff schemes, i.e., all preferential and MFN tariffs. In the past, reporting compliance could be significantly raised when bad reporters were presented with the threat of temporary suspension of access to WITS (Hoda 2001, p. 207). The WTO is the ideal institution to collect tariffs because it also has the most reliable information on the presence of RTAs. Leveraging information on RTAs that are in force, the body could make it obligatory to report all tariff schemes for each of the RTAs currently in place. This procedure would ensure full coverage of tariff data and eliminate any potential reason for relying on (false) interpolation.

References

Aghion, P, J Cai, M Dewatripont, and L Du (2015), "Industrial policy and competition." *American Economic Journal: Macroeconomics*.

Ahn, Jae Bin, Amit K. Khandelwal, and Shang Jin Wei (2011), "The Role of Intermediaries in Facilitating Trade." *Journal of International Economics* 841, 73–85. URL: http://dx.doi.org/10.1016/j.jinteco.2010.12.003.

Aichele, R and I Heiland (2018), "Where Is the Value Added? Trade Liberalization and Production Networks." *Journal of International Economics*.

Albornoz-Crespo, Facundo, Irene Brambilla, and Emanuel Ornelas (2021), "Firm Export Responses to Tariff Hikes." *CEPR Discussion Papers 16455*. URL: https://ideas.repec.org/p/cpr/ceprdp/16455.html%20https://ideas.repec.org//p/cpr/ceprdp/16455.html.

Alfaro, Laura, Paola Conconi, Harald Fadinger, and Andrew F. Newman (2016), "Do Prices Determine Vertical Integration?" *Review of Economic Studies* 833, 855–888.

Alquist, R, N Berman, R Mukherjee, and L L Tesar (2019), "Financial Constraints, Institutions, and Foreign Ownership." *Journal of International Economics*.

Amiti, M and L Cameron (2012), "Trade liberalization and the wage skill premium: Evidence from Indonesia." *Journal of International Economics*.

Amiti, M and A K Khandelwal (2013), "Import Competition and Quality Upgrading." *Review of Economic Statistics*.

Anderson, J E and Y V Yotov (2020), "Short Run Gravity." *Journal of International Economics*.

Anderson, James E. and J. Peter Neary (1994), "Measuring the Restrictiveness of Trade Policy." *The World Bank Economic Review* 82, 151–169.

Anderson, James E. and Eric Van Wincoop (2003), "Gravity with Gravitas : A Solution to the Border Puzzle." *The American Economic Review* 931, 170–192.

— (2004), "Trade Costs." *Journal of Economic Literature* 423, 691–751.

Ardelean, A and V Lugovskyy (2010), "Domestic Productivity and Variety Gains from Trade." *Journal of International Economics*.

Arezki, R, T Fetzer, and F Pisch (2017), "On the Comparative Advantage of US Manufacturing: Evidence from the Shale Gas Revolution." *Journal of International Economics*.

Arkolakis, Costas, Sharat Ganapati, and Marc Andreas Muendler (2021), "The extensive margin of exporting products: A firm-level analysis." *American Economic Journal: Macroeconomics* 134, 182–245.

Arkolakis, Costas, Natalia Ramondo, Andrés Rodríguez-Clare, and Stephen Yeaple (2018), "Innovation and Production in the Global Economy." *American Economic Review* 1088, 2128–73.

Artuc, E, P Bastos, and B Rijkers (2023), "Robots, Tasks, and Trade." *Journal of International Economics*.

Artuc, E, G Porto, and B Rijkers (2019), "Trading Off the Income Gains and the Inequality Costs of Trade Policy." *Journal of International Economics*.

Baccini, L, G Impullitti, and E J Malesky (2019), "Globalization and State Capitalism: Assessing Vietnam's Accession to the {WTO}." *Journal of International Economics*.

Bagwell, Kyle and Robert W. Staiger (2011), "What Do Trade Negotiators Negotiate About? Empirical Evidence from the World Trade Organization." *American Economic Review* 1014, 1238–1273.

Bagwell, Kyle, Robert W. Staiger, and Ali Yurukoglu (2021), "Quantitative Analysis of Multiparty Tariff Negotiations." *Econometrica* 894, 1595–1631.

Baier, Scott L. and Jeffrey H. Bergstrand (2004), "Economic Determinants of Free Trade Agreements." *Journal of International Economics* 641, 29–63.

— (2007), "Do Free Trade Agreements actually Increase Members' International Trade?" *Journal of International Economics* 711, 72–95. URL: https://ideas.repec.org/a/eee/inecon/v71y2007i1p72-95.html.

Baier, Scott L., Jeffrey H. Bergstrand, and Michael Feng (2014), "Economic Integration Agreements and the Margins of International Trade." *Journal of International Economics* 932, 339–350. URL: http://dx.doi.org/10.1016/j.jinteco.2014.03.005.

Baier, Scott L., Yoto V. Yotov, and Tom Zylkin (2019), "On the Widely Differing Effects of Free Trade Agreements: Lessons from Twenty Years of Trade Integration." *Journal of International Economics*.

Baldwin, R E and E Seghezza (2010), "Are Trade Blocs Building or Stumbling Blocs?" *Journal of International Economics*.

Barari, Soubhik Barari and In Song Kim (2020), "Democracy and Trade Policy at the Product Level."

Barattieri, A, M Cacciatore, and F Ghironi (2021), "Protectionism and the business cycle." *Journal of International Economics*.

Bas, M, T Mayer, and M Thoenig (2017), "From Micro to Macro: Demand, Supply, and Heterogeneity in the Trade Elasticity." *Journal of International Economics*.

Bas, Maria and Vanessa Strauss-Kahn (2015), "Input-trade liberalization, export prices and quality upgrading." *Journal of International Economics* 952, 250–262.

Bastos, P, O R Straume, and J A Urrego (2013), "Rain, agriculture, and tariffs." *Journal of International Economics*.

Bergstrand, Jeffrey H., Stephen R. Cray, and Antoine Gervais (2023), "Increasing Marginal Costs, Firm Heterogeneity, and the Gains from "Deep" International Trade Agreements." *Journal of International Economics*.

Bernard, A B, M Grazzi, and C Tomasi (2015), "Intermediaries in International Trade: Products and Destinations." *Review of Economic Statistics*.

Besedes, Tibor, Tristan Kohl, and James Lake (2020), "Phase out tariffs, phase in trade?" *Journal of International Economics* 127.

Beshkar, M and E W Bond (2017), "Cap and escape in trade agreements." *American Economic Journal: Microeconomics*.

Beshkar, M and R Lee (2022), "How Does Import Market Power Matter for Trade Agreements?" *Journal of International Economics*.

Beverelli, C, M Fiorini, and B Hoekman (2017), "Services Trade Policy and Manufacturing Productivity: The Role of Institutions." *Journal of International Economics*.

Blanchard, Emily J. and Xenia Matschke (2015), "U.S. Multinationals and Preferential Market Access." *Review of Economics and Statisticss* 964, 839–854.

Boehm, Christoph E., Andrei A. Levchenko, and Nitya Pandalai-Nayar (2023), "The Long and Short (Run) of Trade Elasticities." *American Economic Review* 1134, 861–905.

Boehm, Johannes, Swati Dhingra, and John Morrow (2022), "The Comparative Advantage of Firms." *Journal of Political Economy* 13012, 3025–3100.

Bombardini, M and B Li (2020), "Trade, pollution and mortality in China." *Journal of International Economics*.

Bombardini, M, G Orefice, and M D Tito (2019), "Does Exporting Improve Matching? Evidence from French Employer-Employee Data." *Journal of International Economics*.

Bouët, Antoine, Yvan Decreux, Lionel Fontagné, Sébastien Jean, and David Laborde (2008), "Assessing Applied Protection across the World." *Review of International Economics* 165, 850–863.

Bound, John, Charles Brown, Greg J. Duncan, and Willard L. Rodgers (1994), "Evidence on the Validity of Cross-Sectional and Longitudinal Labor Market Data." *Journal of Labor Economics* 123, 345–368.

Bound, John, Charles Brown, and Nancy Mathiowetz (2001), "Measurement Error in Survey Data." *Handbook of Econometrics*. Vol. 5. Elsevier Masson SAS, 3705–3843. URL: http://dx.doi.org/10.1016/S1573-4412(01)05012-7.

Bound, John and Alan B. Krueger (1991), "The Extent of Measurement Error in Longitudinal Earnings Data: Do Two Wrongs Make a Right?" *Journal of Labor Economics* 91, 1–24.

Bown, Chad P. and Meredith A. Crowley (2016), "The Empirical Landscape of Trade Policy." *Handbook of Commercial Policy*. Ed. by Kyle Bagwell and Robert W. Staiger. Vol. 1A. Elsevier B.V. Chap. 1, 3–108.

Brancaccio, Giulia, Myrto Kalouptsidi, and Theodore Papageorgiou (2020), "Geography, Transportation, and Endogenous Trade Costs." *Econometrica* 882, 657–691.

Brandt, L, B Li, and P M Morrow (2021), "Processing Trade and Costs of Incomplete Liberalization: The Case of China." *Journal of International Economics*.

Brandt, Loren, Johannes Van Biesebroeck, Luhang Wang, and Yifan Zhang (2019), "WTO Accession and Performance of Chinese Manufacturing Firms." *American Economic Review* 1094, 1616–1621.

Breinlich, Holger, Anson Soderbery, and Greg C. Wright (2018), "From selling goods to selling services: Firm responses to trade liberalization." *American Economic Journal: Economic Policy* 104, 79–108.

Buono, I and G Lalanne (2012), "The effect of the Uruguay Round on the intensive and extensive margins of trade." *Journal of International Economics*.

Burstein, Ariel, Eduardo Morales, and Jonathan Vogel (2019), "Changes in between-group inequality: Computers, occupations, and international trade." *American Economic Journal: Macroeconomics* 112, 348–400.

Burstein, Ariel and Jonathan Vogel (2017), "International Trade, Technology, and the Skill Premium." *Journal of Political Economy* 1255, 1356–1412.

Bustos, Paula (2011), "Trade Liberalization, Exports and Technology Upgrading: Evidence on the Impact of MERCOSUR on Argentinean Firms." *American Economic Review* 1011, 304–340.

Caliendo, Lorenzo, Robert C. Feenstra, John Romalis, and Alan M. Taylor (2023), "Tariff Reductions, Heterogeneous Firms, and Welfare: Theory and Evidence for 1990-2010." *IMF Economic Review* forthcomin.

Caliendo, Lorenzo and Fernando Parro (2015), "Estimates of the Trade and Welfare Effects of NAFTA." *Review of Economic Studies* 821, 1–44.

Caliendo, Lorenzo, Fernando Parro, Luca David Opromolla, and Alessandro Sforza (2021), "Goods and Factor Market Integration: A Quantitative Assessment of the EU Enlargement." *Journal of Political Economy* 12912, 3491–3545.

Carballo, J, A G Graziano, and G Schaur (2022), "The Effects of Transit Systems on International Trade." *Review of Economic Statistics*.

Carluccio, J and M Bas (2015), "The Impact of Worker Bargaining Power on the Organization of Global Firms." *Journal of International Economics*.

Carrere, C and A Grujovic (2020), "Trade and Frictional Unemployment in the Global Economy." *Journal of the European Economic Association*.

Chakraborty, P and O Raveh (2018), "Input-trade liberalization and the demand for managers: Evidence from India." *Journal of International Economics*.

Chen, B, M Yu, and Z Yu (2017), "Measured skill premia and input trade liberalization: evidence from Chinese firms." *Journal of International Economics*.

Chen, Cheng and Claudia Steinwender (2021), "Import Competition, Heterogeneous Preferences of Managers, and Productivity." *Journal of International Economics*.

Chen, M X and M O Moore (2010), "Location decision of heterogeneous multinational firms." *Journal of International Economics*.

Chen, N and L Juvenal (2022), "Markups, quality, and trade costs." *Journal of International Economics*.

Coelli, F, A Moxnes, and K H Ulltveit-Moe (2022), "Better, Faster, Stronger: Global Innovation and Trade Liberalization." *Review of Economic Statistics*.

Conconi, Paola, Laura Puccio, Manuel García-Santana, Roberto Venturini, Laura Puccio, and Venturini (2018), "From Final Goods to Inputs: the Protectionist Effect of Rules of Origin." *American Economic Review* 1088, 2335–2365.

Coşar, A K, P L E Grieco, S Li, and F Tintelnot (2018), "What Drives Home Market Advantage?" *Journal of International Economics*.

Costinot, Arnaud and Andrés Rodríguez-Clare (2014), "Trade Theory with Numbers: Quantifying the Consequences of Globalization." *Handbook of International Economics*. Vol. 4. Elsevier B.V., 197–261. URL: http://dx.doi.org/10.1016/B978-0-444-54314-1.00004-5.

Dasgupta, K and J Mondria (2018), "Inattentive Importers." *Journal of International Economics*.

Davidson, C, F Heyman, S Matusz, and F Sjöholm (2014), "Globalization and Imperfect Labor Market Sorting." *Journal of International Economics*.

Davies, R B, J Martin, and M Parenti (2018), "Knocking on Tax Haven's Door: Multinational Firms and Transfer Pricing." *Review of Economic Statistics*.

De Melo, Jaime and Alessandro Nicita (2018), "Non-Tariff Measures: Data and Quantitative Tools of Analysis." *F.R.E.I.T. Working Paper*. URL: UNCTAD/DITC/TAB/2017/2.

Defever, F, M Imbruno, and R Kneller (2020), "Trade liberalization, input intermediaries and firm productivity: Evidence from China." *Journal of International Economics*.

Diez, F J (2014), "The asymmetric effects of tariffs on intra-firm trade and offshoring decisions." *Journal of International Economics*.

Disdier, Anne-Célia and Keith Head (2008), "The Puzzling Persistence of the Distance Effect on Bilateral Trade." *The Review of Economics and Statistics* 901, 37–48.

Dix-Carneiro, R and B K Kovak (2019), "Margins of Labor Market Adjustment to Trade." *Journal of International Economics*.

Dix-Carneiro, R, R R Soares, and G Ulyssea (2018), "Economic shocks and crime: Evidence from the brazilian trade liberalization." *American Economic Journal: Applied Economics*.

Dix-Carneiro, Rafael and Brian K. Kovak (2017), "Trade Liberalization and Regional Dynamics." *American Economic Review* 10710, 2908–2946.

Dür, Andreas, Leonardo Baccini, and Manfred Elsig (2014), "The Design of International Trade Agreements: Introducing a New Dataset." *Review of International Organizations* 93, 353–375.

Duval, R, N Li, R Saraf, and D Seneviratne (2016), "Value-added Trade and Business Cycle Synchronization." *Journal of International Economics*.

Egger, Peter H. and Mario Larch (2008), "Interdependent Preferential Trade Agreement Memberships: An Empirical Analysis." *Journal of International Economics* 762, 384–399. URL: http://dx.doi.org/10.1016/j.jinteco.2008.08.003.

Erten, B and P Keskin (2021), "Trade-offs? The Impact of {WTO} Accession on Intimate Partner Violence in Cambodia." *Review of Economic Statistics*.

Erten, B and J Leight (2021), "Exporting out of Agriculture: The Impact of {WTO} Accession on Structural Transformation in China." *Review of Economic Statistics*.

Esposito, F (2022), "Demand Risk and Diversification Through International Trade." *Journal of International Economics*.

Estevadeordal, A and A M Taylor (2013), "Is the Washington Consensus Dead? Growth, Openness, and the Great Liberalization, 1970s–2000s." *Review of Economic Statistics*.

Facchini, G, M Y Liu, A M Mayda, and M Zhou (2019), "China's "Great Migration": The Impact of the Reduction in Trade Policy Uncertainty." *Journal of International Economics*.

Fajgelbaum, Pablo D., Pinelopi K. Goldberg, Patrick J. Kennedy, and Amit K. Khandelwal (2020), "The Return to Protectionism." *The Quarterly Journal of Economics* 1351, 1–55. URL: https://doi.org/10.1093/qje/qjz036.

Fan, H, Y A Li, S Xu, and S R Yeaple (2020), "Quality, Variable Markups, and Welfare: A Quantitative General Equilibrium Analysis of Export Prices." *Journal of International Economics*.

Fan, H, F Lin, and S Lin (2020), "The Hidden Cost of Trade Liberalization: Input Tariff Shocks and Worker Health in China." *Journal of International Economics*.

Fan, H, Y Liu, L D Qiu, and X Zhao (2020), "Export to Elude." *Journal of International Economics*.

Feenstra, R C, H Ma, and Y Xu (2019), "US exports and employment." *Journal of International Economics*.

Felbermayr, G, F Teti, and E Yalcin (2019), "Rules of origin and the profitability of trade deflection." *Journal of International Economics*.

Felbermayr, Gabriel J., Jasmin Gröschl, and Inga Heiland (2022), "Complex Europe: Quantifying the Cost of Disintegration." *Journal of International Economics* 138, 103647. URL: https://www.sciencedirect.com/science/article/pii/S0022199622000794.

Fieler, A C and A E Harrison (2023), "Escaping Import Competition in China." *Journal of International Economics*.

Fink, Carsten, Aaditya Mattoo, and Ileana Cristina Neagu (2005), "Assessing the Impact of Communication Costs on International Trade." *Journal of International Economics* 672, 428–445.

Flach, L and F Unger (2022), "Quality and Gravity in International Trade." *Journal of International Economics*.

Fontagné, L, H Guimbard, and G Orefice (2022), "Tariff-based product-level trade elasticities." *Journal of International Economics*.

Fontangé, Lionel, Houssein Guimbard, and Gianluca Orefice (2022), "Tariff-Based Product-Level Trade Elasticities." *Journal of International Economics* 137, 103593.

Forero-Rojas, Alejandro, Hibret Maemir, Ana M. Fernandes, and Aaditya Mattoo (2018), "Developing Countries' Trade and Market Access in the European Union and the United States: Introducing Two New Databases." *mimeo*.

Fuchs, A and N H Klann (2013), "Paying a Visit: The Dalai Lama Effect on International Trade." *Journal of International Economics*.

Fugazza, Marco and Alessandro Nicita (2013), "The direct and relative effects of preferential market access." *Journal of International Economics*.

Garred, J (2018), "The Persistence of Trade Policy in China after {WTO} Accession." *Journal of International Economics*.

Gaulier, Guillaume and Soledad Zignago (2010), *BACI: International Trade Database at the Product-Level. The 1994-2007 Version*. Working Papers 2010-23. CEPII. URL: http://www.cepii.fr/CEPII/en/publications/wp/abstract.asp?NoDoc=2726.

Gnutzmann-Mkrtchyan, A and C Henn (2018), "Peeling Away the Layers: Impacts of Durable Tariff Elimination." *Journal of International Economics*.

Goldberg, Pinelopi K. and Nina Pavcnik (2016), "The Effects of Trade Policy." *Handbook of Commercial Policy*. Vol. 1A. Elsevier B.V. Chap. 3, 161–206.

Greenland, Andrew and John Lopresti (2022), "Trade Policy as an Exogenous Shock: Focusing on the Specifics." *Mimeo*.

Han, J, R Liu, B U Marchand, and J Zhang (2016), "Market Structure, Imperfect Tariff Pass-through, and Household Welfare in Urban China." *Journal of International Economics*.

Han, J, R Liu, and J Zhang (2012), "Globalization and Wage Inequality: Evidence from Urban China." *Journal of International Economics*.

Handley, Kyle and Nuno Limão (2015), "Trade and Investment under Policy Uncertainty: Theory and Firm Evidence." *American Economic Journal: Economic Policy* 74, 89–122.

— (2017), "Policy Uncertainty, Trade and Welfare: Theory and Evidence for China and the United States." *American Economic Review* 1079, 2731–2783. URL: http://pubs.aeaweb.org/doi/10.1257/aer.20141419%20http://www.nber.org/papers/w19376.pdf.

Hanson, G and C Xiang (2011), "Trade Barriers and Trade Flows with Product Heterogeneity: An Application to US Motion Picture Exports." *Journal of International Economics*.

Harding, T and B S Javorcik (2012), "Foreign Direct Investment and Export Upgrading." *Review of Economic Statistics*.

Harrison, Ann and Andres Rodriguez-Clare (2010), "Trade, Foreign Investment, and Industrial Policy for Developing Countries." *Handbook of Development Economics*. Ed. by Dani Rodrik and Mark Rosenzweig. Vol. 5. Elsevier. Chap. Chapter 63, 4039–4214. URL: http://econpapers.repec.org/RePEc:eee:devchp:v:5:y:2010:i:c:p:4039-4214.

Hayakawa, K, J Ishikawa, and N Tarui (2020), "What goes around comes around: Export-enhancing effects of import-tariff reductions." *Journal of International Economics*.

Head, Keith and Thierry Mayer (2014), "Gravity Equations: Workhorse,Toolkit, and Cookbook." *Handbook of International Economics*. Ed. by Gita Gopinath, Elhanan Helpman, and Kenneth Rogoff. Vol. 4. Elsevier B.V. Chap. 3, 131–195.

— (2019), "Brands in motion: How Frictions Shape Multinational Production." *American Economic Review* 1099, 3073–3124.

Heid, B and M Larch (2016), "Gravity with Unemployment." *Journal of International Economics*.

Hillberry, Russell and David Hummels (2013), *Trade elasticity parameters for a computable general equilibrium model*. Vol. 1. Elsevier, 1213–1269. URL: http://dx.doi.org/10.1016/B978-0-444-59568-3.00018-3.

Hoda, Anwarul (2001), *Tariff Negotiations and Renegotiations under the GATT and the WTO*. Second Edi. Cambridge University Press.

Hsu, W T, Y Lu, and G L Wu (2020), "Competition, Markups, and Gains from Trade: A Quantitative Analysis of China between 1995 and 2004." *Journal of International Economics*.

Imai, Kosuke and Teppei Yamamoto (2010), "Causal Inference with Differential Measurement Error: Nonparametric Identification and Sensitivity Analysis." *American Journal of Political Science* 542, 543–560.

Imbs, Jean and Isabelle Mejean (2015), "Elasticity Optimism." *American Economic Journal: Macroeconomics* 73, 43–83.

Jiao, Y, Z Liu, Z Tian, and X Wang (2022), "The Impacts of the US Trade War on Chinese Exporters." *Review of Economic Statistics*.

Johri, A and M M Rahman (2022), "The rise and fall of India's relative investment price: A tale of policy error and reform." *American Economic Journal: Macroeconomics*.

Jordà, Òscar (2005), "Estimation and Inference of Impulse Responses by Local Projections." *American Economic Review* 951, 161–182.

Kee, Hiau Looi, Cristina Neagu, and Alessandro Nicita (2013a), "Is Protectionism on the Rise? Assessing National Trade Policies During the Crisis of 2008." *Review of Economic Statistics*.

— (2013b), "Is Protectionism on the Rise? Assessing National Trade Policies During the Crisis of 2008." *The Review of Economics and Statistics* 951, 342–346.

Kee, Hiau Looi, Alessandro Nicita, and Marcelo Olarreaga (2008), "Import Demand Elasticities and Trade Distortions." *Review of Economics and Statistics* 904, 666–682.

Keller, Wolfgang and Stephen Ross Yeaple (2013), "The Gravity of Knowledge." *American Economic Review* 1034, 1414–1444.

Ketterer, T D, D Bernhofen, and C Milner (2014), "Preferences, Rent Destruction and Multilateral Liberalization: The Building Block Effect of {CUSFTA}." *Journal of International Economics*.

Khanna, G, K Shih, A Weinberger, and M Xu (2023), "Trade Liberalization and Chinese Students in US Higher Education." *Review of Economic Statistics*.

Krishna, P, J P Poole, and M Z Senses (2014), "Wage Effects of Trade Reform with Endogenous Worker Mobility." *Journal of International Economics*.

Krolikowski, P M and A H McCallum (2021), "Goods-Market Frictions and International Trade." *Journal of International Economics*.

Lake, J and D L Millimet (2016), "An Empirical Analysis of Trade-Related Redistribution and the Political Viability of Free Trade." *Journal of International Economics*.

Lashkaripour, A (2021), "The Cost of a Global Tariff War: A Sufficient Statistics Approach." *Journal of International Economics*.

Lashkaripour, Ahmad and Volodymyr Lugovskyy (2023), "Profits, Scale Economies, and the Gains from Trade and Industrial Policy." *American Economic Review* 11310, 2759–2808.

Li, B (2018), "Export expansion, skill acquisition and industry specialization: Evidence from China." *Journal of International Economics*.

Limão, N and P Tovar (2011), "Policy choice: Theory and evidence from commitment via international trade agreements." *Journal of International Economics*.

Lind, Nelson and Natalia Ramondo (2023), "Trade with Correlation." *American Economic Review* 1132, 317–353.

Liu, C and X Ma (2023), "Migration, tariffs, and China's export surge." *Journal of International Economics*.

Liu, Q and L D Qiu (2016), "Intermediate input imports and innovations: Evidence from Chinese firms' patent filings." *Journal of International Economics*.

Liu, Q, L D Qiu, and C Zhan (2019), "Trade liberalization and domestic vertical integration: Evidence from China." *Journal of International Economics*.

Lu, Y, Z Tao, and L Zhu (2017), "Identifying {FDI} Spillovers." *Journal of International Economics*.

Ludema, R D, A M Mayda, and P Mishra (2018), "Information and Legislative Bargaining: The Political Economy of US Tariff Suspensions." *Review of Economic Statistics*.

Ludema, R D, A M Mayda, Z Yu, and M Yu (2021), "The Political Economy of Protection in GVCs: Evidence from Chinese Micro Data." *Journal of International Economics*.

Ludema, R D and Z Yu (2016), "Tariff pass-through, firm heterogeneity and product quality." *Journal of International Economics*.

Ludema, Rodney D. and Anna Maria Mayda (2013), "Do Terms-of-Trade Effects Matter for Trade Agreements? Theory and Evidence from WTO Countries." *Quarterly Journal of Economics* 1284, 1837–1893. URL: http://qje.oxfordjournals.org/cgi/doi/10.1093/qje/qjt016.

Manova, K and Z Yu (2016), "How Firms Export: Processing vs. Ordinary Trade with Financial Frictions." *Journal of International Economics*.

Nicita, Alessandro, Marcelo Olarreaga, and Peri Silva (2018), "Cooperation in WTO's tariff waters?" *Journal of Political Economy* 1263, 1302–1338. URL: https://www.journals.uchicago.edu/doi/10.1086/697085.

Nunn, N and D Trefler (2010), "The structure of tariffs and long-term growth." *American Economic Journal: Macroeconomics*.

Ornelas, E and P Tovar (2022), "Intra-bloc tariffs and preferential margins in trade agreements." *Journal of International Economics*.

Ossa, R. (2016), *Quantitative Models of Commercial Policy*. 1st ed. Vol. 1. Elsevier B.V., 207–259. URL: http://dx.doi.org/10.1016/bs.hescop.2016.04.003.

Parro, Fernando (2013), "Capital-skill complementarity and the skill premium in a quantitative model of trade." *American Economic Journal: Macroeconomics*.

Romalis, John (2007), "NAFTA's and CUSFTA's Impact on International Trade." *The Review of Economics and Statistics* 893, 416–435.

Schennach, Susanne (2022), "Measurement Systems." *Journal of Economic Literature* 604, 1223–1263. URL: https://www.aeaweb.org/articles?id=10.1257/jel.20211355.

Schmidt, J and W Steingress (2022), "No Double Standards: Quantifying the Impact of Standard Harmonization on Trade." *Journal of International Economics*.

Shapiro, Joseph S. (2021), "The Environmental Bias of Trade Policy." *The Quarterly Journal of Economics* 1362, 831–886. URL: https://doi.org/10.1093/qje/qjaa042.

Shu, Pian and Claudia Steinwender (2019), "The impact of trade liberalization on firm productivity and innovation." *Innovation Policy and the Economy* 191, 39–68.

Soderbery, A (2018), "Trade elasticities, heterogeneity, and optimal tariffs." *Journal of International Economics*.

Soderbery, Anson (2018), "Trade elasticities, heterogeneity, and optimal tariffs." *Journal of International Economics* 114, 44–62.

Spearot, A (2016), "Unpacking the long-run effects of tariff shocks: New structural implications from firm heterogeneity models." *American Economic Journal: Microeconomics*.

Tian, Y (2022), "International Trade Liberalization and Domestic Institutional Reform: Effects of {WTO} Accession on Chinese Internal Migration Policy." *Review of Economic Statistics*.

Tombe, T (2015), "The missing food problem: Trade, agriculture, and international productivity differences." *American Economic Journal: Macroeconomics*.

Tong, H and S J Wei (2014), "Does Trade Globalization Induce or Inhibit Corporate Transparency? Unbundling the Growth Potential and Product Market Competition Channels." *Journal of International Economics*.

UNCTAD (2011), "WITS User's Manual (Version 2.01)." January. URL: http://wits.worldbank.org/data/public/WITS%7B%5C_%7DUser%7B%5C_%7DManual.pdf.

Wong, W F (2022), "The round trip effect: Endogenous transport costs and international trade." *American Economic Journal: Applied Economics*.

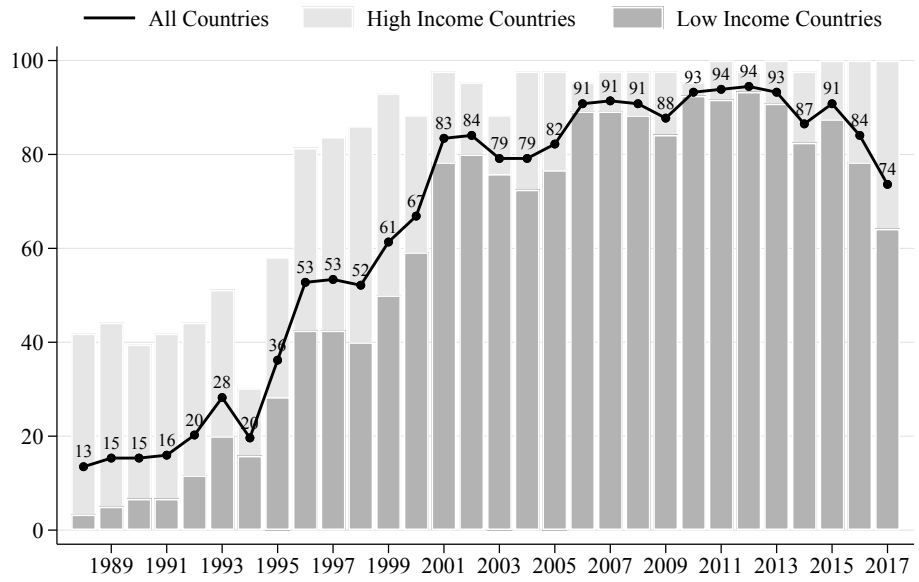
World Bank (2019), *WITS Dataprovider*. URL: https://wits.worldbank.org/dataproviders.html (visited on 08/08/2019).

— (2018), *World Integrated Trade Solution (WITS)*. URL: https://wits.worldbank.org/ (visited on 10/29/2018).

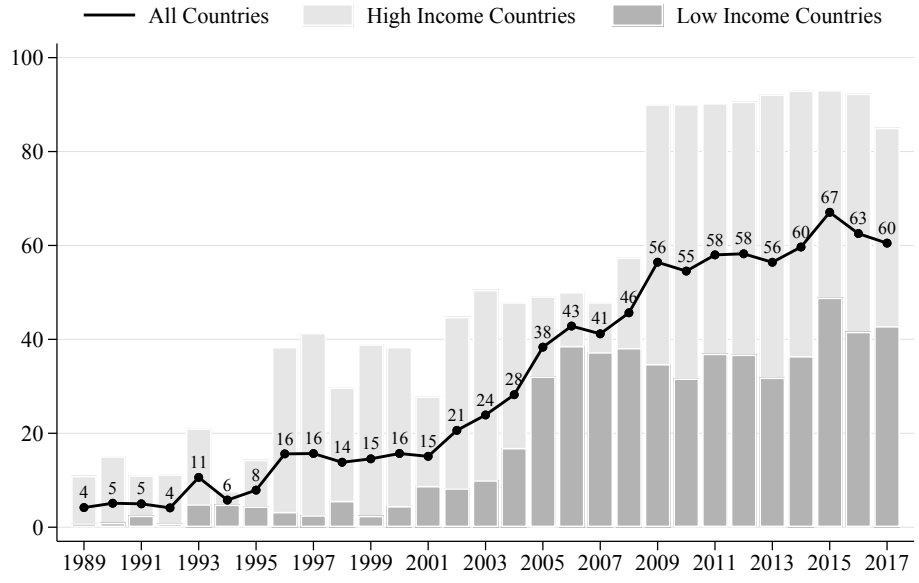
A Appendix: Additional Material

Figure A1: Share of Countries Reporting Tariffs by Year (in %)

(a) MFN Tariffs (Trains and IDB)



(b) Preferential Tariffs (Trains)



Note: The graphs show the share of countries reporting MFN (Panel (a)) and preferential (Panel (b)) tariffs. The solid line gives the overall share, and the bars split the sample by income. For the MFN tariffs, the number of countries reporting tariffs at least once in the period of observation equals 163. For the preferential tariffs, I use information on existing RTAs to determine the number of pairs for which preferential tariffs should be reported.

Table A1: Summary of Differences between the GTD and Other Sources

Source	covered years	deals w/ missings	deals w/ phasing-in	deals w/ AVEs	checks RTAs
CEPII's *MAcMap-HS6*	2001, 04, 07, 10, 13, 16	no	partially	yes	no
Caliendo, Robert C. Feenstra, et al. (2023)	1984–2011	yes	partially	no	yes
GTD	1988–2017	yes	yes	no	yes

Note: The table compares the new tariff database with other databases with comparable country coverage.

A.1 Stylized Facts on Pattern of Spikes

The false imputation in the WITS data leads to non-classical measurement error as it affects only preferential tariffs and, hence, only country pairs with an RTA. Next, I want to understand better whether there are systematic patterns in the missing observations when I use the data provided by WITS. To do so, I identify the observations that lead to the troublesome spikes in the time series and code them with an indicator that I use as a dependent variable to analyze patterns. The spikes $t^{\wedge}_{ijkt}$ are a dummy variable and defined as

$$t^{\wedge}_{ijkt} = \begin{cases} 1, & \text{if } t^{Pref}_{ijkt} = \text{missing } \cap\, RTA_{ij,t} = 1 \cap t^{Trains}_{ijkt} = t^{MFN}_{ikt} \cap t^{MFN}_{ikt} \neq 0 \\ 0, & \text{otherwise.} \end{cases} \tag{12}$$

The first part of the logical expression, $t^{Pref}_{ijkt} = \text{missing} \cap RTA_{ij,t} = 1$, determines whether the preferential tariff of importer i on product k from exporter j at time t is missing although an RTA is in place. As mentioned above, the imputation method used by WITS leads to spikes only when the MFN tariffs are available but preferential tariffs are missing. Hence, missing preferential tariffs are not per se a problem; the issue arises only if WITS imputes the MFN tariff to fill the missing observations, which is captured by the second part of the expression, $t^{Trains}_{ijkt} = t^{MFN}_{ijkt}$. Furthermore, only MFN tariffs that are greater than zero lead to bias $\left(t^{MFN}_{ikt} \neq 0\right)$.

I correlate $t^{\wedge}_{ijkt}$ with various explanatory variables to analyze more systematically for which country pair–product observations the troublesome spikes occur. All importer–exporter–product–year combinations after the year of entry into force of the respective RTA are included if a preferential tariff was reported at least once between 1988 and 2017. To keep the number of observations manageable, I include only observations in 5-year intervals; i.e., the years 1989, 1997, 2002, 2007, 2012, and 2017 are included.

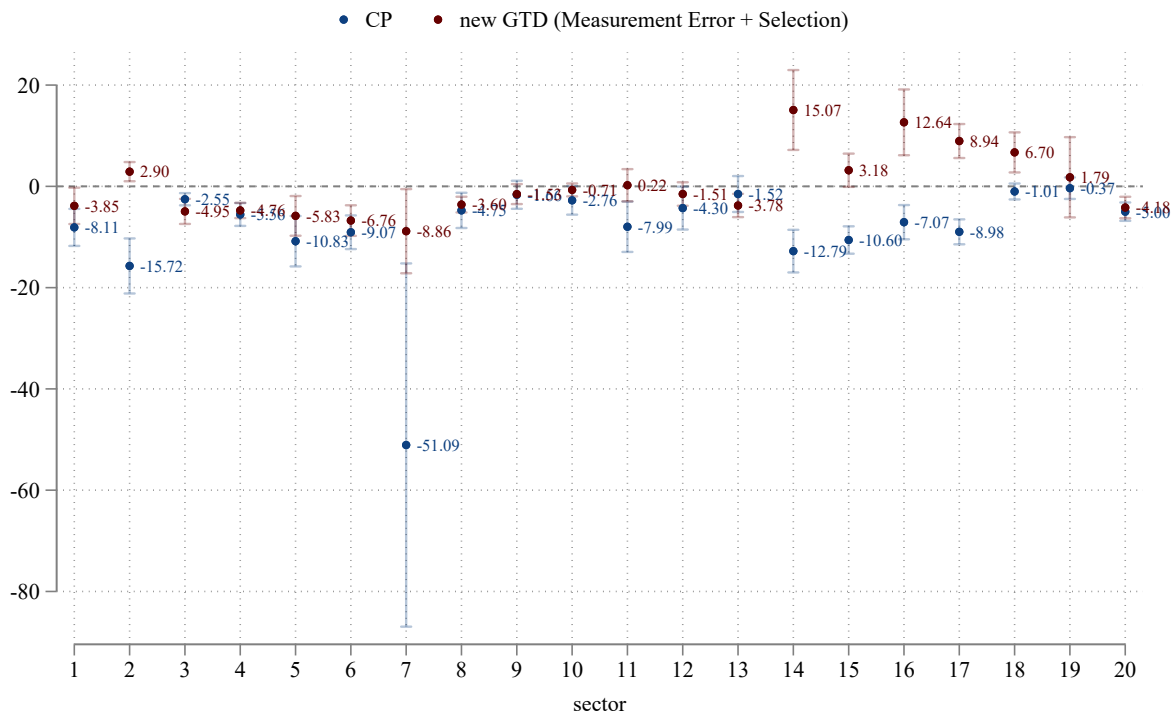
Table A2: Understanding the Measurement Error

	High Income			Low Income		
	(1)	(2)	(3)	(4)	(5)	(6)
Customs Union	0.15***	0.13***		0.04***	0.04***	
	(0.00)	(0.00)		(0.00)	(0.00)	
Deep FTA	-0.01***	-0.01***		-0.03***	-0.03***	
	(0.00)	(0.00)		(0.00)	(0.00)	
Partial Scope TA	0.23***	0.23***		0.07***	0.07***	
	(0.00)	(0.00)		(0.00)	(0.00)	
Non-Reciprocal TA	-0.01***	-0.01***		0.19***	0.20***	
	(0.00)	(0.00)		(0.00)	(0.00)	
zero		-0.11***			-0.07***	
		(0.00)			(0.00)	
Constant (Shallow FTAs)	0.18***	0.28***	0.19***	0.58***	0.64***	0.60***
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Obs. (in Mio.)	118.39	118.39	118.38	41.47	41.47	41.31
R^2	0.02	0.03	0.56	0.01	0.01	0.92
FE			i-hs6-year			i-hs6-year

Note: The table shows the results of the correlation analysis of the pattern of the measurement error. See main text for details. R^2 is adjusted.

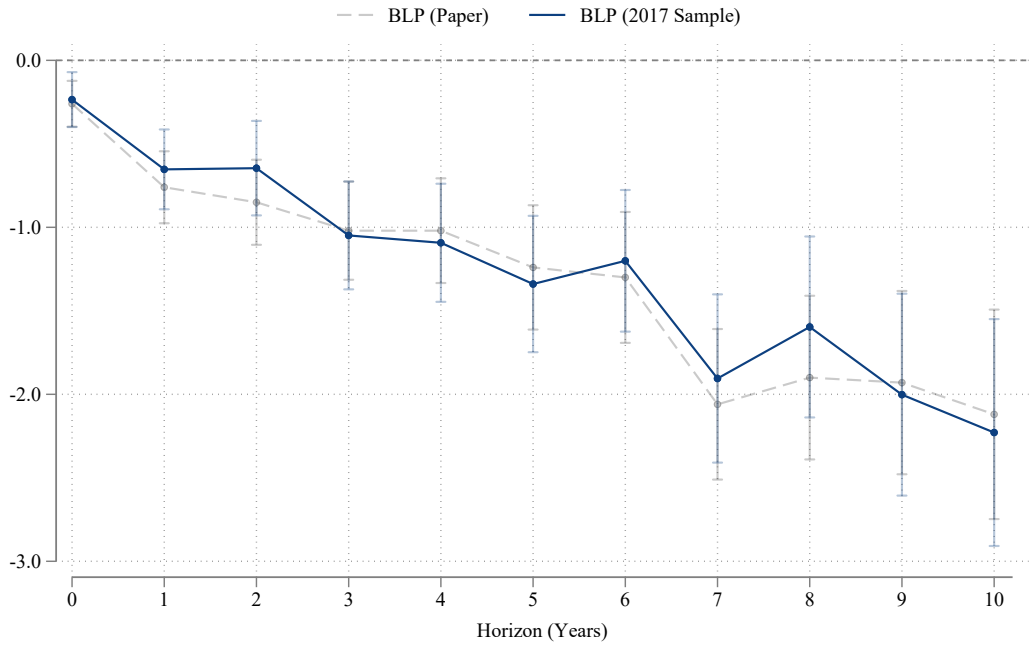
The results are presented in Table A2 and show four interesting facts that inform us on which observations are most affected by the false imputation. First, the probability of misreporting and the resulting spikes correlates negatively with income.Second, misreporting cannot be explained by zero trade flows (cf. columns (2) and (5)). Instead, the negative and significant coefficients for $zero_{ijkt}$, which equals one for zero trade flows and zero otherwise, suggest positive selection—i.e., that countries that trade more have a higher probability of an RTA in the first place and therefore also a higher probability their data series displaying the problematic spikes. Third, the pattern of misreporting is not driven by one particular type of agreement, but instead the troublesome spikes are ubiquitous: although the shallower types of RTAs, such as partial scope agreements and nonreciprocal trade arrangements, have a higher probability of troublesome spikes, missing tariffs are present for every type of RTA.Last, I want to analyze whether countries misreport all their tariff schedules or only specific ones. Put differently, whenever Mexico fails to report its tariff schedule for NAFTA, are all its other preferential tariff schedules for partners with an RTA also missing? To answer this question, I decompose the variance. The dimensions are importer i, exporter j, product k, and year t. If countries either perfectly report in a year, i.e., information on all tariff schedules is available, or have missing observations for all preferential tariffs, all of the variation in the dependent variable can be explained by importer–product–year FE because then misreporting does not vary across partners. Columns (3) and (6) show the results across income groups. Again, rather stark differences across income groups are observable: while the i-k-t FE explain most, but not all, of the variation for low-income countries, this is not the case for high-income countries.

Figure A2: Elasticities by WIOD Sector: Caliendo and Parro (2015)



Note: The graph shows the results of estimating Equation (7) separately for all WIOD sectors. The replicated coefficients from CP are displayed in blue and give the estimates when I use the tariff data that contains both the false interpolation and positive selection ($-\tilde{\sigma}_s$). The results when I use the GTD and hence correct for both issues present in the WITS data are shown in red ($-\sigma$).

Figure A3: Replication of Boehm et al. (2023) with Smaller Sample



Note: The gray line shows the coefficients that result when I replicate BLP's analysis using their provided data and code, which are identical to the results reported in the paper. The dark blue line reports the coefficients when I rerun the same analysis on the same data but trim the sample by dropping the year 2018 and keeping only HS6 products in the 1988/92 nomenclature. The differences are minimal.

Table A3: First Stage in Boehm et al. (2023)

	$h=1$				$h=10$			
z	0.82***	0.72***	0.71***	0.67***	0.71***	0.62***	0.50***	0.56***
	(0.005)	(0.003)	(0.002)	(0.005)	(0.009)	(0.007)	(0.003)	(0.06)
	BLP	GTD	LoM	HICs	BLP	GTD	LoM	HICs

Note: The table shows the how the first stage changes when I correct for false interpolation and sample selection.

Table A4: $Cov(\Delta lnX_{ijkt}, u_{ijkt})$ Boehm et al. (2023)

	$h=1$	$h=10$
u	-0.31***	-0.79***
	(0.010)	(0.216)

Note: The table regresses the changes in trade ΔlnX_{ijkt} on the measurement error u, pretrend controls, and importer–HS4–year, exporter–HS4–year, and importer–exporter–HS4 fixed effects. Standard errors are clustered at the country–pair–product level.

B Appendix: Data

B.1 Existing Papers Using WITS Tariff Data

To maximize replicability and objectivity, I first built a corpus of full texts of recent papers that might have used the WITS data. In November 2023, I used Google Scholar to search for "tariffs" among papers published in leading journals (*American Economic Review*, *Quarterly Journal of Economics*, *Journal of Political Economy*, *Econometrica*, and *Review of Economic Studies*) since 2010. This query gives more than 600 distinct papers. In a first attempt, I followed Disdier and Head (2008) and used EconLit to generate the corpus of full texts that contained "tariffs". Unfortunately, for reasons beyond my understanding, this approach missed many papers, which is why I ended up choosing Google Scholar instead. The list of papers still contained many false positives, i.e., papers that do not use the WITS data such as pure theory papers. Therefore, in a second step, I searched the full text for at least one the following words: "WITS", "world integrated trade solution", "W.I.T.S.", "UNCTAD-Trains", "Trains", "T.R.A.I.N.S.", "Trade Analysis and Information System", and "UNCTAD". Last, I checked all the remaining papers one by one to eliminate any remaining false positives (for example, Shapiro (2021) uses WITS to extract information on nontariff barriers).

Published Papers using WITS Data
Erten and Leight (2021)
Tian (2022)
Amiti and Khandelwal (2013)
Davies et al. (2018)
Harding and Javorcik (2012)
Kee, Neagu, et al. (2013a)
Estevadeordal and Taylor (2013)
Carballo et al. (2022)
Erten and Keskin (2021)
Bernard et al. (2015)
Coelli et al. (2022)
Khanna et al. (2023)
R D Ludema, A M Mayda, and Mishra (2018)
Jiao et al. (2022)
Han, R. Liu, Marchand, et al. (2016)
Esposito (2022)
R D Ludema, A M Mayda, Yu, et al. (2021)

Continued on next page

Table B1 – continued from previous page

Published Papers using WITS Data
Dasgupta and Mondria (2018)
Fugazza and Nicita (2013)
Ornelas and Tovar (2022)
Johri and Rahman (2022)
Duval et al. (2016)
Tombe (2015)
Dix-Carneiro, Soares, et al. (2018)
Han, R. Liu, and Zhang (2012)
G. Felbermayr et al. (2019)
Gnutzmann-Mkrtchyan and Henn (2018)
Bombardini, Orefice, et al. (2019)
C. Chen and Steinwender (2021)
Krishna et al. (2014)
Facchini et al. (2019)
Dix-Carneiro and B K Kovak (2019)
Wong (2022)
J E Anderson and Y V Yotov (2020)
Bergstrand et al. (2023)
Baccini et al. (2019)
Hayakawa et al. (2020)
Alquist et al. (2019)
Ardelean and Lugovskyy (2010)
Carluccio and Bas (2015)
Aghion et al. (2015)
Ahn et al. (2011)
G. J. Felbermayr et al. (2022)
Ketterer et al. (2014)
Bombardini and B. Li (2020)
Parro (2013)
N. Chen and Juvenal (2022)
Artuc, Bastos, et al. (2023)
Fan, Y. A. Li, et al. (2020)
Bas et al. (2017)
Davidson et al. (2014)

Continued on next page

Table B1 – continued from previous page

Published Papers using WITS Data
R D Ludema and Yu (2016)
Q. Liu and Qiu (2016)
Beshkar and Lee (2022)
Fieler and A. E. Harrison (2023)
Beverelli et al. (2017)
Artuc, Porto, et al. (2019)
Nunn and Trefler (2010)
Lake and Millimet (2016)
Brandt et al. (2021)
Bastos et al. (2013)
Fan, Y. Liu, et al. (2020)
Defever et al. (2020)
Arkolakis, Ganapati, et al. (2021)
Barattieri et al. (2021)
Flach and Unger (2022)
Burstein, Morales, et al. (2019)
Beshkar and Bond (2017)
Lashkaripour (2021)
Q. Liu, Qiu, and Zhan (2019)
Garred (2018)
Diez (2014)
Aichele and Heiland (2018)
Hsu et al. (2020)
Chakraborty and Raveh (2018)
Tong and Wei (2014)
Spearot (2016)
Fuchs and Klann (2013)
B. Chen et al. (2017)
Lake and Millimet (2016)
Limão and Tovar (2011)
Soderbery (2018)
Baldwin and Seghezza (2010)
Amiti and Cameron (2012)
Maria Bas and Strauss-Kahn (2015)

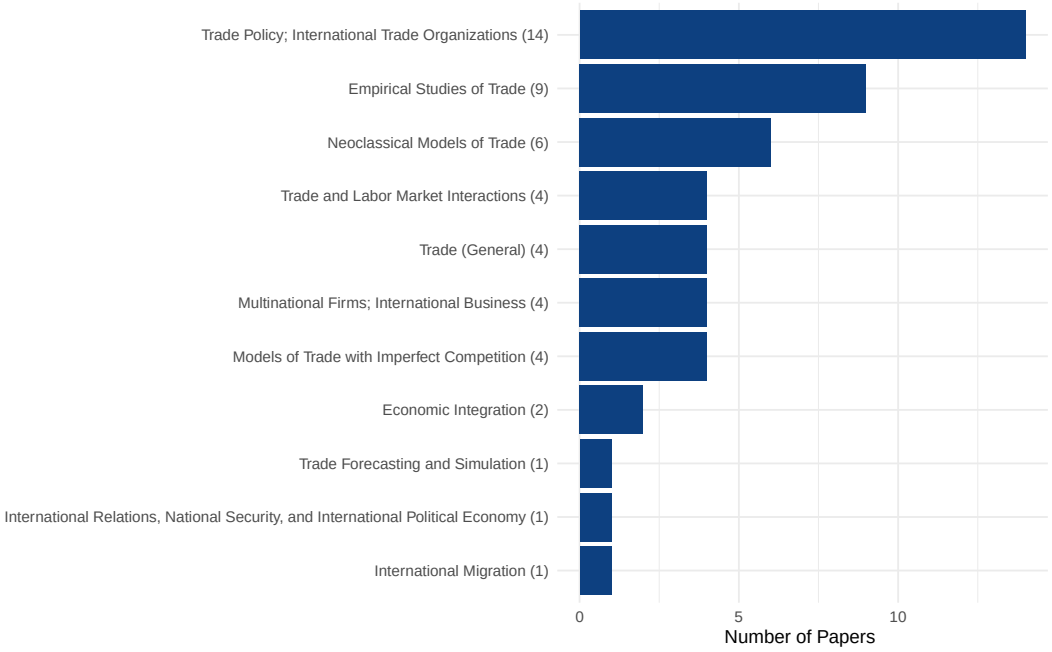
Continued on next page

Table B1 – continued from previous page

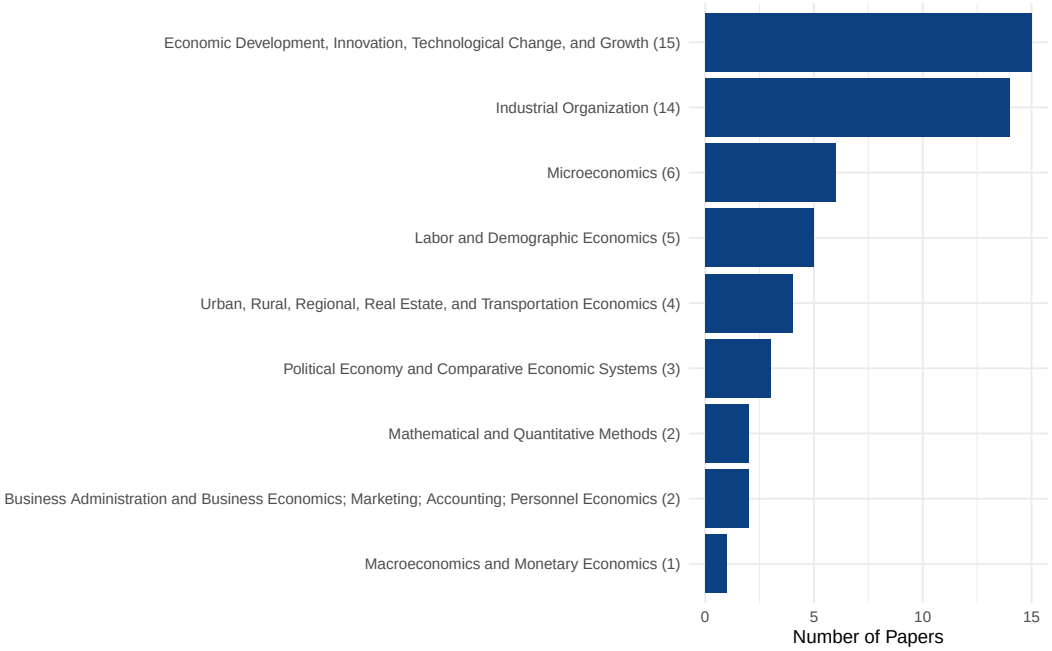
Published Papers using WITS Data
Schmidt and Steingress (2022)
B. Li (2018)
M. X. Chen and Moore (2010)
Arezki et al. (2017)
R C Feenstra et al. (2019)
Baier, Yoto V. Yotov, et al. (2019)
Buono and Lalanne (2012)
Krolikowski and McCallum (2021)
Fan, Lin, et al. (2020)
C. Liu and Ma (2023)
Carrere and Grujovic (2020)
Lu et al. (2017)
Heid and Larch (2016)
Manova and Yu (2016)
Hanson and Xiang (2011)
Breinlich et al. (2018)
Fontagné et al. (2022)
Coşar et al. (2018)

Figure B1: Analysis of JEL Codes of Top 5

(a) Within International Economics



(b) Other Fields



Note: The graphs show the distribution of JEL codes of the 21 papers that use WITS data; the total number of papers with the respective JEL code is in parentheses.