Missing Tariffs

FEODORA A. TETI*

November 20, 2023

Abstract

Many studies use tariffs to measure changes in trade policy. This paper shows that standard sources for tariffs suffer from substantial measurement error due to misreporting and the resulting false imputation: Countries fail to report tariffs every year and missing data are more prevalent for preferential than for most favored nation (MFN) tariffs. WITS, the main data provider for tariffs, falsely interpolates missing preferential tariffs with MFN tariffs. This practice leads to artificial spikes in bilateral time series data and, hence, induces massive measurement error. I introduce a new global tariff dataset at the six-digit product level for 197 countries and 30 years that combines five different sources for tariffs and proposes a new interpolation algorithm taking the misreporting into account. Lastly, I show using gravity that correcting for the messy data increases the estimates of the trade elasticity by 2.89 times.

Keywords: Tariffs, MFN, Preferences, Trade Elasticity

JEL-Classification: F13, F14

*ifo Institute – Leibniz Institute for Economic Research at the University of Munich, Poschinger Str. 5, 81679 Munich, Germany; Phone: +49 89 9224 1389; E-mail: teti@ifo.de

I am very thankful to my advisors, Gabriel Felbermayr, Andrew Bernard, and Monika Schnitzer for their invaluable guidance and support. I would also like to thank Andrea Ariù, Martin Braml, Volker Nitsch, and 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 (project number 280092119), and the scholarship from the Bernt Rohrer Foundation. All errors are my own.

1 Introduction

Many studies in economics, both within the field of international economics but also outside, such as research tackling questions on innovation, labor, inequality, or development, rely on tariffs to measure changes in trade policy (c.f., Goldberg and Pavcnik 2016; Harrison and Rodriguez-Clare 2010; Shu and Steinwender 2019). In international trade, tariffs are often used for trade costs because they vary over time and sectors, and—in contrast to most non-tariff barriers—they are quantifiable and easily to compare across countries. Tariffs have been leveraged to identify the trade elasticity, that 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 for many other research questions in trade such as work that focuses on quality (Amiti and Khandelwal 2013), the environment (Shapiro 2021), labor markets (Greenland and Lopresti 2022), and non-tariff barriers (De Melo and Nicita 2018; Kee et al. 2013; Kuenzel 2022).

Despite their central role, data on tariffs are hard to find. As Anderson and Van Wincoop (2004) state *"the grossly incomplete and inaccurate information on policy barriers available to researchers is a scandal and a puzzle"* (p. 693); almost twenty years later the situation is not substantially better. Why is collecting data on tariffs so challenging? First, the sheer number of observations—197 importer, 196 exporter, and roughly 5,0000 products amounts to 200 Million observations per year—makes tariff data very hard to handle, confusing, and therefore error prone. Second, the large scale trade integration through regional trade agreements (RTAs) and the resulting numerous different tariff schemes increase complexity. Lastly, the data collecting institutions rely on countries to report their tariffs but have very little tools to enforce reporting compliance (Hoda 2001).

This paper shows that standard sources for tariffs suffer from substantial measurement error due to misreporting and the resulting false imputation: Countries fail to report tariffs every year and missing data are more prevalent for preferential than for most favored nation (MFN) tariffs. The World Bank's World Integrated Trade Solution (WITS), the main data provider for tariffs, falsely interpolates missing preferential tariffs with MFN tariffs. This practice leads to artificial spikes in bilateral time series data and, hence, induces massive measurement error. I introduce a new global tariff dataset at the six-digit product level for 197 countries and 30 years that combines five different sources for tariffs and proposes a new interpolation algorithm taking the misreporting into account. With the new data at hand and using gravity, I quantify the bias in estimates of the trade elasticites due to the missing and falsely imputed data.

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

In Section 2, I start by describing the problems of existing tariff databases and discuss for which countries the resulting measurement error is particularly pronounced. To understand the roots of the measurement error in the available tariff data, it is helpful to review the different types of tariffs: (applied) most favored nation tariffs (MFN tariffs) and preferential tariffs. In line with the non-discriminatory principle of the WTO, all WTO-members impose for each tariff line the same tariff against all other WTO-members—the (applied) MFN tariff.[2] RTAs are the only exemption to the WTO's central concept of the non-discrimination rule as they grant the participating countries preferential access. In most cases the exporting firms do not have to pay any tariffs if they can prove that their product originates in one of the contracting parties. For example, the United-States-Mexico-Canada Agreement (USCMA, formerly NAFTA) allows Mexican truck-exporters to enter the U.S.-market for free while German exporters have to pay a tariff of 25%.[3]

The goal of WITS is to provide the public with easily attainable information of the tariff that is imposed for each country pair. WITS is a platform to facilitate the download of tariff data made available by the UNCTAD and the WTO.[4] To determine the applied tariff t_{ijtk} of importer i on product k from exporting country j at time t, WITS assumes that it equals the preferential tariff t^*_{ijkt},whenever an RTA is in force, and the MFN tariff t_{itk} otherwise, hence, $t_{ijkt} = min\left\{t_{ikt}; t^*_{ijkt}\right\}$. Whenever one downloads the so-called effectively applied tariff t_{ijkt} through WITS, this minimum-operation is performed in the background. For example, for aggregated data, WITS will first determine the effectively applied tariff taking the minimum of the MFN and preferential tariff and then aggregate to the requested level.

If information on tariffs were complete, the minimum-procedure would not have any consequences. Unfortunately, in many cases countries fail to report, and missing tariffs are more prevalent for preferential tariffs. Therefore, the interpolation technique used by WITS with the goal to find out the effectively applied tariffs leads to artificial spikes in time-series data: Whenever an RTA is in place and a country fails to report a preferential tariff scheme, the effectively applied tariff jumps up to the the level of the MFN tariff. Hence, the artificial changes in tariffs are not due to changes in trade policy and should therefore also not be used for any empirical analysis.

---

[2] Applied MFN tariffs should not be confused with bound tariffs. The latter are specific commitments made by individual WTO members and are the maximum MFN tariff level for a given product k. Bound tariffs are typically not the tariffs that are applied in practice, instead most countries impose a lower level, the applied MFN tariff, therefore, the focus of this paper lies on the applied MFN tariffs.

[3] The reported tariffs refer to the product code 870421: Motor vehicles for the transport of goods: Other, with compression ignition internal combustion piston engine (diesel or semidiesel): Of a gross vehicle weight not exceeding 5 tonnes.

[4] The UNCTAD provides information through the Trade Analysis Information System (TRAINS) and the WTO, through the Integrated Data Base (IDB) and the Consolidated Tariff Schedules (CTS).

Analyzing the pattern of spikes, I can show that the probability of falsely imputed missing tariffs correlates negatively with income, the spikes cannot be explained by zero trade, the spikes are not particular to one specific type of trade agreement (FTA, customs union, partial scope agreement or nonreciprocal trade arrangement), and lastly, instead of not reporting all tariff schedules, countries failure to report is partner-specific.

To address the unsatisfactory data situation with respect to tariffs, in Section 3, I introduce a new global tariff dataset (new GTD) that accounts explicitly for the missing tariffs. I proceed in three steps: First, I minimize the number of missing observations by combining five different sources on tariffs. Besides including information provided by UNCTAD, the WTO, the ITC, and national authorities, I have also hand-collected the phase-in schedules for 149 FTAs that provide flawless information for the respective FTAs. In the second step, I propose a new algorithm to fill the observations that are still missing. For MFN tariffs, I use the first non-missing preceeding tariff. If no data from earlier years are available, I use the first succeeding observation. For preferential tariffs, the imputation is performed identically, but conditional on the RTA being in force. Furthermore, I account for phased-out tariffs using information on the phase-out regime. The DESTA database (Dür et al. 2014) has information on more than 600 RTAs on the respective phase-out regime. This information is combined with the panel of tariffs to interpolate preferential tariffs as closely as possible to reality. In the last step, I retrieve the effectively applied tariff using the fully filled MFN and preferential tariffs.

With the new tariff data at hand, I want to analyze how much the missing tariffs and the measurement error due to the false imputation biases the estimates of the trade elasticity. To do so, I run regressions guided by gravity, to identify how trade flows change when tariffs increase. The resulting coefficient can be interpreted as the trade elasticity. First, I use information on tariffs provided by WITS, then I compare the estimates with the results when using the new GTD. I estimate 32 distinct specification that differ with respect to the level of aggregation (pair or pair-sector level), the estimator (OLS or PPML), the control variables, and the type of variation used (cross-sectional or over time). The results show a very strong attenuation bias, independently of the exact specification: the median of the 32 σ-coefficients that result when using the data available through WITS is by a factor of 2.86 smaller than when the new GTD is used for estimation. Hence, the bias caused by the measurement error due to false imputation leads to a substantial underestimation of the trade elasticity.

My paper makes three important contributions: first, it is the first to systematically uncover the mystery of messy tariff data. Second, it introduces a new global tariff database that can be used in many applications, also outside of the trade literature. Lastly, the paper shows that not addressing the missing tariffs and the resulting false imputation, leads to substantial attenuation bias in one of the most important parameters in international economics, i.e., the trade elasticity.

I relate to to the extensive literature on determining the trade elasticity using gravity 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 of its magnitude. Estimates vary between -10 and -1 and depend heavily on the used method as well as the exact setting and sample, respectively (**Boehm2022**; Caliendo and Parro 2015; Fajgelbaum et al. 2020; Fontangé et al. 2022; Imbs and Mejean 2015; Romalis 2007). To the best of my knowledge, nobody has discussed the issue of messy tariff data in this context.

There have been previous efforts in putting together new global tariff datasets (Bouët et al. 2008; Bown and Crowley 2016; Caliendo, Feenstra, et al. 2015). However, the new GTD is unique in terms of country and time coverage. I also cross-validate all tariffs with external sources, which is another new feature of the data. In addition, the level of detail that is used to fill the missing observations, especially for preferential tariffs, i.e., gathering information on detailed tariff schedules and accounting explicitly for phased-out tariffs, is unique. Lastly, the papers also adds to the literature on how to deal with missing data, for example Allison (2001) and Enders (2022).

The rest of this paper is structured as follows. Section 2 discusses how missing data and the false imputation by standard sources for tariffs lead to severe measurement error in tariffs. Second 3 introduces the new GTD and carefully explains all steps taken to address the issue of missing tariffs. Lastly, Section 4 presents the revisited estimates of the trade elasticity when accounting adequately for mismeasured tariffs and Section 5 concludes.

2 Missing Data and False Imputation in Standard Sources for Tariffs

In this section, I will discuss how misreporting and false imputation in the standard sources of tariff data lead to systematic measurement error in tariffs, one of the most important variable in the international trade literature.

2.1 Status Quo: What are the Problems with Existing Tariff Data?

When conducting any type of analysis, one is interested in the tariff that each firm has to pay when exporting to all trading partners, i.e., the preferential tariff t^*_{ijkt} that importer i imposes on exporter j for product k at time t whenever an RTA is in place and the MFN tariff t_{ikt} otherwise. In theory, the World Bank's World Integrated Trade Solution (WITS) provides exactly these data. WITS is the key source for global panel tariff data that is publicly available.[5] It is a download tool and pools data from the United Nations Conference on Trade and Development Trade Analysis Information System (TRAINS) and the WTO, namely the Integrated Data Base (IDB) and Consolidated Tariff Schedules (CTS). Since 2010 most of the raw data used in TRAINS come from the International Trade Center (ITC).[6] The data include information for almost 200 countries on the 6-digit product level of the common HS system with some of the data dating back to 1988.[7] Information about preferential and MFN tariffs are derived from both TRAINS/ITC as well as the IDB, while the CTS is the only provider of data concerning bound tariffs.

To determine the tariff that countries impose on each other, WITS uses the concept of "effectively applied tariffs" which is defined as the lowest available tariff, i.e., if a preferential tariff exists, it will be used as the effectively applied tariff, otherwise, it equals the applied MFN tariff (UNCTAD 2011, p. 95). Hence, WITS assumes for every country that the effectively applied tariff t_{ijkt} equals the minimum of the applied MFN and the preferential tariff, $t_{ijkt} = min\left\{t_{ikt}; t^*_{ijkt}\right\}$. The same logic applies when users download aggregated data, for example when interested in sectoral tariffs: first, WITS determines the effectively applied tariff t_{ijkt}, then it will aggregate tariffs to the requested level.

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

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

[7] For a few countries tariffs are even available at the tariff line. These can be accessed when downloading the data for single countries instead of using the bulk download option.

The effectively applied tariffs t_{ijkt} that WITS reports, hinge on a crucial assumption, namely complete information: the reported effectively applied tariff is only correct when the tariffs for both, MFN and preferential tariffs, are available and complete. Whenever countries only report MFN tariffs, although an RTA is in place, the procedure implemented by WITS to determine the effectively applied tariff leads to mismeasured tariffs: Instead of reporting a missing preferential tariff, the imputing method leads to a jump in the tariff to the level of the MFN tariff that is higher than the preferential tariff. Therefore, the resulting spikes in the time series of tariffs are not due to changes in trade policy but instead are an artifact created by wrong interpolation techniques.

Figure 1 graphically illustrates how missing tariffs and the and false imputation by WITS lead to substantial measurement error. It shows the tariff relationship for selected pairs that sign an RTA during the period of observation. The red dots are the simple average of the preferential tariffs over all products k for a given year, the blue dots give the simple average of the MFN tariffs, respectively. The figure shows averages of more than 5,0000 HS6-products. The solid line is the effectively applied tariff as downloaded from WITS and is defined as the minimum between the observed MFN and preferential tariff for a given country-pair. Panel (a) displays Mexican tariffs on trade with the United States. Until NAFTA came into effect on January 1 1994, Mexico imposed the MFN tariff on U.S. imports, afterwards the preferential tariffs negotiated in NAFTA were in place.[8] As Figure 1 shows, Mexico does not report MFN and preferential tariffs for all years: from 1995 onwards, the MFN tariffs are available for every year, before, information is only available for 1991. The availability of the preferential tariffs is even worse as the few red dots illustrate. The practice of defining the effectively applied tariff as the minimum of the MFN and preferential tariff without accounting for the spotty reporting pattern leads to massive variation over time in the effectively applied tariff with spikes whenever the preferential tariffs are missing and the MFN tariffs are available—most certainly not the type of variation anyone wants to use in empirical analyses.

The problem of false imputation due to missing tariffs is not specific to Mexico but instead also prevalent for other important RTAs around the world. Trade within Mercosur, the largest RTA in Latin-America, is free for almost all products. After a four year long phase-in period, the customs union entered into force in 1995. Again, the misreporting leads to measurement error (c.f., Figure 1, Panel (b)). Similar issues can also be found for the tariffs imposed by Chile on Japanese exports illustrated in Panel (c). The eastern enlargement of the EU led to thirteen new member states with the biggest accession-wave in 2004 when ten countries joined.

[8] As Besedes et al. (2020) show, Mexico agreed to cut 51% of all tariff lines vis-à-vis Canada and the United States immediately, the remaining tariff lines were phased-out following different schedules, the maximum phase-outs being 15 years long.

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

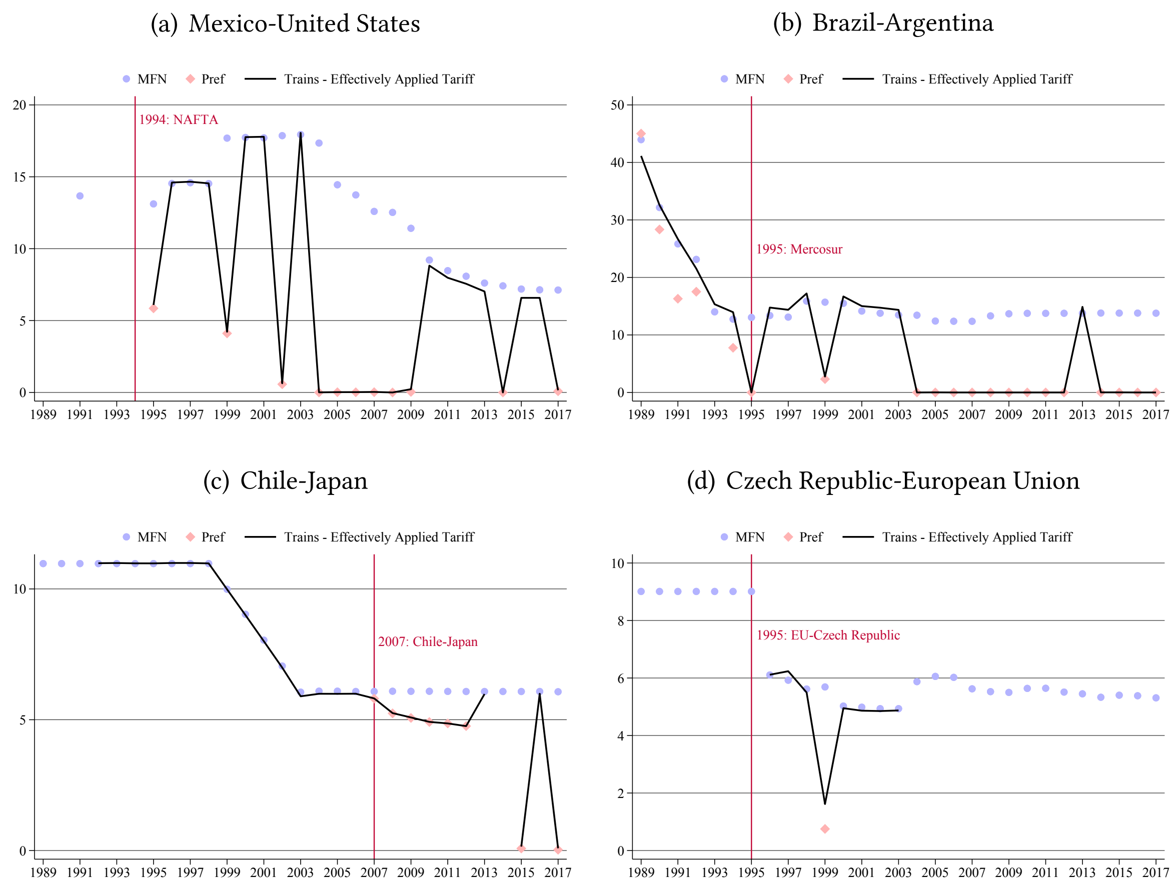

Note: The Figure shows the simple average over all HS6 products k of the preferential, MFN, and effectively applied tariff reported by WITS for selected country pairs. The spikes indicate missing preferential tariffs while the respective country reports MFN tariffs. The entry into force years are from the database on Economic Integration Agreements put together by Scott Baier and Jeffrey Bergstrand (Baier, Bergstrand, and Feng 2014).

However, most of the tariff cuts were already realized long before the accession because the new member states had 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 exports to the Czech Republic were mostly eliminated due to this agreement, a fact that is hard to recognize using data from WITS as they only exist for 1999 (Panel (d)).

Can the observed spikes be rationalized by any events related to trade policy? I will examine two potential explanations that are prime suspects because they both lead to temporary tariff hikes: the suspension of trade agreements and anti-dumping tariffs. Especially in the case of unilateral trade arrangements, sometimes the granted preferences are temporary revoked.[9]

[9] Albornoz-Crespo et al. (2021) provide an example of the United States suspending the General System of Preferences for Argentine due to a dispute over intellectual property.

In these cases, the effectively applied tariff would rightfully jump up to the level of the applied MFN. However, suspensions are a rare event and can therefore not explain the observed spikes.[10] Furthermore, if suspensions were the sole explanation of the spikes, countries should always report preferential tariffs 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 (c.f., Figure A1 (b)). Lastly, for the four examples discussed in Figure 1, I could not find any confirming evidence of the respective agreements being suspended. Anti-dumping tariffs are targeting just a handful of products. Hence, it is highly unlikely that the few anti-dumping tariffs let the average effectively applied tariff of more than 5,0000 products shoot up—let alone to the same level of applied MFN tariffs.[11] The discussion shows that explanations other than false imputation fail to fully explain the observed pattern of spikes.

2.2 Stylized Facts on the Pattern of Spikes

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

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

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 will only lead to spikes when the MFN tariffs are available but preferential tariffs are missing. Hence, missing preferential tariffs are not per se a problem but only if WITS has imputed them with the MFN

[10] When crosschecking for all 241 RTAs that have been in force in 2019 using Mario Larch's RTA database (Egger and Larch 2008) only for 4.72% of country pairs a suspension could be found. A suspension is defined as the dummy variable indicating the presence of an RTA switching more than once within a country pair between the entry into force year and 2019. This definition is overstating the presence of suspensions because it also includes RTAs that were only in place for a subset of time. Therefore, the above reported share is most likely an upper bound of the true number of suspensions, i.e., the real share is even smaller.

[11] When looking at product-level tariff data, it becomes apparent the discrepancies between applied MFN and effectively applied tariff downloaded from WITS vanish.

tariff, which is captured by the second part of the expression, $t_{ijkt}^{Trains} = t_{ijkt}^{MFN}$. Furthermore, only MFN tariffs that are greater than zero lead to bias $\left(t_{ikt}^{MFN} \neq 0\right)$.

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

Table 1: 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 1 and show four interesting facts that inform us on which observations are mostly 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 (c.f., columns (2) and (5)). Instead, the negative and significant coefficients for $zero_{ijkt}$ that equals one for zero trade flows and zero otherwise, suggests positive selection, i.e., countries that trade more, have a higher probability of an RTA in the first place and therefore also a higher probability for 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 more shallow RTA-types, such as partial scope agreements and non-reciprocal trade arrangements, have a higher probability of troublesome spikes, missing tariffs are present for every type of RTA.[12] Lastly, I want to analyze, whether countries misre-

[12] Details on the different types of RTAs can be found in the Appendix.

port all tariff schedules or only specific ones. Put differently, whenever Mexico does not report the tariff schedule for NAFTA, are also all other preferential tariff schedules for partners with an RTA 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 missings for all preferential tariffs, all of the variation in the dependent variable can be explained by an importer-product-year FE because then, misreporting does not vary across partner. Column (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.

3 Introducing the New Global Tariff Database

In this section, I will present how the new global tariff database (GTD) deals with missing data to alleviate the problems associated with them. The goal is to have 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 One*, I combine the data available through WITS with other databases containing information on tariffs to minimize the number of missing observations and I end up using five different primary sources for tariffs. In *Step Two*, 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 one and two are separately conducted for MFN and preferential tariffs, as the goal is to fill the missing observations for both types of tariffs. In *Step Three*, I combine MFN and preferential tariffs that now do not have any missing observations anymore, and use them to retrieve the effectively applied tariffs t_{ijt}.

3.1 *Step One:* Combining Different Sources to Minimize 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, *3)* the IDB-database (both accessed via WITS), and *4)* the International Trade Center (ITC); *5)* U.S. and EU tariff schedules from national authorities kindly provided by researchers from the World Bank (Forero-Rojas et al. 2018). Table 2 summarizes the coverage of the respective sources and shows the available years, number of importing countries, and the type as well as form of tariffs they contain. The number of the rows gives the order of replacing 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 phased-out. For NAFTA, 51% of all tariff lines were cut immediately, for the remaining tariff lines 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 lead to understated tariffs, i.e., the imputed tariffs are too low. Hence, dealing with missing observations for preferential tariffs is one major challenge and one goal of the new GTD is to address the issue.

Table 2: 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 & U.S.

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

For 149 FTAs, I have gathered information on the exact phase-out schedules for all tariff lines and participating countries; when they are available, no imputation is necessary in the first place. They can be found through the WTO's RTA database and cover especially more recent FTAs.[13] The phase-out schedules are the digitized version of the long appendices every FTA has 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, in a first step, the layout of the hundreds of files had to be manually standardized. Then, I have aggregated the information to HS6-products by taking the average of the corresponding tariff lines.[14] The big 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, they serve as the primary source of information—only when they are not available I use data provided by the other sources.

While the WTO's phase-out schedules only cover a subsample of all FTAs currently in force, Trains gives all other preferential tariffs that are offered through any type of RTA (FTAs, customs unions, partial scope agreements, and nonreciprocal arrangements are included), too. 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. Unlike the IDB, Trains separately shows preferential and MFN tariffs, which is necessary when the goal is to correct for the troublesome spikes; Instead of giving preferential tariffs, the IDB only has the effectively applied tariffs that include the spikes, making it im-

[13] 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 Article XXIV and V, respectively. As a result so-called "factual presentations" have to be distributed among WTO members. One part of the factual presentation is the tariff schedule that includes all phasing-in schemes (cf. https://www.wto.org/english/tratop_e/region_e/trans_mecha_e.htm for details). The data are available at https://rtais.wto.org/.

[14] Only ad valorem tariffs are included.

possible to distinguish between imputed MFN and the reported preferential tariff. Therefore, I only use the information on MFN tariffs from the IDB and discard the effectively applied tariff altogether. Trains is therefore for most pairs the core of the new GTD in terms of primary sources. Tariffs collected by Trains as well as the IDB are available through WITS.[15]

The ITC's Market Access Map (MAcMap) is another established source for tariff data.[16] It incorporates bound, applied MFN and preferential tariffs from 2007 onwards for 197 countries on the tariff line. The ITC's MAcMap provides raw data, thus, unless countries report perfectly, similar problems as in WITS can be expected, i.e. missing observations, especially for preferential tariffs. Also the ITC offers a built-in imputation method that is supposed to yield the effectively applied tariff, which is, again, the minimum of the MFN and preferential tariff. So, also for the ITC data, the troublesome spikes bias tariff data whenever a country reports the MFN tariff but not the preferential. 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 the ones one can download at the ITC's 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.

Lastly, I use data provided by national authorities of the United States and the EU. For the United States, the U.S. International Trade Commission gathers and publishes tariffs, and the European Commission 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] Similar to the phase-out schedules, the original data give tariffs at the most disaggregated level of disaggregation, i.e., the national tariff line, and have to be aggregated to HS6-products to be comparable across countries and matched to tariffs from WITS.

For preferential tariffs, the phase-out schedules from the WTO's RTA database serve as the primary source of information, missing preferential tariffs are sequentially filled with information provided by Trains, the ITC and, lastly, for the United States and the EU, with data from national sources. For MFN tariffs, Trains serves as the primary source, missing data is sequentially filled-up with information provided by IDB, ITC, and the national authorities.

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

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

[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.

Another complication arises due to non ad valorem tariffs. Regardless of the type of tariff—bound, MFN or preferential—a tariff can take two forms, i.e., ad valorem tariffs (for example, 8%) and non ad valorem tariffs (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, c.f., 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] Instead, neither for the IDB, the phase-out schedules nor the data provided by national authorities AVEs are available.[20]

For most empirical analysis 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 new GTD flags all AVEs based on information provided by Trains and the ITC. For 2017, the available data suggests for less than 2% of pair-products to be non-ad valorem tariffs. However, this is most likely understating the true number because Trains/ITC only record the non-ad valorem tariffs, for which Trains/ITC calculate AVEs. Whenever countries do the calculation of AVEs themselves and end up only reporting ad valorem tariffs, it is impossible to trace back for which product the reported tariff is a true ad valorem tariff and when a non ad valorem tariff. The EU is a prominent example for 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 the extreme values of tariffs, i.e., that suggest rates high as 150% and higher, as they have a high probability of being AVEs.

3.2 *Step Two:* Imputing the Remaining Missing Observations

After 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 missing data are much more pronounced for preferential tariffs, also for MFN tariffs t_{ikt} of importer i, for product k, in year t reporting-compliance is imperfect. Rather than replacing missing MFN tariffs by linearly interpolating between available observations, I set them equal to the nearest preceding observation. The procedure accounts for anecdotal evidence that countries are more likely to update schedules after a significant tariff change. If

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

[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 tariffs of 1.22 USD + 8% equals 8%.

there is no preceding observation, missing MFN tariffs are set equal to the nearest succeeding observation. Caliendo, Feenstra, et al. (2015) 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 could only fill missing tariffs within the same nomenclature or for those product codes that do not change over the 30 years. In addition, the raw data does not allow 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 only update to the new nomenclature with a significant lack—independently which original source data is used. For example, in WITS, Guyana reports MFN tariffs in 2017 using the outdated HS2002. Using one consistent nomenclature deals with that issue.[21]

Once all 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 it is crucial as otherwise the imputation will do more harm than good. To account for any sort of RTA, I combine Mario Larch's RTA Database (Egger and 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 the DESTA data (Dür et al. 2014). Only when for a country pair ij at least one of the three mentioned databases on RTAs confirms that there is in fact an RTA in force, I impute the preferential tariff.

The imputation method should account for tariff cuts that are 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, i.e., RTAs, for which all tariffs are cut immediately. The DESTA database (Dür et al. 2014) has gathered information for more than 600 RTAs on the phase-out regime 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 entry into force year 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 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.

21 Concordance tables are available through WITS.

The imputation algorithm splits the available data into segments and distinguishes between *(i)* the time before the RTA enters into force marked in gray in Figure 2, *(ii)* the implementation period is depicted in blue, and *(iii)* the years when the RTA is fully in force (red). The time is displayed on the x-axis of the schematic figure. When no phase-out is stipulated in the RTA, the entry into force year is identical to the year when full implementation is reached ($EiF_{ij} = Impl_{ij}$, c.f., Panel (a)).

Figure 2: Imputation Algorithm for Preferential Tariffs

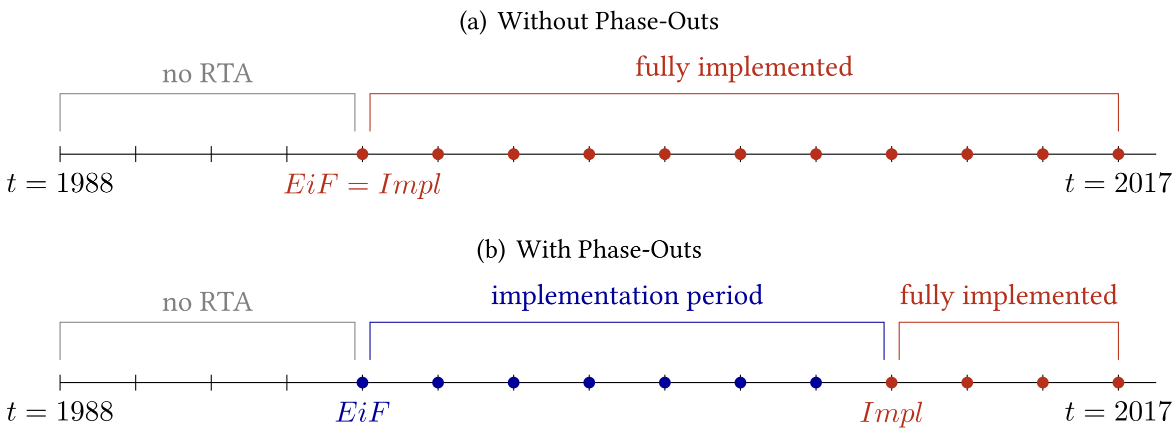

I start with filling the red segment, i.e., when the RTA is fully implemented. For this period, any tariff that can be observed equals the final one, typically zero. Using the year of full implementation as the start date and 2017 as the end date, I sequentially fill the missing observations of the red segment with the first succeeding non-missing observation. If no succeeding 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 already if the importing country only reports 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]. For the last segment (gray in Figure 2), i.e. the time before the RTA was in place, no preferential tariffs are in place and the effectively applied tariff equals the MFN tariff.

[22] C.f., Besedes et al. (2020) for details on NAFTA

In some cases multiple preferential tariff schemes are available, for example, least-developed countries qualify for the Everything but Arms arrangement as well as the General System of Preferences of the EU. Whenever this is the case, I assume that the lowest tariff applies.

3.3 *Step Three:* Retrieving the effectively applied tariff

Now, both, the MFN and the preferential tariffs have no missing observations anymore. Next, I set up a matrix that consists of every pair-product-year combination $ijkt$. For every year t, I combine 197 importer i with 196 exporter 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 nor report tariffs for the whole product range. Figure 3 shows how the imputation method changes the time series the examples given in Section 2.1. For all country pairs the troublesome spikes have vanished.

Figure 3: Imputed Effectively Applied Tariff for Selected Country Pairs

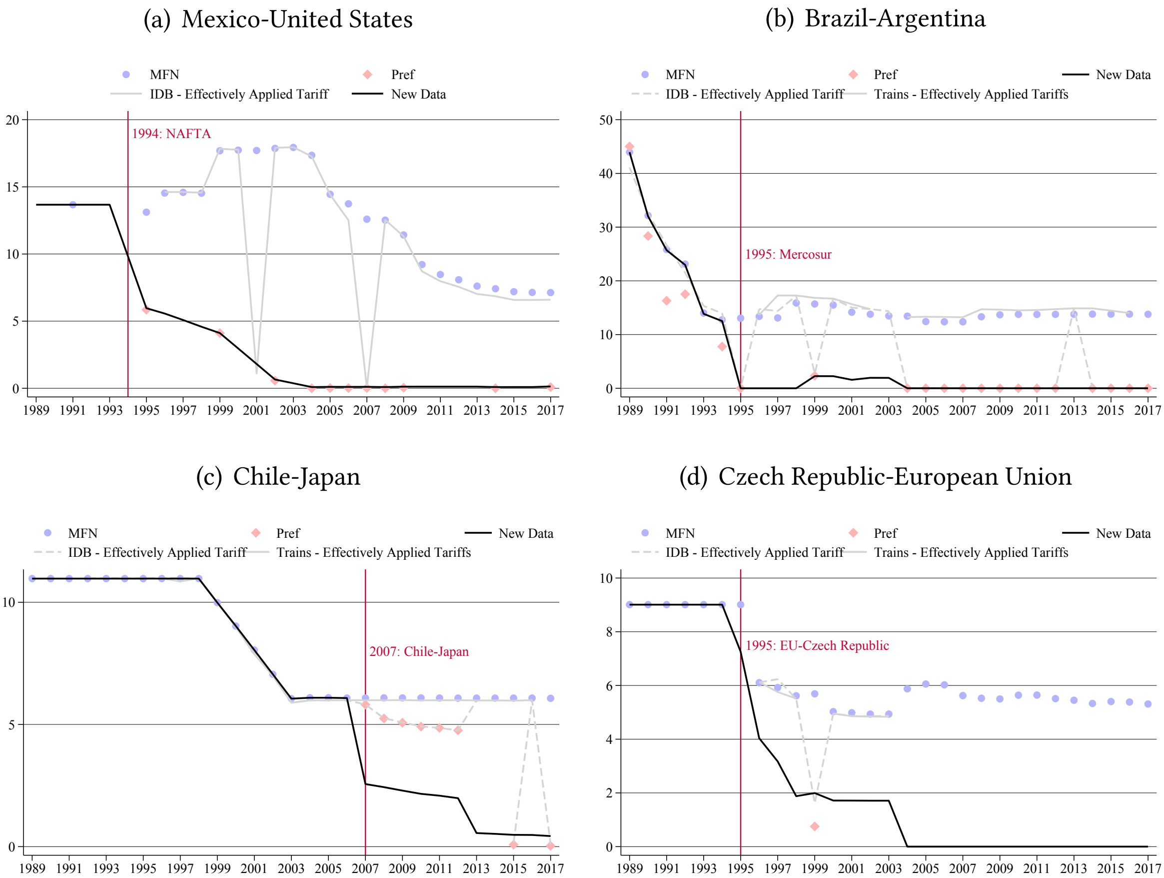

Note: The graph shows the simple average of the imputed effectively applied tariff for selected country pairs.

3.4 Comparison to Other Existing Sources

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

CEPII's *MAcMap-HS6* is based on tariffs reported by the ITC's MAcMap and improves it especially with respect to the non ad valorem tariffs. CEPII first converts all HS6-products into one nomenclature to make inter temporal 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 cover three years (2001, 2004, and 2007), and are publicly available through CEPII's website.[23] Similar to the ITC, the problems of missing observations in the raw data are not addressed.

Table 3: Summary of the Differences to Other Existing Sources

Source	covered years	deals w/ missings	deals w/ phasing-in	deals w/ AVEs	checks RTAs
CEPII's MAcMap-HS6	2001, 04, 07	no	no	yes	no
Caliendo, Feenstra, et al. (2015)	1984-2011	yes	partially	no	no
New GTD	1988-2017	yes	yes	no	yes

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

Caliendo, Feenstra, et al. (2015) have constructed a similar database to the one presented here. However, their dataset differs with respect to covered years (1984–2011) and in terms of the degree of precision of the preferential tariffs. Additionally to the tariffs provided by WITS, they add data from three other sources: manually collected tariff schedules published by the International Customs Tariffs Bureau, U.S. tariff schedules from the US International Trade Commission, and U.S. tariff schedules derived from detailed U.S. 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 phasing-in of preferential tariffs Caliendo, Feenstra, et al. (2015) include information on approximately 100 FTAs and their phasing-in regimes, i.e. whether most tariff lines

[23] The available years correspond with the releases of the GTAP database which CEPII's data on tariffs. The data can be downloaded at http://www.cepii.fr/CEPII/en/bdd_modele/presentation.asp?id=12.

are cut immediately or if phasing-in is common. In my database, I implement a considerable improvement by including detailed phasing-in schedules on the tariff line level for 149 FTAs. For the agreements, for which this information is not available, similar to Caliendo, Feenstra, et al. (2015) the information on the phasing-in regime is used to construct missing preferential tariffs.

Table 3 summarizes the three other existing sources with a comparable country coverage and compares them to the new tariff database. To the best of my knowledge, the data presented here are the first dealing simultaneously with the missing tariffs, accounting explicitly for the phasing-in schedules of RTAs, and cross-validating 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 and years available is unique. Second, the level of precision of the preferential tariffs is much higher than in other existing databases.

4 The Trade Elasticity Revisited

The elasticity of trade flows to trade barriers that is often referred to as "trade elasticity" is a central parameter in international trade because it is crucial when quantifying the gains from trade. Intuitively, a large trade elasticity means that it is easy to substitute with other foreign suppliers, hence, a change in trade costs, for example a reduction in tariffs due to an RTA, will have little impact on welfare. Tariffs are a natural trade barrier that can be used to determine the trade elasticity σ because they vary over time and trade partner and are, at least in principle, easily available for a large set of countries (Hillberry and Hummels 2013). In this section, I will show how much the missing tariffs and the measurement error due to false imputation of missing tariffs biases the estimates of the trade elasticity σ. To do so, I will compare the coefficients for σ when using the falsely imputed tariffs from WITS with the estimates resulting from the new GTD, respectively.

4.1 Identification Strategy

A common approach to estimate σ is derived from the gravity equation and results in some variation the following estimation equation:

$$ln\left(X_{ijt}\right) = -\sigma ln\left(\tau_{ijt}\right) + \gamma RTA_{ijt} + \sum \beta Z_{ij} + \nu_{it} + \nu_{jt} + \epsilon_{ijt} \tag{2}$$

X_{ijt} equals the value of exports of country i to country j at time t; τ_{ijt} is the gross ad valorem tariff $(1 + t_{ijt})$ that i imposes on j, RTA_{ijt} indicates the presence of an RTA (bilateral one as well as nonreciprocal trade arrangements), $\sum Z_{ij}$ is a vector of time-invariant standard gravity controls (distance, colonial ties, common language, contiguity), and ν_{it} and ν_{jt} are included to control for multilateral resistance (Anderson and Van Wincoop 2003). Depending on the exact research question, existing work includes a slightly varying set of control variables, as discussed by Head and Mayer (2014). To check for sensitivity, I will have one specification without any controls and one with. The coefficient σ can then be interpreted as the estimated trade elasticity.

I will follow the two main approaches in the literature and estimate Equation 2 leveraging cross-sectional as well as time-series variation. The two estimation strategies differ with respect to the set of included fixed effects: while most existing work that uses gravity to determine the trade elasticity only includes importer-time and exporter-time fixed effects, ν_{it} and ν_{jt}, i.e., Anderson and Marcouiller (2002), Arkolakis et al. (2018), Fontangé et al. (2022), and Tharakan et al. (2014), more recent work deviates and also adds pair fixed effects μ_{ij} (Felbermayr et al. 2022; Heid et al. 2021). The latter approach can account for all time-invariant omitted variables in a very flexible way, however, the downside is that it takes out much of the available variation making identification hard. The estimation equation changes to

$$ln\left(X_{ijt}\right) = -\sigma ln\left(\tau_{ijt}\right) + \gamma RTA_{ijt} + \nu_{it} + \nu_{jt} + \mu_{ijt} + \epsilon_{ijt}. \quad (3)$$

It is possible to estimate Equations 2 and 3 with OLS as well as PPML. Recent developments recommend PPML because it allows to include zero trade flows and deals with heteroscedasticity (Santos Silva and Tenreyro 2006). I will do both, to show for which empirical method the false imputation leads to the largest bias in the import demand elasticity. Standard errors are always cluster at the pair-level, and pair-sector level when using more disaggregated data.

Intuitively, the false imputation and the resulting measurement error in the standard sources attenuate the estimates of the trade elasticity. When the identifying variation stems from differences in the levels of tariffs, i.e., cross-sectional variation, the false imputation will lead to attenuated results because high MFN tariffs are wrongly associated with high levels of trade while the true, underlying tariff is much lower. With pair fixed effects identification relies on variation in tariffs over time within country pairs, hence, changes from MFN to preferential tariffs are the source of identification. Here, one would falsely assume many more switches from MFN to preferential tariffs to happen, that are created due to the flawed imputation. Most likely the artificial tariff changes are not associated with large increase in trade flows. Hence, also for this empirical approach one would expect the estimates of the trade elasticity to be attenuated.

4.2 Data

The goal is to relate the new estimates using the new GTD to the ones the literature has reported in the past and get a better understanding of how severe the bias due to mismeasured tariff data is. To do so, I download the effectively applied tariff from WITS.[24] Using Trains as a primary source, I download the effectively applied tariff for ISIC Rev. 3 at the four-digit level and aggregate it to the sector level. I use the ISIC Rev. 3 classification instead of the HS classification because it involves less steps of concording across different classifications and nomenclatures to get to the sectors that are used in many standard sources for trade data, i.e. ITPDE and WIOD. I aggregate the data that are available for goods trade at the HS6-digit level to country pairs and for the disaggregated analysis to 18 sectors (agriculture and manufacturing), respectively. For the baseline results, I use simple averages, results using

[24] WITS is available at https://wits.worldbank.org/.

trade-weighted averages are available upon request.The tariff data downloaded through WITS is then combined with the tariff data from the new GTD.

For trade, I use CEPII's BACI (Gaulier et al. 2010) and for robustness WIOD (Timmer et al. 2014), and the first release of the ITPDE (Borchert et al. 2020). The advantage of the BACI data is that is does not include freight and insurance. As **Boehm2022** discuss extensively, this means that the estimated coefficient σ can be directly interpreted as the trade elasticity without of subtracting one. The included gravity controls are from CEPII (distance, colonial ties, common language, and contiguity). Lastly, information on RTAs is included. For partial scope agreements, free trade agreements, and customs unions, we draw information from Mario Larch's RTA database (Egger and Larch 2008). For non-reciprocal arrangements the database on Economic Integration Agreements maintained by Jeffrey Bergstrand and Scott Baier is the main source (Baier, Bergstrand, and Feng 2014), we update it to 2014 ourselves using information of the WTO's PTA database.[25] With these two datasets at hand we know whether there was any RTA in force in the first place and we can distinguish between agreements that were notified under the Enabling Clause, i.e. non-reciprocal trade arrangements and partial scope agreements, and those that were notified under Article XXIV, i.e., free trade agreements and customs unions.

4.3 Results

Figure 4 summarizes the results of 32 different empirical specifications comparing the elasticity σ with tariff data provided by WITS and tariff data using the new GTD using cross-sectional (Panel (a) and (b)) and time-variation (Panel (c) and (d)). Each panel shows the coefficient for σ and the 95% confidence interval for different levels of aggregation (pair and pair-sector), with and without controls, and for different estimators (OLS and PPML). The red dots give the coefficients when using data from Trains available through WITS, the estimates colored gray and blue use the new GTD.

The new GTD does two things, first it cleans the data provided through WITS from the measurement error caused by the false imputation and second, it increases the number of available observations substantially. The latter matters because low-income countries misreport more frequently, hence, not having any information on their tariffs might induce a substantial selection bias. To differentiate between the two biases, I start by keeping the sample constant compared to the available set of observations when using the data from WITS (gray dots). In a second step, I use all available information of the new GTD (blue dots).

[25] http://ptadb.wto.org/.

Figure 4: Import Demand Elasticity Revisited: WITS and new GTD

(a) Cross-Section (t=2005)

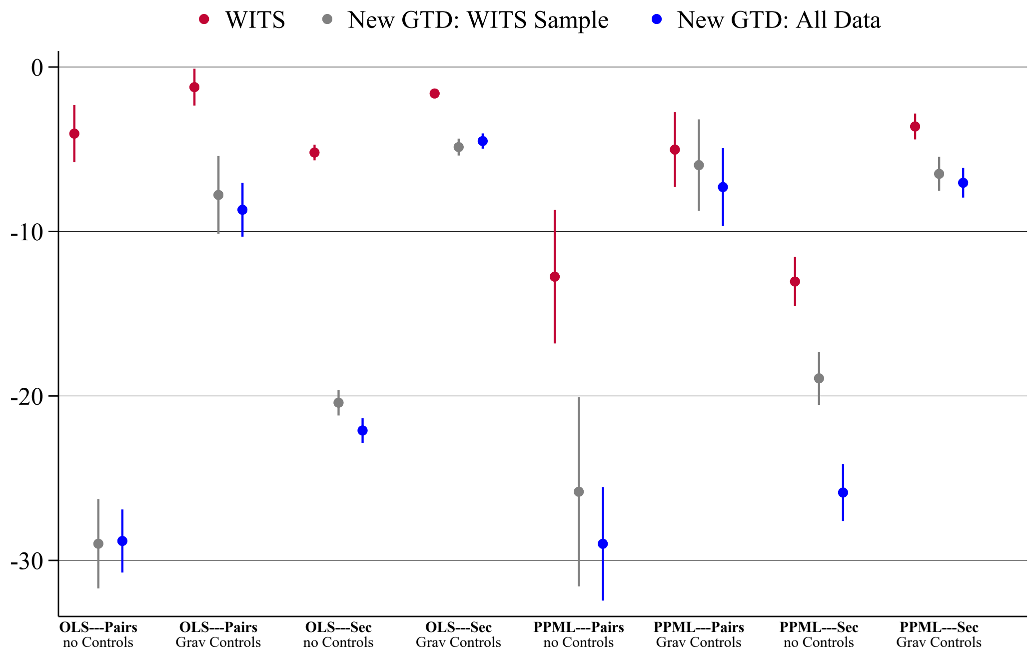

(b) Pooled Cross-Section (t=2000-2014)

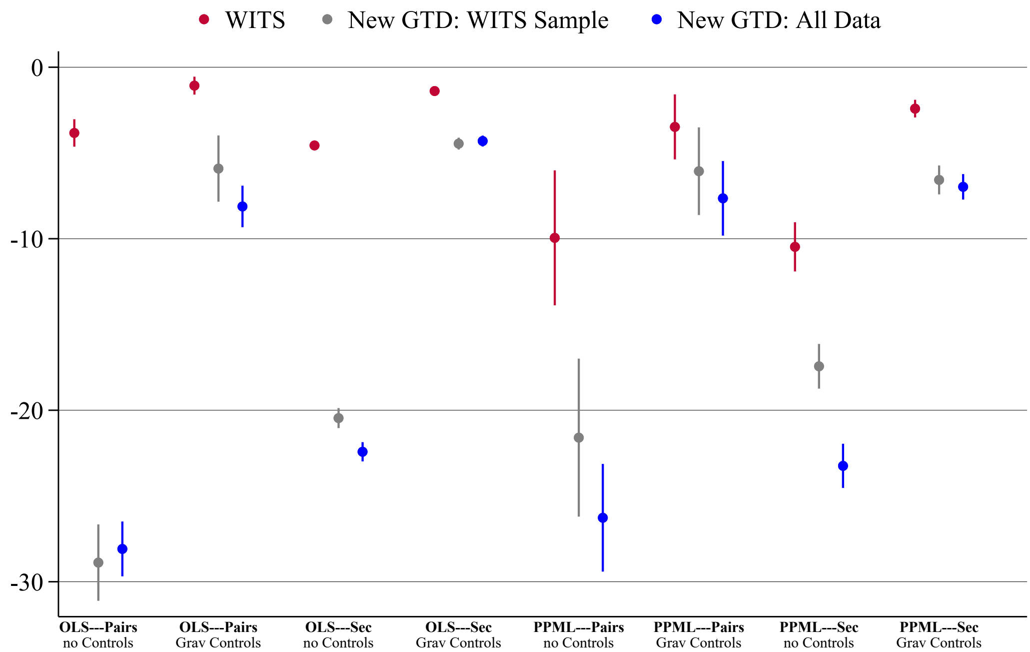

Note: The figure shows the estimated coefficients and the corresponding 95% confidence intervals that result when estimating Equations (2) and (3) using tariffs provided by WITS (red), and the new GTD (gray and blue). Details can be found in the main text. Controls include the standard gravity variables. All standard errors are clustered by pairs and by sector-pairs for the sectoral regressions.

Figure 4: Import Demand Elasticity Revisited: WITS and new GTD — *continued*

(c) Panel (t=2000-2014)

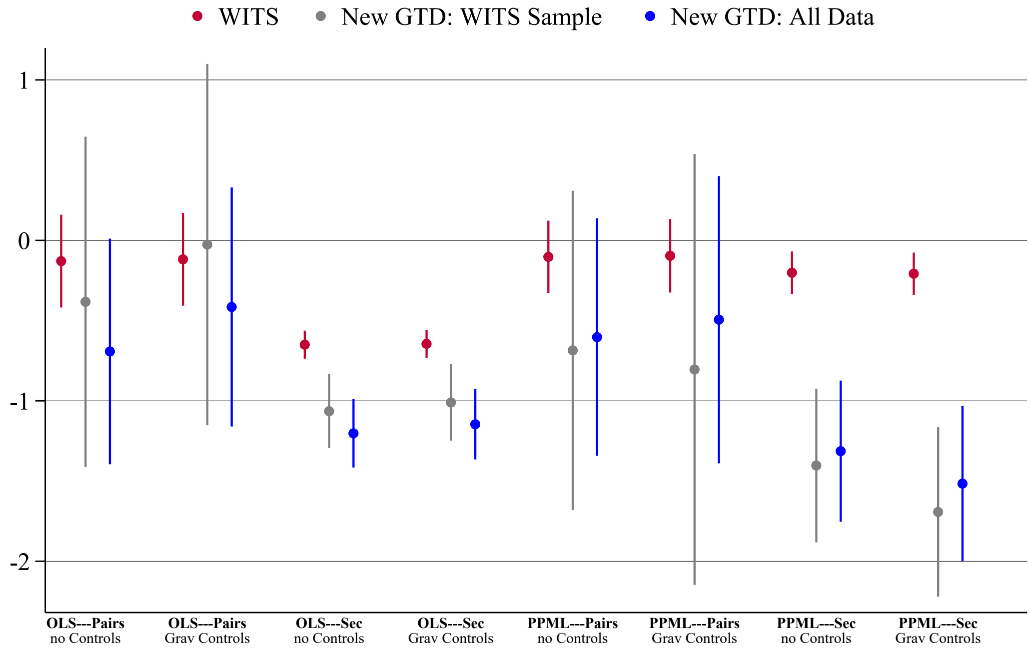

(d) Long Difference (t=2000 & 2014)

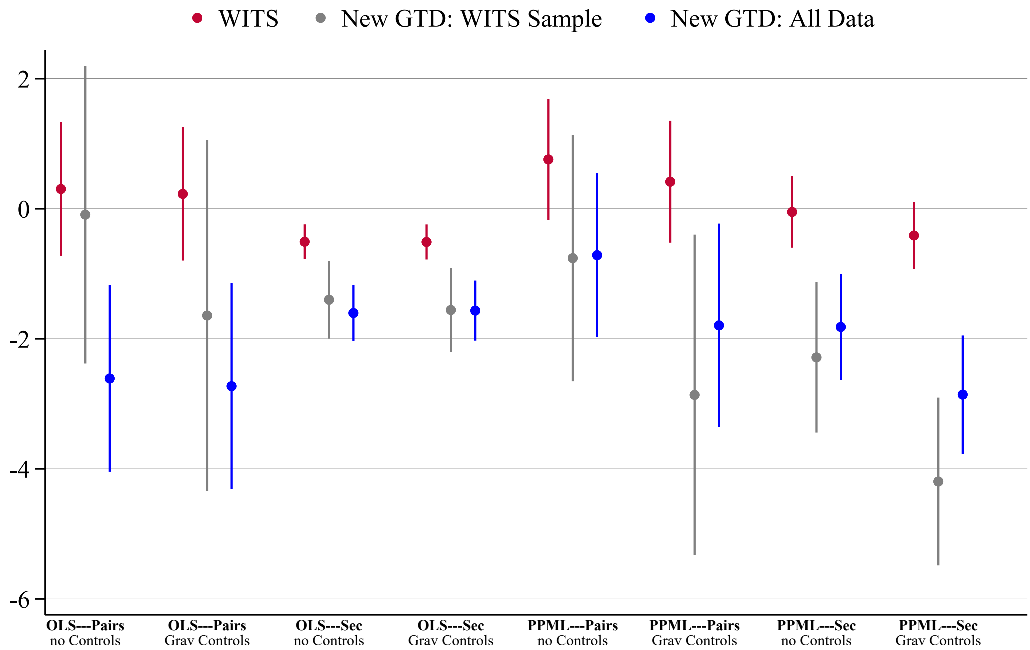

Note: The figure shows the estimated coefficients and the corresponding 95% confidence intervals that result when estimating Equations (2) and (3) using tariffs provided by WITS (red), and the new GTD (gray and blue). Details can be found in the main text. Controls include the standard gravity variables. All standard errors are clustered by pairs and by sector-pairs for the sectoral regressions.

The results show a very strong attenuation bias, independently of the exact specification: the median of the 32 σ-coefficients that result when using the data available through WITS is by a factor of 2.86 smaller than when the new GTD is used for estimation. Hence, the bias caused by the measurement error due to false imputation leads to a substantially smaller trade elasticity. The false imputation is the main cause for differences in the σ's: When correcting for the selection bias caused by the underrepresentation of poor reporters in WITS, I find the elasticity to be at the median 8% larger than when only including the sample of reporters that is included in WITS. Unsurprisingly, the differences due to the selection bias are exclusively driven by the estimates using cross-sectional variation (c.f., Panel (a) and (b)). Low-income countries with high tariffs and little trade are underrepresented in the WITS sample, hence, including them will increase the estimates for σ when the identifying variation comes from comparing different levels in tariffs across importer. However, low-income countries have very few trade agreements. Therefore, when focusing the analysis on within pair variation most of the bad reporter do not have any variation in tariffs, and are therefore dropping out when including pair fixed effects. The same pattern holds when using other trade data (c.f., Appendix A1).

Additionally, to the main results three more nuanced facts stand out: first, the cross-sectional results are much larger than the estimates that use variation over time. When pooling over all available years and including all controls, the estimates for σ vary between -4.30 and -8.11, while when adding the pair fixed effects the coefficients drop to -0.41 to -1.51. This result is in line with existing findings (**Boehm2022**; Goldberg and Pavcnik 2016). Second, when including zero trade flows, which is only possible when using PPML as the estimator, the coefficients increase substantially. Lastly, more disaggregated data yields larger estimates as well.

5 Conclusion

Despite the omnipresence of tariffs in the existing literature of international trade as well as other fields of economics such as innovation, labor, inequality, and development, standard sources of tariffs are severely flawed. This paper is the first to document how failure of reporting tariffs every year leads not only to missing observations but also due to false imputation by WITS, the standard data provider for tariffs, to sever measurement error. Many countries, especially low-income countries, fail to report tariffs every year and missing data are more prevalent for preferential than for MFN tariffs. WITS, the main data provider for tariffs, falsely interpolates missing preferential tariffs with MFN tariffs. This practice leads to artificial spikes in bilateral time series data and, hence, induces massive measurement error. I introduce a new global tariff dataset at the six-digit product level for 197 countries and 30 years that combines five different sources for tariffs and proposes a new interpolation algorithm taking the misreporting into account. Lastly, I show using gravity that correcting for the messy data increases the estimates of the trade elasticity by a factor of 2.89.

Looking up tariffs worldwide and for all tariff lines is not only important for researchers but is crucial for policy makers as well as for exporting firms, too. This paper stresses the need for action with respect to the quality of tariff data. To resolve the issue of missing data, the international institutions could make the access to WITS conditional on reporting all tariff schemes, i.e., all preferential tariffs and MFN tariffs. In the past, reporting compliance could be significantly raised when bad reporters were faced with the threat of the temporary suspension of access to WITS (Hoda 2001, p. 207). The WTO would be 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, one could make it obligatory to report all tariff schemes for each of the RTAs that are currently in place. This procedure would ensure full coverage of tariffs and eliminate any potential source for false imputation even if WITS continuous to be the facility to download tariffs.

References

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.

Allison, Paul D. (2001), *Missing Data.* Sage publications.

Amiti, Mary and Amit K. Khandelwal (2013), "Import Competition and Quality Upgrading." *Review of Economics and Statistics* 952, 476–490.

Anderson, James E. and Douglas Marcouiller (2002), "Insecurity and the Pattern of Trade: An Empirical Investigation." *Review of Economics and Statistics* 842, 342–352.

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.

Arkolakis, Costas, Natalia Ramondo, Andrés Rodríguez-Clare, and Stephen Yeaple (2018), "Innovation and Production in the Global Economy." *American Economic Review* 1088, 2128–2173.

Baier, Scott L. and Jeffrey H. Bergstrand (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.

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

Borchert, Ingo, Mario Larch, Serge Shikher, and Yoto V. Yotov (2020), "The International Trade and Production Database for Estimation (ITPD-E)." *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.

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.

Caliendo, Lorenzo, Robert C. Feenstra, John Romalis, and Alan M. Taylor (2015), "Tariff Reductions, Entry, and Welfare: Theory and Evidence for the Last Two Decades." *NBER Working Paper No. 21768.*

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

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.

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.

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.

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.

Enders, Craig K. (2022), *Applied Missing Data Analysis.* Second. Guilford Publications.

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.

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.

Fontagé, 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.*

Gaulier, Guillaume, Julien Martin, Méjean Isabelle, and Soledad Zignago (2010), "BACI: International Trade Database at the Product-Level. The 1994-2007 Version." CEPII Working Papers No. 2010-23. URL: http://www.cepii.fr/CEPII/fr/publications/wp/abstract.asp?NoDoc=2726.

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*.

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.

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.

Heid, Benedikt, Mario Larch, and Yoto V. Yotov (2021), "Estimating the Effects of Non-discriminatory Trade Policies within Structural Gravity Models." *Canadian Journal of Economics* 541, 376–409.

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.

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

Kee, Hiau Looi, Cristina Neagu, and Alessandro Nicita (2013), "Is Protectionism on the Rise? Assessing National Trade Policies During the Crisis of 2008." *The Review of Economics and Statistics* 951, 342–346.

Kuenzel, David J. (2022), "Non-tariff Measures: What's Tariffs Got to Do with It?" *Canadian Journal of Economics (forthcoming)*.

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

Santos Silva, J.M.C. and Silvana Tenreyro (2006), "The Log of Gravity." *The Review of Economics and Statistics* 884, 641–658. URL: https://doi.org/10.1162/rest.88.4.641.

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.

Tharakan, P. K. M., Ilke *van* Beveren, and Tom *van* Ourti (2014), "Determinants of India's Software Exports and Goods Exports." *The Review of Economics and Statistics* 874, 776–780.

Timmer, Marcel P., Abdul Azeez Erumban, Bart Los, Robert Stehrer, and Gaaitzen J. De Vries (2014), "Slicing Up Global Calue Chains." *Journal of Economic Perspectives* 282, 99–118.

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.

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)

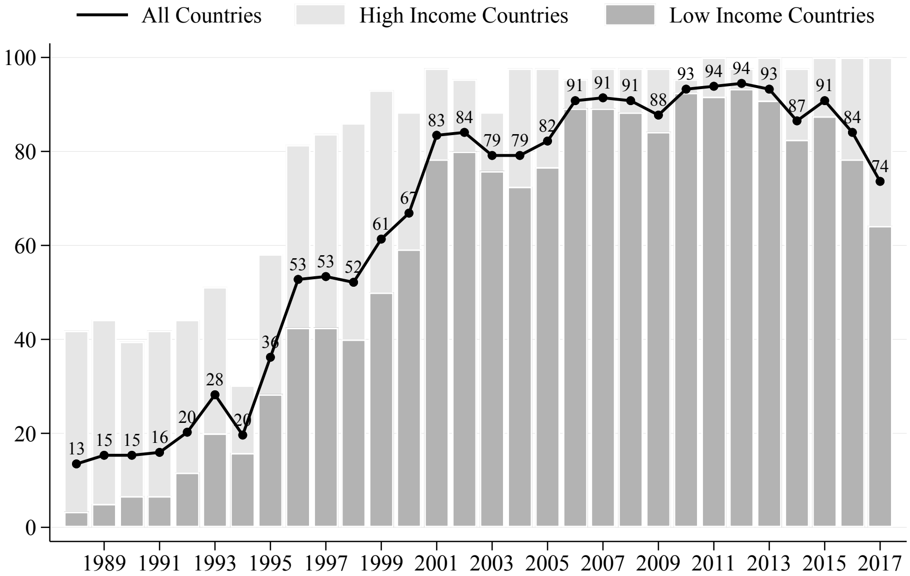

(b) Preferential Tariffs (Trains)

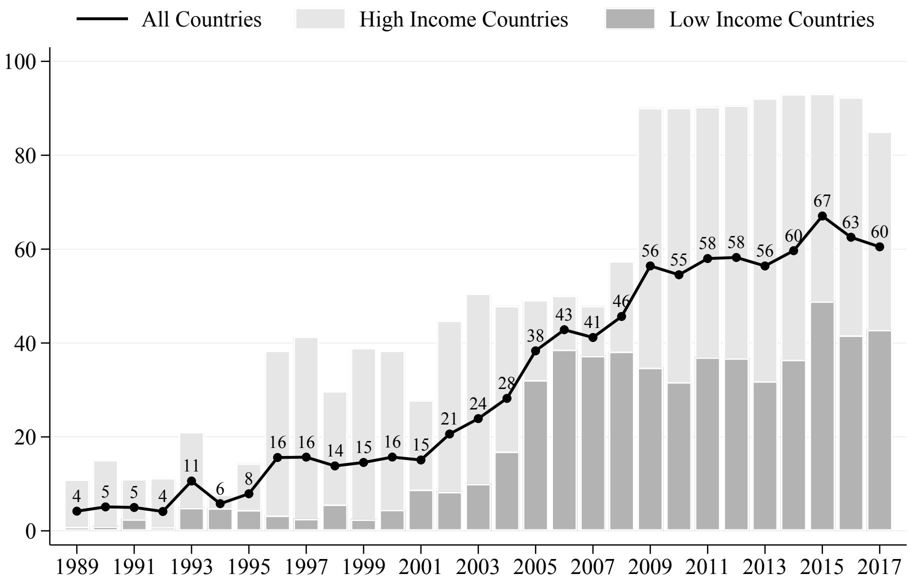

Note: The graphs show the share of reporting countries for MFN (Panel (a)) and preferential (Panel (b)) tariffs. The solid line gives the overall share, the bars split the sample by income. For the MFN tariffs, the number countries reporting tariffs at least once in the period of observation equals 163. 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: Import Demand Elasticity Revisited: WITS and new GTD, Different Trade Data

	Cross-Section (2005)		Pooled C-S (2000-2014)		Panel (2000-2014)		Long Diff. (2000, 2014)	
	(1) WITS	(2) New GTD	(3) WITS	(4) New GTD	(5) WITS	(6) New GTD	(7) WITS	(8) New GTD
OLS, dependent variable ln(X), data aggregated to country pairs								
BACI	-1,22**	-8,68***	-1,07***	-8,12***	-0,12	-0,42	0,23	-2,73***
	(0,57)	(0,83)	(0,27)	(0,62)	(0,15)	(0,38)	(0,52)	(0,81)
N	4.959	11.986	70.698	180.795	70.585	180.614	6.682	22.062
Fixed Effects	it, jt	it, jt	it, jt	it, jt	it, jt, ij	it, jt, ij	it, jt, ij	it, jt, ij
WIOD	-5,07***	-9,84***	-4,94***	-9,85***	-0,91***	-0,52	0,22	-0,54
	(1,74)	(1,84)	(1,43)	(1,54)	(0,28)	(0,62)	(1,53)	(1,16)
N	1.261	1.936	18.163	29.040	18.081	29.040	1.768	3.872
Fixed Effects	it, jt	it, jt	it, jt	it, jt	it, jt, ij	it, jt, ij	it, jt, ij	it, jt, ij
ITPDE	-1,80**	-12,65***	-1,47***	-12,64***	-0,26	-0,79*	-0,17	-3,25***
	(0,83)	(0,99)	(0,50)	(0,79)	(0,16)	(0,42)	(0,57)	(0,85)
N	5.280	12.729	75.400	193.079	75.292	192.931	7.084	23.718
Fixed Effects	it, jt	it, jt	it, jt	it, jt	it, jt, ij	it, jt, ij	it, jt, ij	it, jt, ij
PPML, dependent variable X, data aggregated to country pairs								
BACI	-5,02***	-7,30***	-3,48***	-7,64***	-0,10	-0,49	0,42	-1,79**
	(1,16)	(1,21)	(0,97)	(1,11)	(0,12)	(0,46)	(0,48)	(0,80)
N	5.056	13.806	71.952	207.090	71.809	203.625	6.808	25.570
Fixed Effects	it, jt	it, jt	it, jt	it, jt	it, jt, ij	it, jt, ij	it, jt, ij	it, jt, ij
WIOD	-21,43***	-22,60***	-22,23***	-23,35***	-3,06***	-5,69***	-3,73***	-7,32***
	(2,12)	(2,21)	(1,99)	(2,18)	(0,63)	(0,58)	(0,87)	(0,74)
N	1.261	1.936	18.163	29.040	18.081	29.040	1.768	3.872
Fixed Effects	it, jt	it, jt	it, jt	it, jt	it, jt, ij	it, jt, ij	it, jt, ij	it, jt, ij
ITPDE	-16,86***	-21,82***	-14,01***	-19,47***	-1,14**	-8,49***	-14,46***	-21,68***
	(2,10)	(2,00)	(1,78)	(1,87)	(0,51)	(1,11)	(3,06)	(2,96)
N	5.338	14.387	76.010	215.838	75.908	213.663	7.156	27.212
Fixed Effects	it, jt	it, jt	it, jt	it, jt	it, jt, ij	it, jt, ij	it, jt, ij	it, jt, ij
OLS, dependent variable ln(X), data aggregated to country pairs × sectors								
BACI	-1,61***	-4,50***	-1,39***	-4,30***	-0,64***	-1,15***	-0,51***	-1,57***
	(0,15)	(0,24)	(0,07)	(0,16)	(0,04)	(0,11)	(0,14)	(0,24)
N	105.763	132.867	1.555.017	2.016.788	1.535.006	1.997.489	113.614	219.932
Fixed Effects	it, jt	ikt, jkt	ikt, jkt	ikt, jkt	ikt, jkt, ijk	ikt, jkt, ijk	ikt, jkt, ijk	ikt, jkt, ijk
WIOD	-4,07***	-8,74***	-3,00***	-8,03***	-0,16	-2,20***	-0,76*	-3,26***
	(0,52)	(0,57)	(0,33)	(0,46)	(0,13)	(0,23)	(0,45)	(0,41)
N	22.028	34.479	317.146	517.046	315.628	517.032	30.488	68.732
Fixed Effects	it, jt	ikt, jkt	ikt, jkt	ikt, jkt	ikt, jkt, ijk	ikt, jkt, ijk	ikt, jkt, ijk	ikt, jkt, ijk
ITPDE	-3,12***	-4,02***	-2,70***	-4,00***	-1,14***	-0,96***	-1,09***	-1,33***
	(0,18)	(0,20)	(0,10)	(0,14)	(0,05)	(0,10)	(0,17)	(0,22)
N	121.034	150.286	1.783.483	2.300.411	1.762.693	2.282.443	129.314	250.478
Fixed Effects	it, jt	ikt, jkt	ikt, jkt	ikt, jkt	ikt, jkt, ijk	ikt, jkt, ijk	ikt, jkt, ijk	ikt, jkt, ijk
PPML, dependent variable X, data aggregated to country pairs × sectors								
BACI	-3,61***	-7,04***	-2,41***	-6,97***	-0,21***	-1,52***	-0,41	-2,86***
	(0,40)	(0,46)	(0,26)	(0,38)	(0,07)	(0,25)	(0,26)	(0,46)
N	134.289	246.286	1.956.330	3.693.478	1.821.027	2.914.899	131.932	307.540
Fixed Effects	ikt, jkt	ikt, jkt	ikt, jkt	ikt, jkt	ikt, jkt, ijk	ikt, jkt, ijk	ikt, jkt, ijk	ikt, jkt, ijk
WIOD	-8,38***	-9,30***	-7,61***	-10,04***	-1,60***	-3,13***	-3,66***	-4,45***
	(0,56)	(0,66)	(0,56)	(0,82)	(0,23)	(0,56)	(0,69)	(0,79)
N	22.123	34.584	395.153	654.500	391.692	652.005	37.658	86.596
Fixed Effects	ikt, jkt	ikt, jkt	ikt, jkt	ikt, jkt	ikt, jkt, ijk	ikt, jkt, ijk	ikt, jkt, ijk	ikt, jkt, ijk
ITPDE	-15,26***	-14,42***	-13,81***	-13,16***	-0,98***	-2,19***	-10,47***	-9,02***
	(0,71)	(0,73)	(0,58)	(0,63)	(0,19)	(0,63)	(1,48)	(1,25)
N	135.584	220.510	1.978.342	3.307.747	1.950.469	3.252.939	142.532	352.146
Fixed Effects	ikt, jkt	ikt, jkt	ikt, jkt	ikt, jkt	ikt, jkt, ijk	ikt, jkt, ijk	ikt, jkt, ijk	ikt, jkt, ijk

Note: The table shows the estimated that result when estimating Equations (2) and (3) using tariffs provided by WITS and the new GTD. Columns (1) to (4) include standard gravity controls (distance, colonial ties, common language, contiguity, and RTAs), in columns (5) to (8) the pair fixed-effects absorb all time-invariant controls. All standard errors are clustered by pairs and by sector-pairs for the sectoral regressions.