Do Value-Added Taxes Affect International Trade Flows? Evidence from 30 Years of Tax Reforms

By Youssef Benzarti and Alisa Tazhitdinova*

This paper uses all Value-Added Tax (VAT) changes across EU Member States from 1988 to 2016 to estimate the effect of VATs on trade flows. We find small elasticities of trade flows with respect to VATs, even when VAT changes are large. These elasticities are substantially smaller than the elasticities of trade flows with respect to tariffs estimated in the trade literature. This finding holds across different time periods, countries and types of reforms. Our results imply that VATs are unlikely to distort trade flows.
JEL: F32, F10, H87, H20.
Keywords: value-added taxes, border tax adjustment, international trade.

Value-Added Taxes (VATs) have become the most common form of consumption taxes in the world, rapidly replacing tariffs as well as sales and excise taxes. More than 80% of the world's countries have adopted a form of VAT. The academic consensus is that adjusting VATs at the border – by levying VATs on imports but exempting exports – does not distort trade flows as long as imported goods are subject to the same VAT rate as domestic goods. For this reason, VATs, as they are currently implemented, are considered to be trade neutral, and the World Trade Organization (WTO) allows border adjustment of VATs, while it does not allow border adjustment of direct taxes (i.e., payroll taxes, income taxes, etc.) and strictly regulates tariffs.

This academic consensus, however, has been repeatedly questioned by policy makers in debates going back at least to the 1970s, and intensifying more recently. For example, when the WTO decided to treat direct and indirect taxes differentially in 1970, a number of countries disagreed and argued that exempting VATs (and other indirect taxes) is not trade neutral.[1] As a response, in 1971, the U.S. exempted exports from corporate income taxes on the basis that European

* Youssef Benzarti: University of California Santa Barbara and National Bureau of Economic Research (benzarti@ucsb.edu); Alisa Tazhitdinova: University of California Santa Barbara and National Bureau of Economic Research (tazhitda@ucsb.edu). We thank Pinelopi Goldberg, Jim Hines, Brett McCully, Pau Pujolas, Bob Rijkers, Joel Slemrod, the editor (John Friedman), three anonymous referees and seminar participants at the Conference on the Value Added of Value-Added Taxes and National Tax Association meetings for helpful comments. We also thank the Social Sciences and Humanities Research Council of Canada for financial support.

[1] See §8 of the *Report by the Working Party on Border Tax Adjustments* (available at https://www.wto.org/gatt_docs/English/SULPDF/90840088.pdf). The main reason put forward for allowing exemptions of VATs and other indirect taxes was that these exemptions have been in place for twenty years and "had proved fairly adequate and easy to administer" and "no motive could be found to change them" (see §9).

countries were able to exempt their exports from VATs.[2] This legislation was challenged by the European Union (EU), and the WTO disagreed with the U.S. on the grounds that subsidies to direct taxes are not considered trade neutral, whereas exemptions for exports are.

In spite of the academic and WTO consensus on the trade neutrality of the border adjustment of VATs, there is limited empirical evidence on this question. Grossman (1980) and Feldstein and Krugman (1990) apply the symmetry theorem from Lerner (1936) to provide theoretical arguments in favor of VATs' trade neutrality: they show that an idealized VAT leaves nominal factor prices unchanged and thus does not affect trade flows. Their theory relies on plausible but also strong assumptions, disregarding a number of practical issues associated with the implementation and administration of VATs. First, VATs are often implemented as a substitute to taxes that are likely to distort trade balances, such as corporate, income or payroll taxes. If VAT changes are offset by changes in distortionary taxes, international trade flows might be affected, as shown theoretically in Feldstein and Krugman (1990).[3] Second, while exports should qualify for full VAT rebates, this does not often happen in practice due to administrative difficulties. For this reason, changes in VATs could hinder exports (Chandra and Long (2013)). Third, VATs vary across products, while the assumptions of Feldstein and Krugman (1990) require a uniform tax rate. Finally, another theoretical argument for the trade neutrality of VATs is that tax incentives could be mitigated through changes in real exchange rates. However, since 1999, many European countries have shared a uniform currency, the Euro, yet impose different VATs. While efforts are being made to harmonize VATs, the process has yet to be completed, which implies that, at least from 1999 onward, exchange rates have not been able to mitigate the incentive effects of European VATs. While all four of these caveats are likely to exist, the magnitude of the frictions they introduce is unknown and could possibly be small enough that – in spite of their prevalence – VATs are still approximately trade neutral.

In this paper, we empirically assess the trade neutrality of VATs by using all VAT changes that occurred in Member States of the European Union from 1988 to 2016. We use two data sources. First, we use historical VAT rates compiled by Benzarti et al. (Forthcoming) which contain detailed information on VAT rates by commodity and country, including the exact dates of the VAT rate changes. Second, we use information on export and import flows from the European Commission's Eurostat database. This trade data records monthly trade flows between members of the European Union and their trading partners from 1988 to 2016, broken down by Combined Nomenclature levels.

We estimate the effect of VATs on trade using a fixed effects regression sim-

[2] In 1971, the U.S. Congress legislated the creation of a new corporate entity called a Domestic International Sales Corporation (DISC), which was exempt from U.S. corporate taxes as long as most of its income was the result of exports.

[3] For example, Buettner and Madzharova (2018) document that declines in tariffs in developing countries were accompanied by increases in VATs.

ilarly to Fuest, Peichl and Siegloch (2018). Our identification strategy relies on plausibly exogenous VAT rate changes across EU members, and allows for contemporaneous, lagged, and anticipatory effects of tax rate changes. Our approach contrasts with the trade literature, which has mostly focused on estimating the effect of specific trade agreements, often restricted to a specific set of commodities. In theory, one could collect information on all tariff changes and apply our approach to tariffs (instead of VATs). However, this general approach does not seem practicable for tariffs, since (even if we could systematically collect tariff rates across this thirty-year period) one would still need to know the intermediate input composition of final goods. Since VATs do not apply to intermediate inputs, the data demands are much lighter. To see this, suppose that a country imposes a tariff on steel but not on cars. Then our approach would treat cars as a control group, even though the policy is likely to result in increased imports of cars. Increased steel costs would increase the price of locally produced cars, which would lead to substitution for imported cars. Given this limitation, we do not estimate the effect of tariffs on trade flows using our data, but instead rely on estimates from the vast trade literature on the topic, and implement similar specifications to the ones they use in order to compare our estimates.

We find that changes in VAT rates have little to no effect on imports or exports. We estimate upper bounds on trade flow elasticities with respect to VAT rates that are substantially smaller than the previously estimated trade flow elasticities with respect to tariffs. Our finding holds across different specifications, countries, and time periods, both for small and large VAT changes, and also for VAT increases and decreases. Importantly, our results do not suffer from measurement errors due to timing, because VATs on imports are collected immediately at the border.[4] Our results are also robust to price changes due to the passthrough of VATs, because we consider trade flow measures both in dollars and in tons. Finally, our results suggest that any border adjustment taxes that are theoretically equivalent to value-added taxes are likely to be trade-neutral (Auerbach (2019)). However, theoretically equivalent taxes are not always equivalent in practice due to differences in enforcement and implementation, potentially limiting the applicability of our results to these types of taxes.

A potential concern with our approach is that the reforms we use could be endogenous to economic conditions or could be contemporaneous to other tax changes. We address this concern in two ways. First, using a narrative approach, in the spirit of Romer and Romer (2010), we document the underlying reasons for the VAT changes and select out any changes implemented as part of stimulus packages or austerity measures and find similar trade flow elasticities with respect to VAT rates. Second, we find no significant pre-trends in trade flows prior to these reforms.

Our findings contribute to the literature in three ways. First, in spite of the

[4]The trade flow data we use provides monthly records of imports and exports, allowing us to accurately assign applicable tax rates.

prevalence of VATs in the world, there has been limited work done analyzing them. Our paper helps further our understanding of how VATs affect the economy. This paper is most related to a literature that estimates the effect of VATs on trade, such as Desai and Hines (2005), Keen and Syed (2006), Nicholson (2010) and, more recently, Freund and Gagnon (2017) and Sharma (2020). This literature finds mixed results on the effect of VATs on trade flows. We contribute to this literature by improving on identification: we use detailed tax rate information across all EU member countries matched to detailed trade flow information on various categories of products across all OECD countries to causally estimate the relationship between tax rates and imports/exports. Our approach further differs from the previous work in two other aspects. First, we consider VAT changes rather than VAT adoptions: the latter are likely to occur concurrently with other changes in the economy. Particularly, it is possible that VAT introductions replace other taxes, including tariffs, which are known to distort trade flows. Second, our analysis focuses on developed economies with well-functioning VAT systems that are closest to the "idealized" VATs considered in theoretical models.

Second, our paper contributes to the broader Public Finance literature by considering other margins of distortion imposed by VATs beyond direct distortions to the taxed commodity. Indeed, most estimates of the distortionary effects of taxation focus on direct effects but do not consider fiscal externalities. Because of their ubiquity, VATs could have far reaching effects beyond their direct effect on prices. While other papers have focused on several indirect effects, such as the effect of VATs on wages, on the cost of intermediate outputs, or firm entry and exit, few papers have considered the indirect effect that taxes can have on trade.[5]

Third, our paper contributes to the vast trade literature that estimates the effect of trade barriers on trade flows and the implication of this literature for WTO policies. Our paper supports the WTO policy and the prior in the trade literature that VATs, as they are implemented, are not barriers to trade.

I. Theoretical Framework

Applying the Lerner (1936) symmetry theorem, Feldstein and Krugman (1990) show that, under certain conditions, VATs do not affect trade flows. In Section I.A, we formally lay out their framework and provide intuition for this result. In doing so, we closely follow their notation and framework and, naturally, derive the same results. Then, in Section I.B, we discuss the key underlying assumptions of their model, extend these assumptions by following Costinot and Werning (2019), and show how these assumptions may affect their conclusions.

[5] Keen (2007) considers some of the most important criticisms of the VAT; Pomeranz (2015) and Naritomi (2019) study the effect of VAT on tax evasion; while Carbonnier (2007) and Benzarti and Carloni (2019) explore the effect of VATs on prices and firm behaviors.

A. Are VATs Trade Neutral in Theory?

Consider a small open economy that produces three types of goods: exported goods (X), imported goods (M) and non-traded goods (N). The economy lasts for two periods and the country can borrow and lend at the world interest rate r^*.[6] Further, nominal prices are assumed to be constant over time.[7]

We further assume that the production of X, M and N is perfectly competitive and is given by the transformation functions T^1 and T^2 for each period 1 and 2:

$$T^1(Q_X^1, Q_M^1, Q_N^1) = 0,$$

$$T^2(Q_X^2, Q_M^2, Q_N^2, K_X, K_M, K_N) = 0,$$

where Q_j^t denotes the quantity of commodity $j = \{X, M, N\}$ at time $t = \{1, 2\}$, and K_j is the quantity of commodity j saved from the previous period.

We assume that demand is driven by the preferences of a representative agent over the three commodities X, M and N and a discount factor $\delta \in [0, 1]$, and is given by:

$$W = U(C_X^1, C_M^1, C_N^1) + \delta U(C_X^2, C_M^2, C_N^2). \tag{1}$$

We first solve this model in the absence of a VAT and then introduce a VAT to show that it does not change the equilibrium. In equilibrium, and in the absence of a VAT, the consumption of the non-traded good N should be equal to production minus any investments set aside for the future period:

$$C_N^1 = Q_N^1 - K_N, \quad C_N^2 = Q_N^2. \tag{2}$$

The quantity constraint for traded goods, on the other hand, is relaxed by the fact that countries can trade, and is given by:

$$\begin{aligned} P_X^1(Q_X^1 - K_X) + P_M^1(Q_M^1 - K_M) + (1 + r^*)^{-1}[P_X^2 Q_X^2 + P_M^2 Q_M^2] \\ = P_X^2 C_X^2 + P_M^2 C_M^2 + (1 + r^*)^{-1}[P_X^2 C_X^2 + P_M^2 C_M^2], \end{aligned} \tag{3}$$

where P_j^t denotes the price of commodity $j = X, M, N$ in period $t = 1, 2$.

Because the country is a small open economy, it must be a price taker for goods X and M, hence,

$$P_X^1 = P_X^2 = P_X^*, \quad P_M^1 = P_M^2 = P_M^*, \tag{4}$$

[6] This is an important model feature in the context of a VAT because, in contrast to an income tax, VATs exempt savings.

[7] As argued in Feldstein and Krugman (1990), this assumption does not affect the results but allows for less cumbersome notation.

where P_j^* denotes the world price of commodity j. The price of the non-traded good, on the other hand, is determined by setting domestic supply equal to domestic demand. Supply is given by maximizing the sum of the appropriately discounted value of output in period 1 and 2, taking into account equilibrium prices from (4):

$$\begin{aligned}\Pi = P_X^*(Q_X^1 - K_X) + P_M^*(Q_M^1 - K_M) + P_N^1(Q_N^1 - K_N) \\ +(1+r^*)^{-1}[P_X^* Q_X^2 + P_M^* Q_M^2 + P_N^2 Q_N^2].\end{aligned} \tag{5}$$

The demand is given by maximizing the utility of the representative agent given by equation (1) subject to the budget constraints given by equations (2) and (3).

We now introduce a VAT into this model and show that it does not affect any of the above. Following Feldstein and Krugman (1990), we consider a "textbook" ad-valorem VAT that is remitted by the producers and applies at the same rate to all commodities in the economy, including the imported goods, and fully exempts exports and investments. Furthermore, we assume that tax revenue is returned to the agent via a lump-sum transfer.[8] We define domestic consumer prices as prices inclusive of the VAT and denote them by $\tilde{P}_j^t$, for commodity $j = \{X, M, N\}$ in period $t = \{1, 2\}$.

Price of the imported good. — For a small open economy, the price of the imported good must satisfy the no-arbitrage condition. Hence, it is determined by the international price of the good plus the tax: $\tilde{P}_M^t = P_M^*(1+\tau)$, for $t = 1, 2$.

Price of the exported good. — A similar logic applies to the price of the exported good: a domestic producer of the exported good must be indifferent to selling it domestically or exporting it. Since exports are exempt from tax τ but domestic sales of the good are not, we will again have that the domestic price of the exported good should be equal to its international price plus the tax, i.e., $\tilde{P}_X^t = P_X^*(1+\tau)$, for $t = 1, 2$.

Price of the non-traded good. — As before, the equilibrium price of the non-traded good is determined by setting supply equal to demand. Supply is now determined by maximizing the after-tax value of outputs:

$$\begin{aligned}\Pi &= (1+\tau)^{-1}\{\tilde{P}_X^1(Q_X^1 - K_X) + \tilde{P}_M^1(Q_M^1 - K_M) + \tilde{P}_N^1(Q_N^1 - K_N) \\ &\qquad +(1+r^*)^{-1}[\tilde{P}_X^2 Q_X^2 + \tilde{P}_M^2 Q_M^2 + \tilde{P}_N^2 Q_N^2]\}, \\ &= P_X^*(Q_X^1 - K_X) + P_M^*(Q_M^1 - K_M) + \tilde{P}_N^1/(1+\tau)(Q_N^1 - K_N) \\ &\qquad +(1+r^*)^{-1}[P_X^* Q_X^2 + P_M^* Q_M^2 + \tilde{P}_N^2/(1+\tau) Q_N^2].\end{aligned} \tag{6}$$

Equation (6) implies that as long as $\tilde{P}_N^1/(1+\tau) = P_N^1$ and $\tilde{P}_N^2/(1+\tau) = P_N^2$, the maximization function is identical to that of (5), and therefore production

[8] This is not explicitly stated in Feldstein and Krugman (1990) but is implicitly assumed in their model.

incentives are not distorted by the VAT. A sufficient condition for the latter to be satisfied is for the production of good N to be perfectly competitive, as is assumed here and in Feldstein and Krugman (1990), in which case taxes are fully passed through to consumers, which implies that $\tilde{P}_N^1/(1+\tau) = P_N^1$ and $\tilde{P}_N^2/(1+\tau) = P_N^2$.[9]

Furthermore, the requirement that $\tilde{P}_N^1/(1+\tau) = P_N^1$ and $\tilde{P}_N^2/(1+\tau) = P_N^2$ ensures that the demand side decisions are also undistorted. To see this, note that the maximization of equation (1) subject to budget constraints yields three compensated demand functions for each of the goods $j = \{X, M, N\}$, which we denote by $C_j^t = H_j^t(p, W)$, where p is the vector of prices inclusive of tax and W is the level of utility reached. Under standard assumptions, the functions $H(\cdot)$ are all homogeneous of degree zero, implying that if all prices and income are multiplied by a scalar – in this case $1+\tau$, then demand will remain unchanged. Note that, in this case, W will also increase because all revenue raised from the VAT is given back to the representative agent. Intuitively, since all prices increase by the same proportion and the revenue raised from the VAT is given back to the representative agent, relative and absolute demand levels are unaffected by the introduction of a VAT.

This simple model thus shows that for a "textbook" VAT, and in the presence of full pass-through, the VAT does not distort the relative trade flows of goods. Furthermore, if the tax revenue raised by the VAT is given back to the representative agent, the absolute magnitude of trade flows will also be unaffected. The result relies on the simple intuition that if the prices of all commodities in the economy increase in the same proportion because of the VAT, then the tax should not distort relative production and consumption of these commodities.

B. Discussion of Assumptions

The results in the previous section rely on a number of crucial assumptions which we discuss here.

Small open economy. — While this assumption is important for simplicity and tractability, it implicitly implies that the VAT is fully passed through to consumers in the case of traded goods because domestic prices exclusive of the VAT must equal world prices. In practice, this assumption is likely to fail for two reasons. First, many countries that have implemented a VAT – e.g. Germany or France – are not small economies and are likely to affect at least some world prices. Second, other countries – while small open economies in general – are large exporters of certain commodities. If such a country imposes a VAT, the world price of the highly-exported commodity might be affected by the domestically imposed VAT, potentially resulting in a less than 100% pass-through of the VAT to this commodity's price but a 100% pass-through for other commodities. In this case,

[9] Costinot and Werning (2019) extend the neutrality result to economies with imperfect competition under some fairly stringent assumptions, e.g. their model rules out multinational firms.

domestic prices will not adjust equally for all goods, which would distort trade flows. Costinot and Werning (2019) show that the open economy assumption can be relaxed under some conditions.

"Textbook" VAT assumptions. — In practice, real-world VATs differ from textbook VATs in two main ways: First, VAT rates tend to vary across commodities in most countries. Second, export exemptions are not fully claimed. Similar to the assumption above, if the VAT is not imposed on all commodities equally, or if export rebates are claimed unequally, then the VAT is likely to distort the relative prices of goods, therefore leading to trade distortions.[10]

VAT revenue redistribution. — In the Feldstein and Krugman (1990) and Costinot and Werning (2019) models, revenue collected from the VAT is returned to the representative agent, thereby leaving the representative agent just as well off as in the absence of a VAT. If VAT revenues were instead to be thrown away, then the relative demands for the three goods should remain constant, but the quantity demanded will decrease proportionally, in absolute levels. In practice, tax revenue is rarely "returned" to the consumers; therefore, this assumption is likely to fail.

Exchange rate adjustments. — An important implication of these assumptions is that VATs are fully passed through to consumers, thus ensuring that relative prices across sectors are unchanged. Therefore, any set of assumptions that could guarantee that pass-through is full across all three sectors should predict trade neutrality of VATs. If passthrough is equal across goods but is not full – for example, due to imperfect competition or to the presence of large open economies, then the trade neutrality result can still be achieved via an exchange rate adjustment. To see this, note that with a partial pass-through of VATs to consumers, the price of imports after the imposition of a VAT would be higher than that of domestic goods, leading to a reduction in demand for imports. On the other hand, the price of exports, would be lower than the world price, since exports are exempt from the VAT, but the passthrough is less than 100%, leading to an increased demand for exported goods. Therefore, with less than full passthrough, equilibrium can be reached through an exchange rate adjustment: an increase in domestic currency would eliminate arbitrage opportunities and restore trade balance. These theoretical arguments thus suggest that trade neutrality of VATs in the EU may in part rely on the ability of exchange rates to adjust in response to VAT changes, which is likely to be limited because of a common currency.

While the discussed assumptions could be plausible, they are unlikely to be fully satisfied, which implies that, ultimately, the trade neutrality of VATs is an

[10] Costinot and Werning (2019) show that changing unequal taxes proportionally can be trade-neutral. But since taxes are rarely changed proportionally, even this assumption is unlikely to hold in practice. However, introducing unequal VATs is never trade neutral.

empirical question. For this reason, we turn to an empirical investigation of the trade neutrality of VATs.

II. Data and Institutional Background

Value Added Taxes. — VATs are a form of indirect taxation that applies to the value-added of goods and services sold. Sales taxes, excise taxes and VATs are theoretically equivalent, but there are a few differences in the way they are implemented. First, intermediate inputs are subject to VATs, but firms can claim credits for the VAT that they pay on these inputs, which implies that only value-added is ultimately taxed and the statutory incidence of VATs falls on final consumers, similarly to sales taxes. Second, VATs in the EU are included in consumer prices, while sales taxes in the U.S. are not. Third, EU member countries have several VAT rates in place, including a standard rate that applies to most commodities and reduced rates for food, pharmaceuticals, works of art, newspapers, books and certain agricultural products. There are also some commodities that are not subject to VATs, such as commodities offered by not-for-profit institutions. For imports, value-added taxes are due at the moment of import; this ensures that we are able to correctly assign VAT rates based on the month of import.

Importantly for this analysis, VATs, sales taxes and excise taxes in the EU, the US and around the world are border adjustable, i.e., imports are subject to these taxes while exports are not. The WTO makes a clear distinction between sales taxes, excise taxes and VATs on the one hand, and tariffs on the other; the latter are considered to be trade barriers.

Historical VAT Data. — We use the data on historical VAT rates compiled by Benzarti et al. (Forthcoming). This database contains detailed information on VAT rates by commodity and country, and the exact dates of tax rate changes. The data covers all commodities subject to VATs but does not contain information on VAT re-classifications, i.e. commodities that are re-classified from the standard to the reduced VAT rate and vice versa. These re-classifications are relatively rare, as they are only allowed by the European Commission under special circumstances. During the studied period, VAT rate changes vary between 0.40 and 15 percentage points. Half of the VAT changes in our sample are under 2pp, and 75% are under 4pp.

Trade Data. — The information on export and import flows was obtained from the European Commission's Eurostat database.[11] The trade data records trade flows between members of the European Union (28 countries) and their partners from 1988 to 2016. For this reason, the data availability depends on

[11] The trade data (Table DS-016890) is periodically updated. The version of the data used in this paper was downloaded on April 11, 2017.

the year a given country joined the EU or announced such plans: trade data is available starting from 1988 for Austria, Denmark, France, Finland, Germany, Greece, Ireland, Italy, the Netherlands, Portugal, Spain, Sweden and the United Kingdom; while data for Belgium and Luxembourg is available starting from 1999; and data for Bulgaria, the Czech Republic, Cyprus, Estonia, Hungary, Latvia, Lithuania, Malta, Poland, Romania, Slovakia and Slovenia is available from 2000 onward. Finally, trade flows for Croatia are available starting from 2002. The data is classified following the Combined Nomenclature Levels classification (HS2, HS4, HS6 and CN8), which are described in more detail in Appendix Section A. For each product category, Eurostat records monthly import and export flows in euros and tons, as well as trade quantities in supplementary units for some goods, separately for each partner country. Partner countries include all countries of the world.

Destination and Origin Countries. — We follow the way trade economists refer to trade partners. Most trade economics research considers the effect of trade costs on imports. In this case, a destination country is the country that receives the good imported, whereas the origin country is the country that produces the good being imported by the destination country. In this paper, we also consider exports, in which case the destination country is the country that receives the exported good, whereas the origin country is the one that produces it. Occasionally, we refer to a "partner country", which, when considering imports as an outcome variable, corresponds to the origin country and, when considering exports, to the destination country.

Matching the Trade Data with the Tax Data. — Because VATs only apply to final products, the first step in the matching process is to exclude categories of goods that are likely to be used as intermediate products. To do so, we rely on the United Nations Broad Economic Categories (BEC) classification, which divides products into four categories according to a given good's main use: (1) intermediate, (2) consumption, (3) capital goods and (4) others. We exclude categories 1 and 3 as they are not subject to VATs. Category 4 contains some products that can be used both as a means of production or as consumption goods by individuals, e.g. a car, which can be used by firms or individuals. Our main analysis only includes consumption goods (category 3) since they are the most likely to be affected. We then match the BEC categories to the Harmonized System (HS) categories, which is the classification used in the trade data. After dropping intermediate goods, we are left with 320 product categories at the HS4 level.

Summary Statistics and Sample Construction. — The trade data available from Eurostat is very detailed. However, for many importer-exporter pairs, trade volume in many goods is zero. This happens for two reasons: first, the Combined Nomenclature classifies goods into very narrow categories. Second, imports and

exports are recorded for all potential partners, which include all countries of the world, irrespective of how plausible or frequent such trade relationships are. To minimize the impact of infrequent trade observations and thus reduce the extent of measurement error, we aggregate our data along two dimensions: by trading partners and by commodity type. When aggregated by trading partners, our main sample consists of records of imports and exports at the HS4 level from/to individual EU countries to/from all countries of the world combined. In order to aggregate at the commodity level, we aggregate commodities into seven groups that are subject to the same type of VAT rate: commodities subject to the standard VAT rate, food (subject to the reduced VAT rate), pharmaceuticals (often subject to a "super-reduced" VAT rate), books, newspapers, agricultural goods and works of art. Furthermore, we restrict the set of partner-countries to 37 countries: Australia, Canada, China, Japan, Korea, Mexico, Turkey, the UK, the USA, as well as every "eventual" EU country (i.e., all current 27 EU Member States). The rest of the trading partners are combined into one observation. Hence, when aggregated by commodities, our sample consists of records of imports and exports for 7 groups of commodities from/to individual EU countries to/from the subset of partner countries listed above.

Table 1 shows summary statistics on the aggregated samples and the VAT rate variation we use to estimate how trade flows respond to VAT changes. Even when aggregated (either at the partner or commodity level), the sample size is large and ranges from 400,000 to 700,000 observations. VAT rate changes vary between 0.4 and 15 p.p., with a median of 2 p.p. Roughly one third of these changes represent changes of reduced tax rates rather than standard tax rates; approximately one quarter are tax decreases. Note that the number of tax changes counts the number of observations that experience a VAT change. Therefore, if, for example, the standard VAT rate increases in a given country in a given month and it affects N number of HS4 categories, then we count this as N VAT changes. We also provide further details, in Appendix Section A, on how we construct our dataset and what observations are included in each one of the Figures and Tables.

Table 1—: Summary Statistics and Sample Sizes

	Collapsed by Partner				Collapsed by Commodity			
	Imports		Exports		Imports		Exports	
	Euros	Tons	Euros	Tons	Euros	Tons	Euros	Tons
Panel A: Main Samples – Quarterly data:								
N of Observations	772,789	762,061	739,380	718,083	444,556	418,828	447,356	424,004
N of Observation with Lags	763,033	752,434	730,242	709,421	440,284	414,939	443,375	420,326
N of Tax Changes	19,953	19,453	18,661	18,045	8,198	7,908	7,834	7,527
N of Tax Changes of Reduced Rates	5,258	5,225	4,984	4,922	5,684	5,412	5,353	5,083
N of Tax Change of Standard Rates	14,695	14,228	13,677	13,123	2,514	2,496	2,481	2,444
N of Large Tax Changes (i.e. >4pp)	1,359	1,350	1,246	1,220	1,070	992	869	802
N of Non-Austerity Tax Changes	17,819	17,348	16,626	16,065	7,461	7,185	7,115	6,815
N of Tax Increases	16,010	15,562	14,957	14,429	6,531	6,341	6,341	6,145
N of Tax Decreases	3,943	3,891	3,704	3,616	1,667	1,567	1,493	1,382
Mean Absolute Tax Change (pp)	1.91	1.90	1.89	1.88	2.45	2.41	2.31	2.26
Mean Tax Rate Before Reform (pp)	17.87	17.97	17.99	18.07	14.17	14.25	14.06	14.19
Absolute Tax Change Min (pp)	0.40	0.40	0.40	0.40	0.40	0.40	0.40	0.40
Absolute Tax Change P50 (pp)	2	2	2	2	2	2	2	2
Absolute Tax Change Max (pp)	15	15	15	15	15	15	15	15
Panel B: N of Observations in Subsamples – Quarterly data:								
Reduced Rates Only	275,580	273,467	262,862	258,940	356,544	331,651	360,084	337,750
Standard Rates Only	487,453	478,967	467,380	450,481	83,740	83,288	83,291	82,576
VAT Increases Only	667,378	658,047	638,800	621,026	398,730	376,258	405,262	385,059
VAT Decreases Only	463,428	458,102	443,338	431,106	313,600	293,852	319,559	302,034
Large VAT Changes Only	433,140	428,274	414,137	402,808	305,558	286,844	312,705	295,768
No Austerity Changes	638,539	629,071	611,056	593,170	383,287	360,162	388,632	367,813
Euro-zone Importer	330,582	327,654	316,488	308,836	198,670	188,737	205,891	197,554
Non Euro-Zone Importer	432,451	424,780	413,754	400,585	241,614	226,202	237,484	222,772
Partner Rates Sample					309,058	292,435	313,500	297,092
Panel C: N of Observations in Other Samples:								
Monthly	2,272,576	2,225,471	2,125,759	2,044,844	1,197,518	1,121,795	1,218,291	1,148,220
Disaggregated 1st Differences (uncollapsed)	11,145,332	9,786,492	11,371,016	9,939,283				

Notes: Panel A shows summary statistics for the main samples that were used to estimate the results displayed in Figures 1 and 2 and Table 2. Because lags and leads cannot be calculated for all data points, the number of observations in the distributed lag-lead model is smaller. Panel B shows the number of observations for the robustness subsamples used to generate Figures 3 and 4. Panel C shows the number of observations used to generate Figures B.1 and B.2.

III. The Effect of VATs on Trade

A. Empirical Approach

To estimate the effect of VAT rate changes on trade flows, we follow the approach of Evans, Ringel and Stech (1999), Suárez Serrato and Zidar (2016) and Fuest, Peichl and Siegloch (2018) by running the following fixed effects regression and including leads and lags of tax rates:

$$\ln\left(X^j_{nit}\right) = \sum_{k=-K}^{k=K} \theta_k \ln(\tilde{\tau}^j_{n,t+k}) + \delta^j + \mu_n + \kappa_i + \lambda_t + C_{nt} + Y_{it} + \varepsilon^j_{nit}, \tag{7}$$

where X^j_{nit} measures imports in euros or tons for commodity j imported by destination country n from origin country i at time t, $\tilde{\tau}^j_{nt} - 1$ measures the VAT rate in the destination country n at time t for commodity j.[12] Time is measured monthly or quarterly, depending on the specification. Destination country controls C_{nt} include the destination country's GDP and whether the country is a member of the Eurozone. Y_{it} measures the origin country's GDP.

While most of our specifications measure the effect of VATs on imports, a few of our specification estimate the effect on exports. In these cases, X^j_{nit} measures exports in euros or tons for commodity j exported by origin country n to destination country i at time t, $\tilde{\tau}^j_{nt} - 1$ measures the VAT rate in the origin country n at time t for commodity j. Origin country controls C_{nt} include the origin country's GDP and whether the country is a member of the Eurozone. Y_{it} measures the destination country's GDP.

The main coefficient of interest is θ_0. It is the elasticity of trade flows (imports or exports) with respect to the VAT rate at the time of the VAT rate change: for example, if $\theta_0 = 0$, then trade flows do not respond to VAT changes, and if $\theta_0 = 1$, then a one percentage point change in VAT rates leads to a one percent change in trade flows.[13] For $k < 0$, the θ_k coefficients estimate the lagged response of trade flows, k months after the VAT rate change takes place. Conversely, for $k > 0$, the θ_k coefficients estimate anticipatory responses, k months before the VAT rate changes take place.

This fixed effects regression with leads and lags generalizes a difference-in-differences specification with multiple periods, commodities and countries. The main identification assumption is the same as that for difference-in-differences regressions: absent the tax change, there would have been no change in trade flows of the treated relative to the untreated commodities. Figures 1 and 2, which plot the leads and lags coefficient estimates, show no pre-trends in the response of

[12] Throughout the paper, we follow the trade notation, whereby τ measures the ad valorem tax, and $\tilde{\tau} = 1 + \tau$.

[13] Under the assumption that $\log(1 + \tau) \approx \tau$ for small values of τ.

trade flows to VAT rate changes, which supports this identification assumption. The identification is obtained from within-country specific-commodity variation in VAT rates over time.

As described in Section II, to reduce the effect of measurement error and to ensure that our results are not driven by unit, time or commodity aggregation, we run specification (7) along three different dimensions of the data. First, we aggregate the data in two ways: across trading partners and across commodities. When aggregating by trading partners, we focus on trade flows at the HS4 level, using specification

$$\ln\left(X_{nt}^{j}\right) = \sum_{k=-K}^{k=K} \theta_k \ln(\tilde{\tau}_{n,t+k}^{j}) + \delta^{j} + \mu_n + \lambda_t + C_{nt} + \varepsilon_{nt}^{j}. \tag{8}$$

When aggregating by commodities, we collapse by seven categories of VAT rate type, and focus on the following 37 partner-countries: the USA, China, Japan, Mexico, Canada, Korea, Turkey, Australia, the UK, as well as all current 27 EU Member States. The rest of the trading partners are combined into one observation.[14] This results in specification:

$$\ln\left(X_{nit}\right) = \sum_{k=-K}^{k=K} \theta_k \ln(\tilde{\tau}_{n,t+k}^{j'}) + \zeta_{j'} + \mu_n + \kappa_i + \lambda_t + C_{nt} + Y_{it} + \varepsilon_{nit}, \tag{9}$$

where $\zeta_{j'}$ denote VAT rate type fixed effects. Second, we use two different measures of trade, which are volume (tons) and value (euros). And third, we consider both monthly and quarterly trade flows. We include leads and lags within three years of reforms, thus setting $K = 12$ for quarterly data and $K = 36$ for monthly data. Overall, this amounts to running eight different specifications. Standard errors are clustered by country of destination (for imports) or origin (for exports).

Figure 1 plots the results of running specification (8) on quarterly data, for both tons and euros, aggregated across partner countries. Similar results based on specification (9) are available in Figure 2. Monthly level results are reported in Appendix Figure B.2.

Consistently across all specifications, we find that trade flows – be it exports or imports – are barely affected by changes in VAT rates. The estimates of θ_0 range between -1.61 to +0.81. Our estimates imply a lower 95% confidence bound of the elasticity of trade flows, θ_0, in value (euros), with respect to VAT rates of -2.58 and -3.82 for imports and exports, respectively. The bounds on the estimates in tons are equal to -3.5 and -2.95 for imports and exports, respectively. Specifications in tons control for possible price effects if VATs are passed through

[14] We do this because the data is extremely detailed and will include very small countries with no or very little trade, making the bilateral trade matrix very large.

Figure 1. : Distributed Lag-Model: Partner Aggregates

(a) Outcome: Imports in Euros

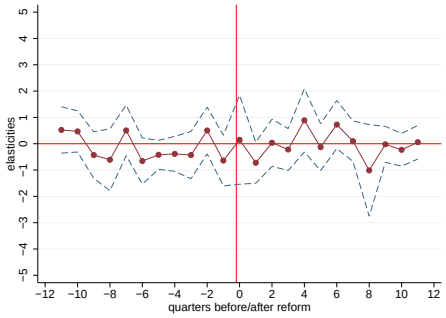

(b) Outcome: Imports in Tons

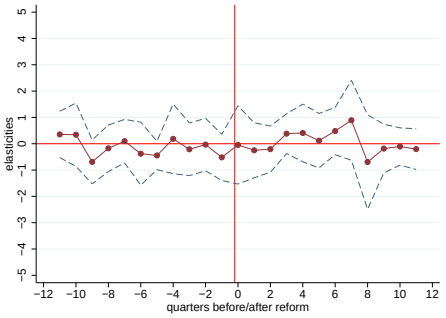

(c) Outcome: Exports in Euros

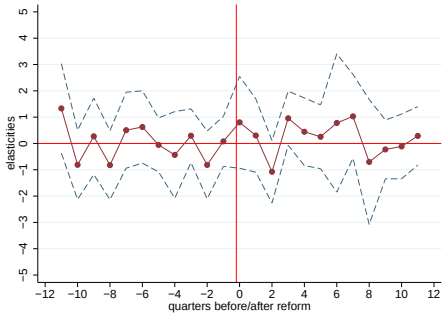

(d) Outcome: Exports in Tons

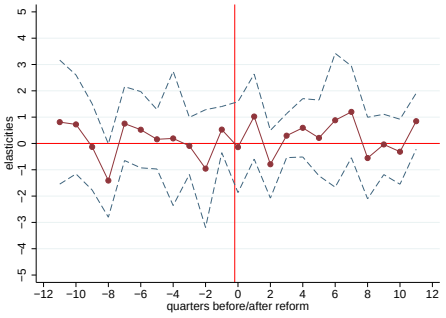

Notes: This figure plots the estimates of θ_k (where θ_0 is the elasticity of trade flows with respect to the VAT rate of the country of destination (for imports) or origin (for exports)) and the corresponding 95% confidence intervals from regression model (8), with 12 quarters of leads and lags included. All specifications include the destination country (for imports) or the origin country (for exports), product category, and year-quarter fixed effects, as well as the following controls: the GDP of the destination country (for imports) or origin country (for exports), whether the destination country (for imports) or the origin country (for exports) is a Eurozone member, whether a tax change is an increase. The outcome variable measures import or export flows in euros or tons. Standard errors are clustered by country of destination (for imports) or origin (for exports). The number of observations are (a) 763,033; (b) 752,434; (c) 730,242; (d) 709,421.

to prices. These elasticities are smaller than trade flow elasticities with respect to tariffs, as discussed in Section IV. Importantly for our identification strategy, we detect no evidence of pre-trends, anticipatory or lagged responses, as can be seen in Figures 1 and 2. This mitigates concerns that VAT changes are implemented as a response to trade flow changes.

Figure 2. : Distributed Lag-Model: Product Category Aggregates

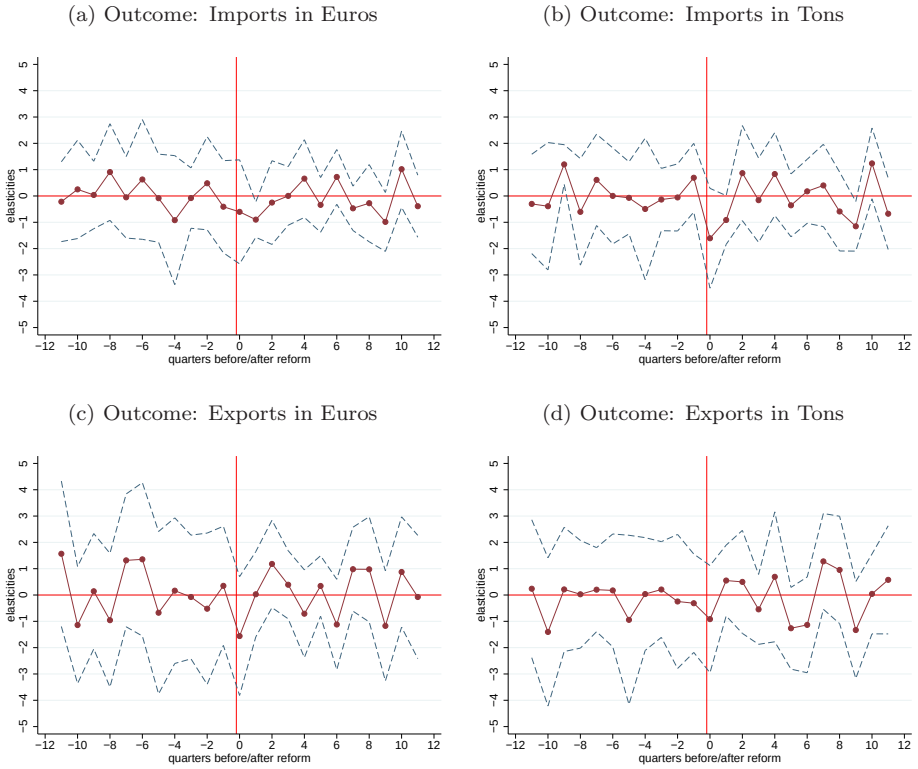

Notes: This figure plots the estimates of θ_k (where θ_0 is the elasticity of trade flows with respect to the VAT rate of the country of destination (for imports) or origin (for exports)) and the corresponding 95% confidence intervals from regression model (9), with 12 quarters of leads and lags included. All specifications include the destination country and origin country, tax rate type, and year-quarter fixed effects, as well as the following controls: the GDP of the destination country and the origin country, whether the destination country (for imports) or the origin country (for exports) is a Eurozone member, whether a tax change is an increase. The outcome variable measures import and export flows in euros or tons. Standard errors are clustered by destination (for imports) or origin (for exports) country. Number of observations: (a) 440,284; (b) 414,939; (c) 443,375; (d) 420,326.

We interpret the lack of trade response as evidence of *relative* trade neutrality: a change in the VAT rate of a given commodity i does not distort the amount of imports or exports of commodity i relative to another untaxed commodity

j. In our view, it remains possible that VATs increase or decrease trade flow *levels* across all commodities in a country. Our interpretation stems from the fact that our fixed effects specification effectively compares changes in imports and exports of commodities in a given country that have experienced a tax change to those that have not, thus differencing out any changes that affect all commodities equally. Our findings therefore do not exclude the possibility that individuals reduce consumption of all goods in response to a VAT increase, or vice versa to a VAT decrease.

B. Robustness Checks

To ensure that our results are not driven by our choice of aggregation, specific reforms or subsets of the data, we perform several robustness checks. For all of the specifications described below, we find coefficients of similar magnitudes as the ones estimated using our main specifications (shown in Figures 1 and 2), which mitigates our concerns that our elasticity estimates are spurious.

We start by ensuring that our results are robust to sample selection. First, we estimate equation (8) for imports measured in euros separately for standard and reduced tax rates (Figures 3(a) and (b)). Second, we run specification (8) separately on VAT increases and decreases in (Figures 3(c) and (d)).[15] Third, to ensure that the absence of response of trade flows to VAT rates is not due to the fact that the VAT rate changes are small, we run specification (8) on the 10% largest VAT changes, i.e. those with tax changes of 4pp or greater (Figure 3(e)). Fourth, to ensure that our reforms do not occur at the same time as recessions, which in turn could affect trade flows, we exclude any reforms that are part of stimulus packages (Figure 3(f)). The excluded reforms are listed in Appendix Section A. Fifth, we consider Eurozone and non-Eurozone countries separately (Figure 3(g) and (h)), to account for potential differences in exchange rate adjustments. Equivalent results for equation (9) are available in Appendix Figure B.3.

Next, we ensure that our results are robust to different estimation approaches. First, we run specification (7) on the disaggregated HS4 categories. Due to the large number of product categories and partners, we perform this analysis in first differences. The results are available in Appendix Figure B.1. Second, we consider different levels of time aggregation and estimate specifications (8) and (9) on monthly data (Appendix Figure B.2). Third, instead of relying on destination and origin country fixed effects, we add year-destination-country and year-origin-country fixed effects (Appendix Figure B.4). Adding such fixed effects significantly reduces our sources of variation, since identification relies on tax changes that happen within the same year in the same origin country (for

[15] We do so because the pass-through of VATs to domestic prices is a key parameter in determining the responsiveness of trade flows to VAT rates, and Benzarti et al. (Forthcoming) show that pass-through rates are different for VAT increases and decreases.

Figure 3. : Distributed Lag-Model: Heterogeneity (Partner Aggregates)

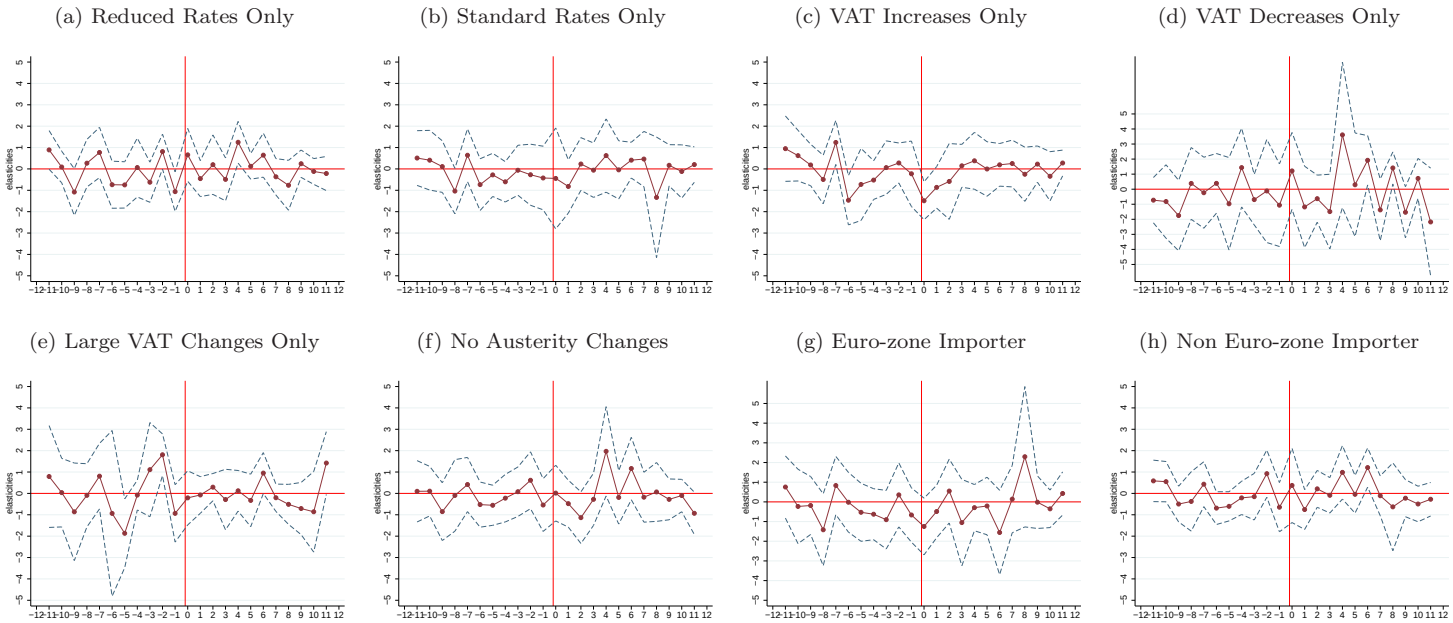

Notes: This figure plots the estimates of θ_k, where θ_0 is the elasticity of imports with respect to the VAT rate of the destination country, and the corresponding 95% confidence intervals from regression model (8), with 12 quarters of leads and lags included. All specifications include destination, product category, and year-quarter fixed effects, as well as the following controls: the GDP of the destination country, whether the destination country is a Eurozone member, whether a tax change is an increase. The outcome variable is import flows in euros. Standard errors are clustered by destination country. The number of observations is shown in Panel B of Table 1.

exports) or destination country (for imports) for different products, which are rare. Nonetheless, all these results are similar to our baseline specification.

Finally, we estimate how imports respond to the VAT rate of the origin country and to the difference between the VAT rate of the destination and origin countries using specification (9) in Figure 4. We find that, similar to when considering the VAT rate of the destination country, imports do not respond to the difference between the VAT rate of the destination and origin countries. We find a small and statistically significant response of imports to the contemporaneous VAT rate of the origin county, but it is still substantially smaller than trade cost elasticities estimated in the trade literature.

Figure 4. : Distributed Lag-Model: Destination and Origin VAT Rates (Product Aggregates)

(a) Difference in Destination and Origin VAT Rates (b) Origin VAT Rates

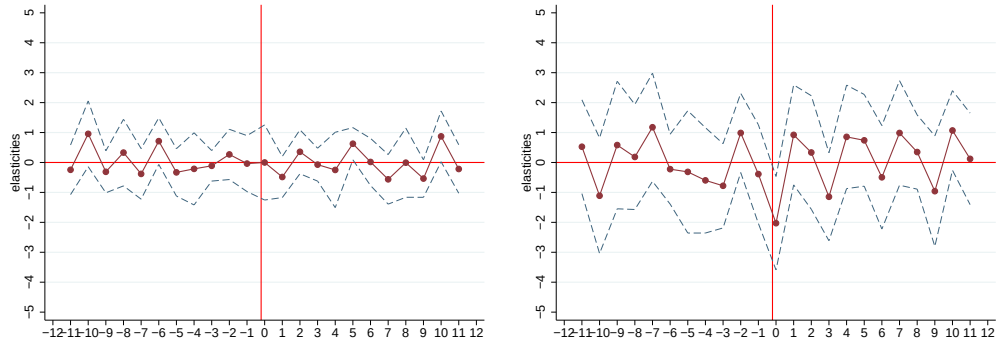

Notes: This figure plots the estimates of θ_k (where θ_0 is the elasticity of imports with respect to VAT rates) and the corresponding 95% confidence intervals from regression model (9), with 12 quarters of leads and lags included. In Figure (a), the tax rate equals the logarithm of one plus the difference between destination and origin VAT rates. In Figure (b), the tax rate measures the logarithm of 1 plus the origin VAT rate. All specifications include destination and origin country, tax rate type and year-quarter fixed effects, as well as the following controls: the GDP of the country of origin and destination, and whether the destination country is a Eurozone member. The outcome variable is import flows in euros. Only origin countries for which tax rates are known are included (i.e. "eventual" EU countries). Standard errors are clustered by destination country. The number of observations for both figures is 309,058.

IV. Relation to Trade Cost Elasticities

The international trade literature estimates elasticities of trade flows with respect to trade costs. Therefore, comparing our VAT elasticity estimates to the elasticity estimates derived in the trade literature allows us to compare the distortionary effects of VATs on trade to the distortionary effects of trade costs overall, and tariffs in particular.

While trade economists are primarily interested in how trade flows respond to all types of trade costs, a substantial share of elasticity estimates are derived using variation in tariff costs. Among these, the most recent elasticity estimate of trade flows with respect to tariffs is from Caliendo and Parro (2015) who find an average elasticity of 4.55 (0.35).[16] Romalis (2007) also uses tariff changes under NAFTA and finds a trade elasticity that ranges between 6.2 and 10.9. Clausing (2001) and Head and Ries (2001) find elasticities between 7 and 11.4. Studies that do not rely on tariff variation typically find similar elasticity estimates: Eaton and Kortum (2002) report elasticity estimates ranging between 3.60 and 12.86, and their preferred estimate is 8.28; Hillberry et al. (2005) find an average elasticity of 17; Broda and Weinstein (2006) find an average elasticity of 17 at the seven-digit (TSUSA), 7 at the three-digit (TSUSA), 12 at the ten-digit (HTS), and 4 at the three-digit (HTS) goods disaggregation. Yi (2003) considers several trade models and shows that elasticities equal to 15 are needed to match the observed bilateral trade flows.

In contrast, our elasticity estimates are systematically smaller than 2, both for imports and exports, in euros and in tons, in most specifications and subsamples of the data we consider. The implied 95% confidence upper and lower bounds from our estimates are smaller than 4 in absolute value. This implies that VATs are unlikely to distort trade flows in the way tariffs do.

That being said, one could argue that our setting and approach are different from the trade literature. We address this concern by discussing the similarities and differences between our approach and the approaches taken in the trade literature. Most trade estimates are derived from gravity models. This approach is summarized in Head and Mayer (2014) Handbook of International Economics Chapter 3: the baseline gravity equation relates bilateral imports, X_{ni} to country n from country i as a function of exporter capabilities S_i, features of the destination market M_n, and the bilateral accessibility of exporter i to importer n, ϕ_{ni}, and takes the following form:

$$\ln(X_{ni}) = \ln(G) + \ln(S_i) + \ln(M_n) + \ln(\phi_{ni}). \tag{10}$$

In this framework, ϕ_{ni} measures the overall impact of trade costs on trade flows and, hence, combines measures of trade costs with the corresponding elasticities (see Definition 1 in Head and Mayer (2014)). Finally, G represents a gravitational constant. As discussed in Head and Mayer (2014), a wide class of models is consistent with this gravity equation, despite differing underlying microstructures of the model.[17]

[16] Note that, while all trade elasticities are reported as positive numbers in the literature, the negative sign is simply omitted (as is often the case). The gravity model implies a negative relationship between trade flows and trade costs, tariffs representing one form of such costs.

[17] For example, in Caliendo and Parro (2015), the elasticity of trade with respect to trade costs measures the dispersion of productivity, rather than the elasticity of substitution as in Armington models, used, for example, in Romalis (2007). As explained by the authors, their parameter of the dispersion of productivity can be related to the elasticity of substitution by restricting producers of the interme-

Earlier studies estimated (10) using countries' GDPs in place of S_i and M_n, but the modern practice estimates gravity equations with time-varying importer $\mu_{n\times t}$ and exporter $\kappa_{i\times t}$ fixed effects, as well as time-varying industry fixed effects $\delta^{j\times t}$ resulting in specification:

$$\ln(X^j_{nit}) = \theta_0 \ln(\tilde{\tau}^j_{nit}) + \delta^{j\times t} + \mu_{n\times t} + \kappa_{i\times t} + \lambda_t + \varepsilon^j_{nit}, \tag{11}$$

where j identifies industries or commodities, n identifies importers, i exporters, and $\tilde{\tau}^j_{ni} - 1$ measures tariffs.

Alternatively, a number of recent studies estimate variants of (11), by taking ratios of imports and therefore estimating multi-differences in logs. In this respect, tariffs and VATs differ in two important dimensions. A given destination country's tariffs vary by product, by time, and by importer, while VATs only differ by product and time. Therefore estimation strategies that rely on differences in tariffs between partner countries (these include Caliendo and Parro (2015) and Romalis (2007)) cannot be replicated using VAT rate variation. This limitation is offset by the fact that VAT changes should not result in spillovers to other commodities, as inputs are effectively tax-exempt. On the other hand, changes in tariffs can result in changes in imports of various commodities since tariffs apply to inputs as well as final goods.

Since VATs do not depend on the exporter i, we cannot include time-varying commodity fixed effects in our estimation procedure. Because we use aggregated data (across commodities and across partners), we estimate two variants of the gravity model (11). For data aggregated across partners, we estimate

$$\ln\left(X^j_{nt}\right) = \theta_0 \ln(\tilde{\tau}^j_{nt}) + \delta^j + \mu_{n\times t} + \lambda_t + C_{nt} + \varepsilon^j_{nt}. \tag{12}$$

and for commodity aggregates

$$\ln\left(X_{nit}\right) = \theta_0 \ln(\tilde{\tau}^{j'}_{nt}) + \zeta_{j'} + \mu_{n\times t} + \kappa_{i\times t} + \lambda_t + C_{nt} + Y_{it} + \varepsilon_{nit}, \tag{13}$$

where j' identify commodity groups based on VAT rate type. These two specifications derived from the gravity equation differ from our main estimation equations (8) and (9) in that they include exporter-year and importer-year fixed effects instead of exporter and importer fixed effects, and do not include lags and leads of tax rates.[18]

diate good aggregate from purchasing goods from the same source, irrespective of trade costs. In this special case, the trade elasticity will be given by the elasticity of substitution as in Armington models. Importantly, both models result in similar gravity-type equations, despite conceptually different effects of tariffs on trade flows: in Caliendo and Parro (2015), trade generates production-side gains, whereas in Romalis (2007), gains are driven by consumption preferences for variety.

[18]Note that these two gravity-model-based estimations are similar to the one we implemented in Appendix Figure B.4, which include year-destination-country and year-origin-country fixed effects but also leads and lags.

Table 2—: Comparing to Trade Literature

	Aggregated by Partner		Aggregated by Product		
	(1)	(2)	(3)	(4)	(5)
$\log(\tilde{\tau}_t)$	-0.852	-0.767	-0.871	-0.784	0.143
	(0.601)	(0.650)	(0.549)	(0.550)	(0.475)
N of observations	772,789	772,789	444,556	444,556	444,556
Quarter-Year FE	Yes	Yes	Yes	Yes	Yes
Product/Tax Rate Type FE	Yes	Yes	Yes	Yes	Yes
Destination FE	Yes	No	Yes	No	Yes
Destination × Year FE	No	Yes	No	Yes	No
Origin FE	No	No	Yes	No	Yes
Origin × FE Year	No	No	No	Yes	No

Notes: This table presents the estimates of θ_0, which is the elasticity of imports with respect to the VAT rate of the country of destination, from $\ln\left(X_{nt}^{j}\right) = \theta_0 \ln(\tilde{\tau}_{nt}^{j}) + \delta^{j} + \mu_{n\times t} + \lambda_t + C_{nt} + \varepsilon_{nt}^{j}$ (regression model 12) in columns (1) and (2) and estimates of θ_0 from $\ln\left(X_{nit}\right) = \theta_0 \ln(\tilde{\tau}_{nt}^{j'}) + \zeta_{j'} + \mu_{n\times t} + \kappa_{i\times t} + \lambda_t + C_{nt} + Y_{it} + \varepsilon_{nit}$ (regression model 13) in columns (3) – (5). The outcome variable is imports in euros in columns (1) – (4), and in column (5) the outcome variable is imports as a share of imports from all countries. The standard errors are clustered at the destination country level.

The results of estimating equations (12) and (13) on imports in euros are shown in Table 2. Note again that these coefficients should be compared to the negative of the ones reported in the trade literature since the trade literature often omits the negative sign when reporting elasticities (as is common in other literatures as well). These estimates are similar to the ones we estimate using our fixed-effect specifications (8) and (9), mitigating concerns that the reason the effect of VATs on trade flows is substantially smaller than that of tariffs or trade costs is due to specification differences.

V. Conclusion

This paper uses a large set of VAT changes to provide empirical evidence on the effect of VATs on trade flows. We find that VATs have very little effect on imports or exports, and our elasticity estimates are substantially lower than the tariff elasticities estimated in the trade literature. Our finding has important implications. First, it fills a gap in the international trade and public finance literatures by providing causal empirical evidence on the effect of VATs on trade flows. Second, it helps settle a long-lasting debate between the U.S., the WTO and EU Member States as to whether VATs should be treated as tariffs.

References

Auerbach, Alan J. 2019. "Tax Equivalences and Their Implications." *Tax Policy and the Economy*, 33: 81–107. Publisher: The University of Chicago Press.

Benzarti, Youssef, and Dorian Carloni. 2019. "Who Really Benefits from Consumption Tax Cuts? Evidence from a Large VAT Reform in France." *American Economic Journal: Economic Policy*, 11(1): 1–26.

Benzarti, Youssef, Dorian Carloni, Jarkko Harju, and Tuomas Kosonen. Forthcoming. "What Goes Up May Not Come Down: Asymmetric Incidence of Value-Added Taxes." *Journal of Political Economy*.

Broda, Christian, and David E Weinstein. 2006. "Globalization and the Gains from Variety." *The Quarterly Journal of Economics*, 121(2): 541–585.

Buettner, Thiess, and Boryana Madzharova. 2018. "WTO membership and the shift to consumption taxes." *World Development*, 108: 197–218.

Caliendo, Lorenzo, and Fernando Parro. 2015. "Estimates of the Trade and Welfare Effects of NAFTA." *The Review of Economic Studies*, 82(1): 1–44.

Carbonnier, Clement. 2007. "Who Pays Sales Taxes? Evidence From French VAT Reforms, 1987–1999." *Journal of Public Economics*, 91(5-6): 1219–1229.

Chandra, Piyush, and Cheryl Long. 2013. "VAT Rebates and Export Performance in China: Firm-Level Evidence." *Journal of Public Economics*, 102: 13–22.

Clausing, Kimberly A. 2001. "Trade Creation and Trade Diversion in the Canada–United States Free Trade Agreement." *Canadian Journal of Economics/Revue Canadienne d'Économique*, 34(3): 677–696.

Costinot, Arnaud, and Ivn Werning. 2019. "Lerner symmetry: A modern treatment." *American Economic Review: Insights*, 1(1): 13–26.

Desai, Mihir, and James Hines. 2005. "Value-Added Taxes and International Trades: The Evidence." *Working Paper*.

Eaton, Jonathan, and Samuel Kortum. 2002. "Technology, Geography, and Trade." *Econometrica*, 70(5): 1741–1779.

Eurostat. 1988-2016. "International Trade Dataset." European Commission. (accessed April 11, 2017).

Evans, William N, Jeanne S Ringel, and Diana Stech. 1999. "Tobacco Taxes and Public Policy to Discourage Smoking." *Tax policy and the Economy*, 13: 1–55.

Feldstein, Martin S., and Paul R. Krugman. 1990. "International Trade Effects of Value-Added Taxation." In *Taxation in the Global Economy*. 263–282. University of Chicago Press.

Freund, Caroline, and Joseph E. Gagnon. 2017. "Effects of Consumption Taxes on Real Exchange Rates and Trade Balances." PIIE Working Paper 17-5.

Fuest, Clemens, Andreas Peichl, and Sebastian Siegloch. 2018. "Do Higher Corporate Taxes Reduce Wages? Micro Evidence from Germany." *American Economic Review*, 108(2): 393–418.

Grossman, Gene M. 1980. "Border Tax Adjustments: Do They Distort Trade?" *Journal of International Economics*, 10(1): 117–128.

Head, Keith, and John Ries. 2001. "Increasing Returns Versus National Product Differentiation as an Explanation for the Pattern of US-Canada Trade." *American Economic Review*, 91(4): 858–876.

Head, Keith, and Thierry Mayer. 2014. "Chapter 3 - Gravity Equations: Workhorse,Toolkit, and Cookbook." In *Handbook of International Economics*. Vol. 4 of *Handbook of International Economics*, , ed. Gita Gopinath, Elhanan Helpman and Kenneth Rogoff, 131–195. Elsevier.

Hillberry, Russell H, Michael A Anderson, Edward J Balistreri, and Alan K Fox. 2005. "Taste Parameters as Model Residuals: Assessing the "Fit" of an Armington Trade Model." *Review of International Economics*, 13(5): 973–984.

Keen, Michael. 2007. "VAT attacks!" *International Tax and Public Finance*, 14(4): 365–381.

Keen, Michael, and Murtaza Syed. 2006. "Domestic Taxes and International Trade: Some Evidence." *Working Paper*.

Lerner, Abba P. 1936. "The symmetry between import and export taxes." *Economica*, 3(11): 306–313. Publisher: JSTOR.

Naritomi, Joana. 2019. "Consumers as Tax Auditors." *American Economic Review*, 109(9): 3031–3072.

Nicholson, Michael W. 2010. "Value-Added Taxes and US Trade Competitiveness." *Working Paper*.

Pomeranz, Dina. 2015. "No Taxation Without Information: Deterrence and Self-enforcement in the Value Added Tax." *The American Economic Review*, 105(8): 2539–2569.

Romalis, John. 2007. "NAFTA's and CUSFTA's Impact on International Trade." *The Review of Economics and Statistics*, 89(3): 416–435.

Romer, Christina D, and David H Romer. 2010. "The Macroeconomic Effects of Tax Changes: Estimates Based on a New Measure of Fiscal Shocks." *American Economic Review*, 100(3): 763–801.

Sharma, Rishi R. 2020. "Does the VAT Tax Exports?" *Economic Inquiry*, 58(1): 225–240.

Suárez Serrato, Juan Carlos, and Owen Zidar. 2016. "Who Benefits from State Corporate Tax Cuts? A Local Labor Markets Approach with Heterogeneous Firms." *American Economic Review*, 106(9): 2582–2624.

UN Trade Statistics. 1988-2016*a*. "Correspondence Tables." United Nations. (accessed April 11, 2017).

UN Trade Statistics. 1988-2016*b*. "Intermediate Goods in Trade Statistics." United Nations. (accessed April 11, 2017).

World Bank. 1988-2016. "Gross Domestic Product Dataset." (accessed April 11, 2017).

Yi, Kei-Mu. 2003. "Can Vertical Specialization Explain the Growth of World Trade?" *Journal of Political Economy*, 111(1): 52–102.

APPENDIX FOR ONLINE PUBLICATION

Sample Selection and Construction

Data Aggregation. — We run our estimation on data that is aggregated either at the partner or at the product category level. We do this in order to reduce noise (some bilateral flows are very infrequent and small), and also because the disaggregated data is very large. In order to aggregate at the partner level, we simply add flows to/from multiple partner countries into one observation. In order to aggregate at the commodity level, we categorize commodities into groups that are subject to the same type of VAT rate and then collapse the data by the VAT type. Practically, we categorize all commodities into seven groups that are subject to the same VAT rate: commodities subject to the standard VAT rate, food (subject to the reduced VAT rate), pharmaceuticals (often subject to a super-reduced VAT rate), books, newspapers, agricultural goods, and works of art, all of which can be subject to reduced rates. We then aggregate our data at this seven-product categories level.

HS Commodity Categories. — The Harmonized Commodity Description and Coding System, often referred to as the Harmonized System (HS), is an international nomenclature for the classification of products, which has been adopted by most countries in the world. It allows countries to categorize their commodities into a unified system, making tracking them (and taxing them) easier across borders. Detailed descriptions and categories of this nomenclature can be found on the United Nations Trade Statistics website.[19].

The HS nomenclature is a six-digit code system that contains approximately 5,300 six-digit categories. The first two digits (HS2) classify a broad category (also called chapter), the next two digits are narrower and the final two digits are even more detailed. For example, the HS6 number for Green tea is 09.02.10, where the HS2 category 09 refers to "coffee, tea, mate and spices", and the HS4 category 09.02 refers to "tea, whether or not flavored". There are 96 HS2 categories. The full list of categories can be downloaded from the the United Nations Trade Statistics website.[20] In this paper we use HS4 categories.

Narrative Approach. — The narrative approach consists of analyzing the historical context of each of the VAT changes using either official documents or press coverage of the reforms. We tag any reforms that are found to have been implemented as part of either fiscal stimulus packages or austerity measures. These

[19] Link: https://unstats.un.org/unsd/tradekb/Knowledgebase/50018/Harmonized-Commodity-Description-and-Coding-Systems-HS

[20] http://unstats.un.org/unsd/tradekb/Attachment439.aspx?AttachmentType=1

reforms are excluded only as a robustness check, in the specification used to estimate Figures 3 (f) and B.3 (h). Specifically, VAT changes that occur in the following countries and year are excluded: Cyprus in March 2012, January 2013 and January 2014; the Czech Republic in January 2008 and January 2012; Estonia in July 2009; Greece in March 2010 and January 2011; Spain in July 2010; Finland in July 2010 and January 2013; France in August 1995; Hungary in July 2009; Ireland in January 2001, July 2011 and January 2012; Lithuania in September 2009; Latvia in January 2009 and January 2011; the Netherlands in October 2012; Poland in January 2011; Portugal in July 2010; Romania in January 2000 and July 2010; the UK in December 2008.

FIGURE 1. — This set of Figures is the result of estimating equation (8) on the full sample of reforms and trade flows where trade flows are aggregated across origin countries (for imports) or destinations countries (for exports). Figure 1 (a) uses the full sample of imports in Euros, Figure 1 (b) the full sample of imports in tons, Figure 1 (c) the full sample of exports in euros and Figure 1 (d) the full sample of exports in tons.

FIGURE 2. — This set of Figures is the result of estimating equation (9) on the full sample of reforms and trade flows where trade flows are aggregated across product categories. Figure 2 (a) uses the full sample of imports in Euros, Figure 2 (b) the full sample of imports in tons, Figure 2 (c) the full sample of exports in euros and Figure 2 (d) the full sample of exports in tons.

FIGURE 3. — This set of Figures plots estimates of equation (8) on different subsamples of the data. We construct all of these subsamples in a similar way. We first identify the "treatment" of interest, i.e., the type of VAT change we want to analyze and then run our specification using only that type of variation on a set of "treated" observations, i.e., commodity*time pairs that experience a VAT change of the type we are considering (VAT increase, reduced rate VAT change etc.), and "control" observations, i.e., commodity*time pairs that do not experience a VAT change. We detail each of the treatment and control observations for each figure below.

Figure 3 (a) considers the effect of changes in the reduced VAT rate on imports. To do this, we estimate (8) on the set of commodities that are subject to reduced VAT rates. Figure 3 (b) similarly estimates (8) on the set of commodities that are subject to standard VAT rates.

The samples for Figures 3 (c), (d), (e) and (f) are all constructed in a similar manner (but different from Figure 3 (a) and (b)). Here, again we only consider specific VAT changes and so restrict the sample of treated observations to the treatment of interest. For the control observations, we include all of them. For

example, when considering the effect of VAT *increases* on imports, we only consider observations that are treated by a VAT increase in our treatment group but include all untreated observations in our control group. An observation is considered "untreated" as long as there was no VAT *decrease* 12 quarters before/after a given date. We proceed similarly for the VAT decreases (Figure 3 (d)), large VAT changes (Figure 3 (e)) and VAT changes that are not part of austerity measures (Figure 3 (f)).

For Figure 3 (g) and (h), only treated and control observations for euro-zone importers are included in Figure 3 (g), and only non-euro-zone importers are included in Figure 3 (h).

FIGURE 4. — This set of figures uses the same underlying dataset as the one used to estimate Figure 2. The only difference is that the estimation uses the VAT rates of both countries of origin and of destination. Therefore, we only include observations in which both countries are EU members. In the case of countries that have joined the EU more recently, trade flow information may be asymmetric. For example, flows from Croatia to France are not included until Croatia joined the EU, but flows from France to Croatia are included, as long as information on Croatian tax rates is available.

TABLE 2. — This table reports the results of several regressions that mirror the approach trade economists have used to estimate the effect of trade costs on imports. It uses the same sample as the one used in Figures 1 and 2.

APPENDIX FIGURE B.1. — This Figure uses a disaggregated version of the data used to estimate Figure 1, estimated using differences (rather than levels).

APPENDIX FIGURE B.2. — This Figure uses the same samples as in Figures 1 and 2 but at monthly intervals.

APPENDIX FIGURE B.3. — This Figure is constructed similarly to Figure 3 but uses data aggregated across the product categories instead of the trading partner level.

APPENDIX FIGURE B.4. — This Figure uses the same samples as in Figures 1 and 2.

Appendix Figures

Figure B.1. : Distributed Lag-Model: HS4 Disaggregated Categories

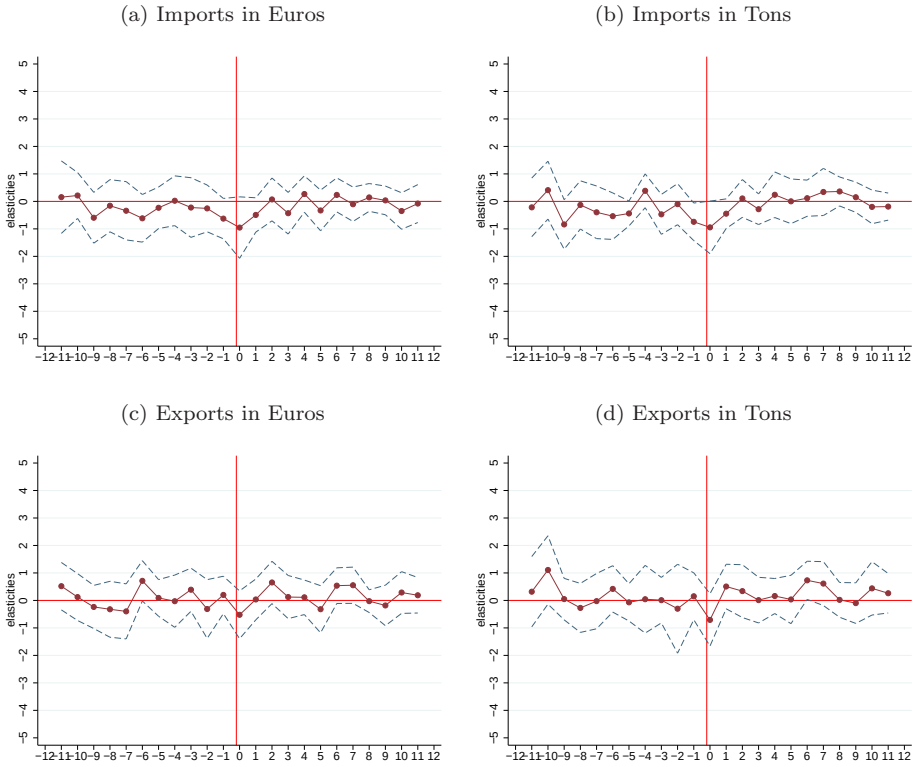

Notes: This figure plots the estimates of θ_k, where θ_0 is the elasticity of trade flows with respect to destination (for imports) or origin (for exports) VAT rates, and the corresponding 95% confidence intervals from regression model (7), with 12 quarters of leads and lags included. The model is estimated in first differences. All specifications include year-quarter fixed effects, as well as the following controls: the GDP of the origin and destination countries, whether the origin country (for exports) or destination country (for imports) is a Eurozone member, whether a tax change is an increase. The outcome variable is imports or exports in euros or in tons, depending on the figure. Standard errors are clustered by destination (for imports) or origin (for exports) country. The number of observations are shown in Panel C of Table 1.

Figure B.2. : Robustness Checks (monthly)

Panel A: Partner Country Aggregates

(a) Imports in Euros (b) Imports in Tons (c) Exports in Euros (d) Exports in Tons

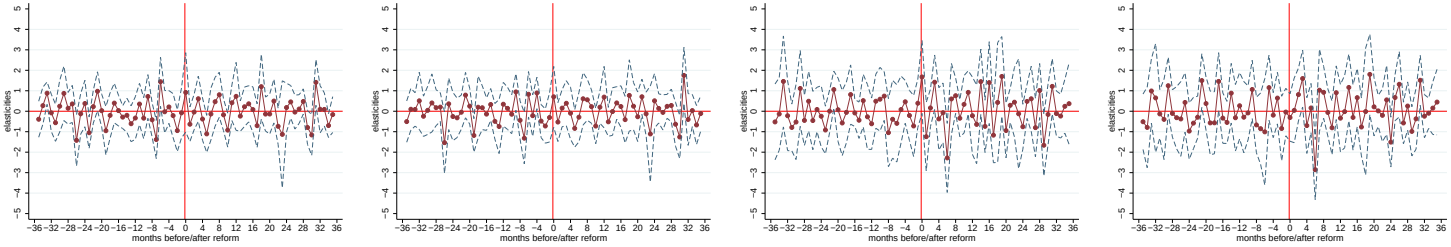

Panel B: Product Category Aggregates

(e) Imports in Euros (f) Imports in Tons (g) Exports in Euros (h) Exports in Tons

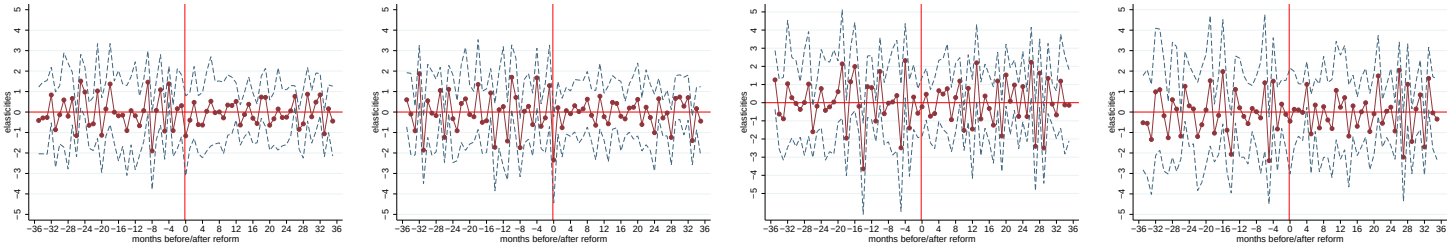

Notes: This figure plots the estimates of θ_k, where θ_0 is the elasticity of trade flows with respect to destination (for imports) or origin (for exports) VAT rates, and the corresponding 95% confidence intervals from regression model (8) and (9), with 36 months of leads and lags included. The outcome variable measures import or export flows in euros or in tons. Standard errors are clustered by destination (for imports) or origin (for exports) country. Specifications in Panel A include the destination country (for imports) or the origin country (for exports), product category, and year-quarter fixed effects, as well as the following controls: the GDP of the destination country (for imports) or origin country (for exports), whether the destination country (for imports) or the origin country (for exports) is a Eurozone member, whether a tax change is an increase. Specifications in Panel B include the destination country and origin country, tax rate type, and year-quarter fixed effects, as well as the following controls: the GDP of the destination country and the origin country, whether the destination country (for imports) or the origin country (for exports) is a Eurozone member, whether a tax change is an increase. The number of observations are shown in Panel C of Table 1.

Figure B.3. : Distributed Lag-Model: Heterogeneity (Product Aggregates)

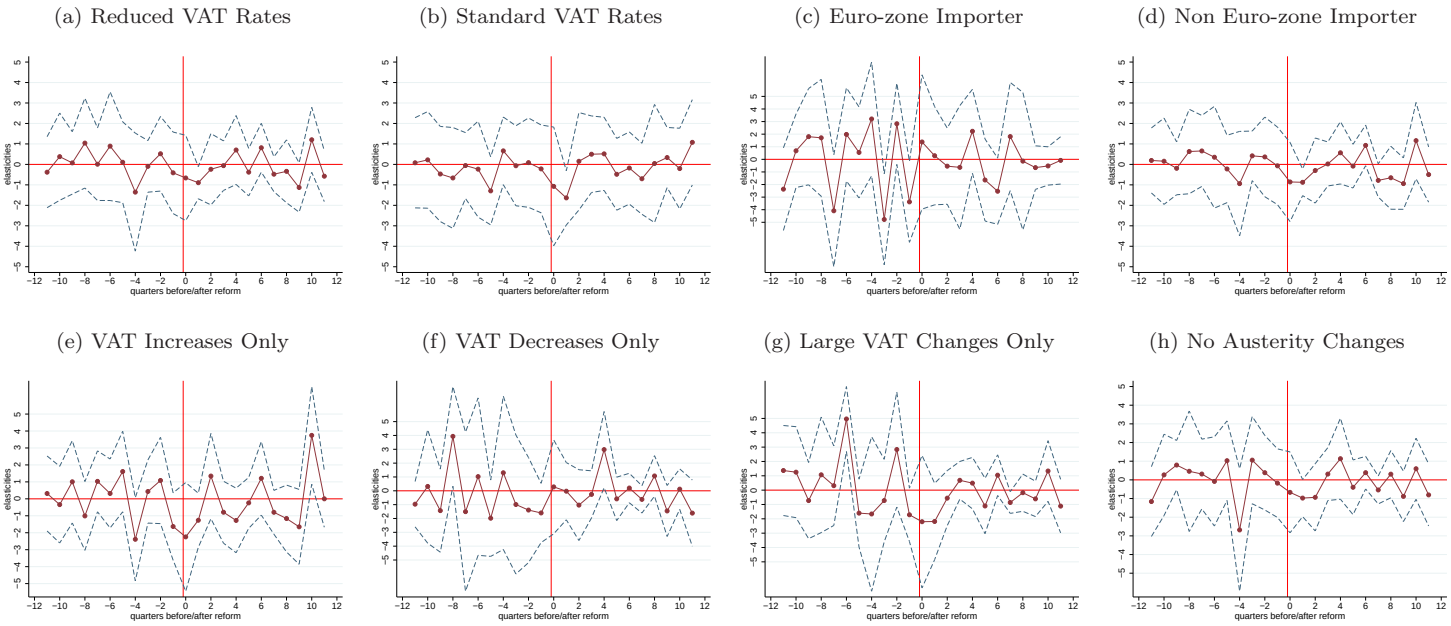

Notes: This figure plots the estimates of θ_k, where θ_0 is the elasticity of imports with respect to destination VAT rates, and the corresponding 95% confidence intervals from regression model (9), with 12 quarters of leads and lags included. All specifications include the destination country and origin country, tax rate type, and year-quarter fixed effects, as well as the following controls: the GDP of the origin and destination countries, whether the destination country is a Eurozone member, whether a tax change is an increase. The outcome variable measures import flows in euros. Standard errors are clustered by destination country. The number of observations are shown in Panel B of Table 1.

Figure B.4. : Robustness Checks: with destination-country-year and origin-country-year fixed effects

Panel A: Partner Country Aggregates

(a) Imports in Euros (b) Imports in Tons (c) Exports in Euros (d) Exports in Tons

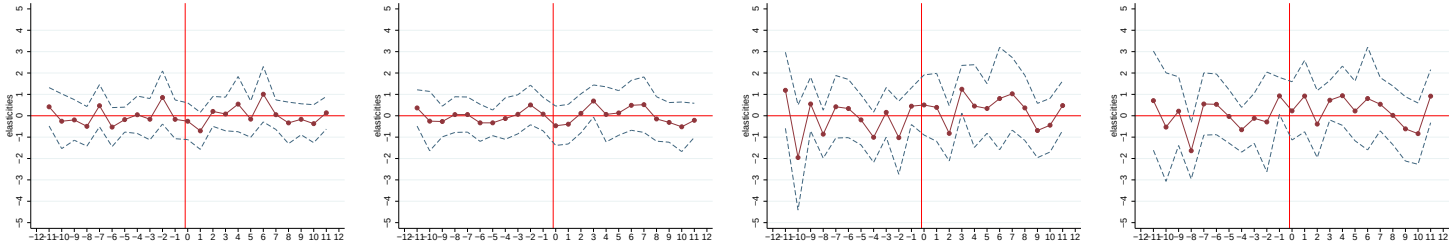

Panel B: Product Category Aggregates

(e) Imports in Euros (f) Imports in Tons (g) Exports in Euros (h) Exports in Tons

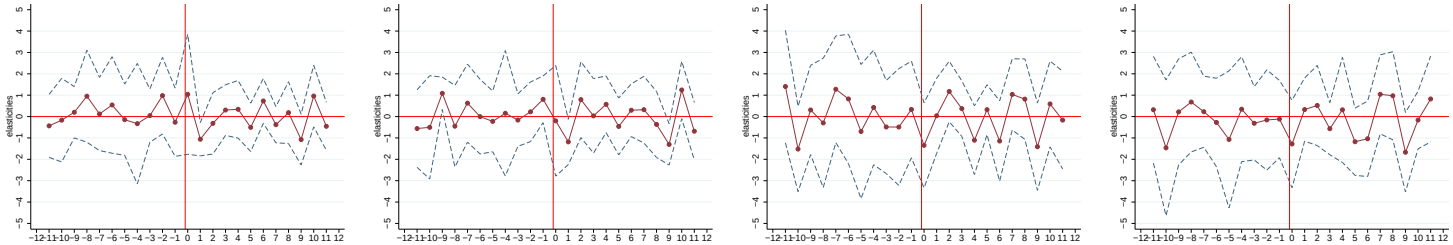

Notes: This figure plots the estimates of θ_k, where θ_0 is the elasticity of trade flows with respect to VAT rate of the country of destination (for imports) or origin (for exports), and the corresponding 95% confidence intervals from regression models (8) and (9), with 12 quarters of leads and lags included. All specifications include year-quarter fixed effects and the following controls: the GDP of the origin country (for exports) or destination country (for imports), whether the origin country (for exports) or destination country (for imports) is a Eurozone member, whether a tax change is an increase. In addition, Panel A includes destination-country-year (for imports) or origin-country-year (for exports) and product category fixed effects, while Panel B includes origin-country-year and destination-country-year, and tax rate type fixed effects. The outcome variable is exports or imports in euros or tons depending on the specification. Standard errors are clustered by destination (for imports) or origin (for exports) country. The number of observations are shown in Panel A of Table 1.