RESEARCH SEMINAR IN INTERNATIONAL ECONOMICS

Gerald R. Ford School of Public Policy
The University of Michigan
Ann Arbor, Michigan 48109-3091

Discussion Paper No. 680

The Long and Short (Run) of Trade Elasticities

Christoph E. Boehm
University of Texas, Austin

Andrei A. Levchenko
University of Michigan, NBEER, & CEPR

Nitya Pandalai-Nayar
University of Texas, Austin & NBER

February, 2021

The Long and Short (Run) of Trade Elasticities*

Christoph E. Boehm
UT Austin

Andrei A. Levchenko
University of Michigan
NBER and CEPR

Nitya Pandalai-Nayar
UT Austin
and NBER

February 2021

Abstract

We propose a novel approach to estimate the trade elasticity at various horizons. When countries change Most Favored Nation (MFN) tariffs, partners that trade on MFN terms experience plausibly exogenous tariff changes. The differential effects on imports from these countries relative to a control group – countries not subject to the MFN tariff scheme – can be used to identify the trade elasticity. We build a panel dataset combining information on product-level tariffs and trade flows covering 1995-2018, and estimate the trade elasticity at short and long horizons using local projections (Jordà, 2005). Our main findings are that the elasticity of tariff-exclusive trade flows in the year following the exogenous tariff change is about -0.76, and the long-run elasticity ranges from -1.75 to -2.25. Our long-run estimates are smaller than typical in the literature, and it takes 7-10 years to converge to the long run, implying that (i) the welfare gains from trade are high and (ii) there are substantial convexities in the costs of adjusting export participation.

Keywords: Trade Elasticity
JEL Codes: F14

*We are grateful to Doireann Fitzgerald for extensive input that improved the paper, and to George Alessandria, Costas Arkolakis, Dominick Bartelme, Chad Bown, Emily Blanchard, Meredith Crowley, Alejandro Cuñat, Matthew Grant, Gene Grossman, Keith Head, Oleg Itskhoki, Amit Khandelwal, Sam Kortum, Kalina Manova, Nina Pavcnik, Andrés Rodríguez-Clare, Sebastian Sotelo, Meredith Startz, Joe Steinberg, as well as seminar and workshop participants at NYU/Columbia, Yale, MIT, Dartmouth, UCSD, UT Austin, VMACS, Virtual ITM, SMU, WCTW (Hawaii), ASSA, Women in International, the NBER ITI Winter Meeting, the Wien Workshop, the IIES, and the ASSA for helpful comments, and to Barthelemy Bonadio, Jaedo Choi, Jonathan Garita and Florian Trouvain for excellent research assistance. Email: `chris.e.boehm@gmail.com`, `alev@umich.edu` and `npnayar@utexas.edu`.

1 Introduction

The elasticity of trade flows to trade barriers – the "trade elasticity" – is the central parameter in international economics. Quantifications of the impact of shocks or trade policies on trade flows, GDP, and welfare hinge on its magnitude. However, there is currently no consensus on the value of this parameter, with a variety of empirical strategies delivering a broad range of estimates.[1]

This paper develops and implements a novel approach to estimating trade elasticities. Following a long tradition in the literature, our empirical strategy exploits variation in tariffs. Our principal contributions are to simultaneously address (i) endogeneity due to possible reverse causality and omitted variables, and (ii) evolution across time horizons. The main results are as follows. First, our estimate of the long-run elasticity of trade values exclusive of tariff payments is -1.75 to -2.25, which is at the lower end of the range of existing estimates. This implies that the welfare-relevant (i.e., tariff-inclusive) long-run elasticity is around 1 in absolute value, and thus the gains from trade implied by most static trade models are large. Second, the trade elasticity in the year following the initial tariff change is -0.76, and it takes several years for it to converge to the long-run value. The trade elasticity point estimates stabilize between years 7 and 10.

To obtain these results, our first contribution is to highlight the role of omitted variables. The theoretical foundations of the gravity equation emphasize the need to control for exporter and importer multilateral resistance terms, structurally (Anderson, 1979; Anderson and van Wincoop, 2003) or with appropriate fixed effects (e.g. Redding and Venables, 2004; Baldwin and Taglioni, 2006). We show that the traditional log-levels gravity specification with multilateral resistance fixed effects yields the conventional wisdom elasticities of -3 to -7. However, multilateral resistance terms do not absorb aggregate or product-specific bilateral taste shocks or other unobserved bilateral gravity variables. For instance, if buyers in a particular importing country have idiosyncratically high tastes for a certain product from a certain country, the policymaker might respond by setting a low tariff. Omitting these unobservables can thus lead to large elasticity estimates. Indeed, once we time-difference the traditional gravity specification to remove bilateral, time-invariant, unobserved gravity variables and taste shifters, conventional OLS estimates fall sharply to around 1 in absolute value.

Our second contribution is to address the residual endogeneity of tariffs. Differencing the data still leaves open the possibility that, for instance, changes in tariffs are caused by changes in trade flows. A surge in imports due to high productivity growth in the exporting country may intensify lobbying for protection and lead to higher tariffs. In this case, estimates that do not account for this reverse causality will be biased towards zero. Our identification strategy relies on the key institutional feature of the WTO system: the MFN principle. Under this principle, a country must apply the

[1] Anderson and van Wincoop (2004) and Head and Mayer (2014) review available estimates.

same tariffs to all its WTO member trade partners. We estimate the trade elasticity based on the response of minor exporters to an importer's MFN tariff change. The identifying assumption is that developments in the minor exporters do not affect a country's decision to change its MFN import tariffs. Our estimation procedure then compares the minor exporters' trade flows to a control group of exporters to the same country to whom MFN tariffs do not apply. These are countries in preferential trade agreements with the importer. Addressing the reverse causality indeed produces larger elasticities in absolute value than OLS.

Our third contribution is to provide estimates over different time horizons, ranging from impact to 10 years. Because tariff changes can be autocorrelated, to estimate elasticities at longer horizons we use time series methods, namely local projections (Jordà, 2005). This approach takes into account the fact that tariffs themselves may have a nontrivial dynamic impulse response structure, implying the elasticities of trade flows at different horizons might depend on the autocorrelation patterns of tariffs. A key advantage of this approach is that we can compare short- and long-run elasticities obtained within the same estimation framework. It is well-known that trade elasticities estimated from cross-sectional variation in tariffs tend to be much higher than the short-run elasticities needed to fit international business cycle moments. Normally, this discrepancy is rationalized by assuming that the elasticities estimated from the cross-section essentially reflect the long run. However, existing estimates either use purely cross-sectional variation (e.g. Caliendo and Parro, 2015), or a time difference over only one horizon (e.g. Head and Ries, 2001; Romalis, 2007). In both cases it is unclear whether what is being estimated is a long-run elasticity, an elasticity over a fixed time horizon, or a mix of short- and long-run elasticities. Our exercise provides mutually consistent estimates of the short- and the long-run elasticities, as well as their full path over time.

Our analysis uses data on global international trade flows from BACI, and tariffs from UN TRAINS. The sample covers 183 economies, over 5,000 HS 6-digit categories, and the time period 1995-2018. These data also allow us to explore sectoral heterogeneity in trade elasticities. Across 11 broad HS sections, the long-run values range from -0.75 to -5.

Our empirical strategy is deliberately not tied to a particular theory, because we expect that our estimates can serve as targets for multiple theories. The mapping between our estimates and structural parameters in theoretical models will then depend on model structure. This is well-understood in the context of static trade models, as multiple microfoundations generate the gravity equation. To illustrate this in a dynamic setting, the final section of the paper presents a simple model, focusing on the minimal common structure required to (i) state the short- and long-run model-implied elasticities and the properties of their time path; (ii) deliver our empirical estimating equations to first order; and (iii) produce a sluggish adjustment of trade to trade cost shocks, consistent with the empirical estimates. The model yields analytical expressions for trade elasticities at all horizons that clarify the determinants of the adjustment dynamics.[2] Our framework nests dynamic versions of the Krugman (1980), Melitz (2003), and Arkolakis (2010) models, as well as extensions with pricing to market (e.g. Burstein, Neves, and Rebelo, 2003; Atkeson and Burstein, 2008).

Finally, we apply our elasticity estimates to the Arkolakis, Costinot, and Rodríguez-Clare (2012) gains from trade formula. To do that, we must account for the fact that our left-hand side variable is trade values exclusive of tariff payments, whereas the elasticity that enters gains from trade formulas is that of tariff-inclusive spending. Our estimates imply an elasticity relevant for computing the welfare gains from trade of about -1. Under this value, the gains from trade are 5-6 times larger than under the commonly used elasticity of -5.

Related Literature Anderson and van Wincoop (2004) and Head and Mayer (2014) review existing trade elasticity estimates. One common approach is to use tariff variation to estimate this elasticity (e.g. Head and Ries, 2001; Romalis, 2007; Caliendo and Parro, 2015; Imbs and Mejean, 2015, 2017). Other methods exploit differences in prices across locations (Eaton and Kortum, 2002; Simonovska and Waugh, 2014; Giri, Yi, and Yilmazkuday, 2020). Existing estimates do not attempt to address the endogeneity of tariffs, and do not distinguish different time horizons.[3] An alternative is to estimate an elasticity of substitution structurally (e.g. Feenstra, 1994; Broda and Weinstein, 2006; Feenstra et al., 2018; Soderbery, 2015, 2018; Fajgelbaum et al., 2020). In some environments the substitution elasticity governs the trade elasticity, but in others it does not. Our empirical strategy is not confined to environments in which the trade elasticity coincides with the elasticity of substitution.

An important recent strand of the literature uses customs data to estimate firm-level elasticities of exports to tariffs, and aggregates firm-level responses to recover macro elasticities (see, among others, Bas, Mayer, and Thoenig, 2017; Fitzgerald and Haller, 2018; Fontagné, Martin, and Orefice, 2018). Often, similar to our strategy, the identifying variation comes from comparisons of MFN and non-MFN destinations. Our approach complements these firm-level analyses. The customs data have the clear advantage of the forensic precision with which different dimensions of firm-level responses to tariffs can be pinned down. On the other hand, this approach normally uses data for a limited set of countries (most often 1) and years, making it challenging to control for multilateral resistance

[2]The recent literature on trade dynamics is rich in both substantive mechanisms and quantification (see, among many others, Costantini and Melitz, 2007; Ruhl, 2008; Burstein and Melitz, 2013; Drozd and Nosal, 2012; Alessandria and Choi, 2014; Alessandria, Choi, and Ruhl, 2014; Ruhl and Willis, 2017; Fitzgerald, Haller, and Yedid-Levi, 2019).

[3]A strand of the literature uses time series methods (most commonly Error Correction Models) to estimate time-varying trade elasticities with trade prices or trade cost changes as independent variables (e.g. Hooper, Johnson, and Marquez, 2000; Gallaway, McDaniel, and Rivera, 2003; Alessandria and Choi, 2019; Yilmazkuday, 2019; Khan and Khederlarian, 2020). Our work builds on this line of research by tackling tariff endogeneity, using flexible local projections, and expanding the scope of the analysis to many more importers, exporters, and products.

terms and/or exploit time series variation in tariffs for identification.[4,5]

Bown and Crowley (2016) describe the empirical features of tariff policy in general, and the MFN system in particular. A property of MFN tariffs important for our purposes is that countries negotiate upper bounds on MFN tariffs, and are then free to set actual MFN tariffs anywhere below those bounds. In the data, a significant fraction of MFN tariffs is actually below the bounds, and thus countries can vary them without violating their WTO commitments. There is a voluminous theoretical and empirical literature on trade policy, both unilateral and within the framework of trade agreements, synthesized most recently in Bagwell and Staiger (2016). This literature emphasizes endogeneity of tariffs to a variety of factors, and thus calls for an effort to overcome that endogeneity in estimation.

The rest of the paper is organized as follows. Section 2 lays out the econometric framework and the identification strategy. Section 3 describes the data, and Section 4 the main results. Section 5 explores the empirical estimates in a number of dimensions, while Section 6 connects the estimates to theory. Section 7 concludes.

2 Estimation Framework

2.1 The horizon-h trade elasticity

As the objective of this paper is to estimate elasticities of trade flows to trade cost shocks at different time horizons, we start with a definition of a horizon-specific trade elasticity. Let i and j index countries, p products, and t time. Let $X_{i,j,p,t}$ be the exports of p from j to i, and $\phi_{i,j,p,t}$ the "iceberg" trade cost. Denote by Δ_h a time difference in a variable between periods $t-1$ and $t+h$: $\Delta_h x_t \equiv x_{t+h} - x_{t-1}$.

Definition. For $\Delta_h \ln \phi_{i,j,p,t} \neq 0$ the *horizon-h trade elasticity* ε^h is defined as

$$\varepsilon^h = \frac{\Delta_h \ln X_{i,j,p,t}}{\Delta_h \ln \phi_{i,j,p,t}}. \tag{2.1}$$

Both conceptually and for the purposes of estimation, it is important to take into account the fact that trade costs follow a stochastic process, and the h-horizon change $\Delta_h \ln \phi_{i,j,p,t}$ is a cumulation

[4] While our work focuses on estimating constant elasticities in line with much of the literature, several contributions also explore environments with non-constant elasticities (e.g. Novy, 2013; Adão, Arkolakis, and Ganapati, 2020; Lind and Ramondo, 2020).

[5] An exception to the common finding of high long-run trade elasticities is Sequeira (2016), who estimates a virtually zero elasticity of trade flows to tariffs for the Mozambique-South Africa preferential trade agreement. The proposed explanation for this result is that high levels of corruption in Mozambique imply that firms rarely pay the tariffs in the first place. This mechanism is unlikely to account for the comparatively low elasticities we find in worldwide data.

of a sequence of period-to-period changes that occur between t and $t+h$. A useful way to think about this horizon-h-specific trade elasticity is as follows. Suppose an unanticipated shock to trade costs occurs at time t. The denominator $\Delta_h \ln \phi_{i,j,p,t}$ captures the effect of this shock on trade costs h periods into the future relative to time $t-1$. It can thus be thought of as a horizon-h impulse response. Similarly, the numerator $\Delta_h \ln X_{i,j,p,t}$ captures the effect of the time-t shock to $\ln \phi_{i,j,p,t}$ *and of the subsequent changes in* $\ln \phi_{i,j,p,t}$ on trade flows h periods into the future.

This discussion makes clear that both the numerator $\Delta_h \ln X_{i,j,p,t}$ and the denominator $\Delta_h \ln \phi_{i,j,p,t}$ can be thought of as sequences following the initial shock. They jointly inform the behavior of dynamic models, which study how trade adjusts to changes in trade costs.

Traditionally, models of international trade are static, representing a metaphor for the long run. Thus, parameterizing these models requires the long-run elasticity ε, defined as the limit:

$$\varepsilon = \lim_{h \to \infty} \varepsilon^h,$$

if it exists. This limit measures the permanent change in trade flows that accompanies a permanent change in trade costs.

2.2 Estimating equations

In practice, we will use tariff variation to estimate ε^h. Let the total trade costs be multiplicative in gross ad valorem tariffs $\tau_{i,j,p,t}$ and non-tariff costs $\kappa_{i,j,p,t}$:[6]

$$\phi_{i,j,p,t} = \kappa_{i,j,p,t} \cdot \tau_{i,j,p,t}.$$

Then $\varepsilon^h = \Delta_h \ln X_{i,j,p,t} / \Delta_h \ln \tau_{i,j,p,t}$.

Consider a change in tariffs $\Delta_0 \ln \tau_{i,j,p,t}$ between $t-1$ and t. We estimate the following equation using local projections (Jordà, 2005):

$$\Delta_h \ln X_{i,j,p,t} = \beta_X^h \Delta_0 \ln \tau_{i,j,p,t} + \delta_{i,p,t}^{X,h} + \delta_{j,p,t}^{X,h} + \delta_{i,j,p}^{X,h} + u_{i,j,p,t}^{X,h}, \quad (2.2)$$

where the δs are fixed effects.

The coefficient β_X^h in (2.2) captures the change in trade flows h periods ahead that follows an initial one-period change in tariffs: $\frac{\Delta_h \ln X_{i,j,p,t}}{\Delta_0 \ln \tau_{i,j,p,t}}$. One might be tempted to use β_X^h as an estimate of the horizon-h trade elasticity. If $\Delta_0 \ln \tau_{i,j,p,t}$ was a one-time change in tariffs (that is, $\Delta_h \ln \tau_{i,j,p,t} = \Delta_0 \ln \tau_{i,j,p,t}$), β_X^h is indeed an estimate of ε^h for all h. We can, and do, estimate β_X^h, but it is often

[6] That is, to streamline notation, $\tau_{i,j,p,t}$ is one plus the ad valorem tariff rate, so that a 5% tariff means $\tau_{i,j,p,t} = 1.05$.

misleading as a measure of the trade elasticity if—following the initial change $\Delta_0 \ln \tau_{i,j,p,t}$—tariffs themselves keep changing during the next h periods. For instance, if a tariff reduction in the initial year tends to be followed by further tariff reductions, we would attribute a large change in trade flows to a small initial tariff change not taking into account the impact of subsequent, dependent, tariff decreases. The opposite would happen if tariffs were mean-reverting, such that initial reductions tend to be followed by increases. The h-period change in trade flows thus conflates the impact of initial and subsequent tariff changes. Below we show that in the data, tariffs do continue to change following an initial impulse.

To account for this, we estimate a local projection of the h-period tariff change on the initial shock in tariffs:

$$\Delta_h \ln \tau_{i,j,p,t} = \beta_\tau^h \Delta_0 \ln \tau_{i,j,p,t} + \delta_{i,p,t}^{\tau,h} + \delta_{j,p,t}^{\tau,h} + \delta_{i,j,p}^{\tau,h} + u_{i,j,p,t}^{\tau,h}, \tag{2.3}$$

where the impact effect of tariffs on tariffs is $\beta_\tau^0 = 1$ by definition. The horizon-h trade elasticity can then be recovered as $\varepsilon^h = \frac{\beta_X^h}{\beta_\tau^h}$.

Further, we estimate the combined specification:

$$\Delta_h \ln X_{i,j,p,t} = \beta^h \Delta_h \ln \tau_{i,j,p,t} + \delta_{i,p,t}^h + \delta_{j,p,t}^h + \delta_{i,j,p}^h + u_{i,j,p,t}^h. \tag{2.4}$$

When $\Delta_0 \ln \tau_{i,j,p,t}$ is used as an instrument for $\Delta_h \ln \tau_{i,j,p,t}$, (2.4) is equivalent to (2.2)-(2.3), and directly identifies the trade elasticity at horizon h using the impulse at time t: $\widehat{\beta}^h$ is an estimator of ε^h. Estimating (2.4) has the advantage that standard errors for the elasticity estimates are easier to compute. To address the potential endogeneity of $\Delta_0 \ln \tau_{i,j,p,t}$, $\Delta_h \ln \tau_{i,j,p,t}$ can instead be instrumented with an exogenous subset of tariff changes, as we describe below.

This estimation is carried out at different horizons $h = 0, ..., H$, to trace the full profile of ε^h over h. In practice, the period length is a year and we use a maximum horizon of $H = 10$. If the estimates of β_X^h and β_τ^h stabilize within 10 years of the shock, we interpret it as convergence of both the numerator and the denominator in (2.1), rendering our estimates informative about the long-run trade elasticity. While Section 6 provides a detailed discussion of the convergence to the long-run elasticity in the context of a conventional class of models, the interpretation of our estimates is not confined to a particular theoretical framework. While the baseline analysis estimates a single elasticity across product categories, below we also implement these specifications for broad product groups to obtain a distribution of β_p^hs.

The estimating equations (2.2)-(2.4) are deliberately not tied to a particular theory. We posit a fairly general estimating equation that can be viewed as time-differenced gravity, and our objective is to develop a set of estimates that can potentially serve as targets for multiple theories. Indeed, it is common in both macroeconomics and trade that multiple microfoundations lead to the same estimating equation. For instance, many business cycle models have a vector autoregressive (VAR) representation (Sims, 1980; Canova and Sala, 2009). In trade, the gravity relationship can be derived from Armington, Ricardian, and monopolistic competition models (Head and Mayer, 2014). We relate the econometric estimates to a tractable dynamic model in Section 6. This model delivers estimating equation (2.2) and illustrates that the fixed effects capture dynamic analogues of multilateral resistance terms.

Conventional estimation A common approach to estimating the long-run elasticity ε starts from a static gravity equation: $\ln X_{i,j,p,t} = \beta \ln \tau_{i,j,p,t} + \delta_{i,p,t} + \delta_{j,p,t} + u_{i,j,p,t}$, and relies on either cross-sectional variation or a single-horizon difference of this equation. The coefficient β is interpreted as an estimate of the long-run elasticity ε.

Conventional approaches cannot speak to the horizon-h trade elasticity. This is immediate for estimates in log-levels, which mostly use cross-sectional variation for identification. However, it is also true for estimates in differences. A research design that estimates an elasticity based on, say, 5-year differences in both tariffs and trade ignores the timing of tariff shocks. A 5-year tariff change of a given magnitude could be due to shocks at the beginning or the end of the five year period. As a result, a 5-year difference specification will estimate a conflation of horizon-0 to horizon-5 trade elasticities. We formalize this argument based on our model in Section 6. Appendix Proposition C.1 shows that estimation in h-period differences does not generally identify the horizon-h trade elasticity. As an example, if tariffs follow a random walk, estimation in h-period differences instead identifies the simple average of horizon-0 to horizon-h trade elasticities, but the result is of course more general than the random walk case. This observation suggests the use of macroeconometric methods to estimate the trade elasticity.

A corollary is that estimation in long differences will not necessarily identify the long-run trade elasticity since many tariff shocks could have occurred close to the end-point of the difference. We will also show below that estimation approaches based predominantly on cross-sectional variation likely suffer from omitted variable problems. Thus, we argue that our long-run estimates are likely preferable to the conventional alternatives even for researchers only interested in the long-run elasticity for calibrating a static trade model.

2.3 Identification

To achieve identification, we control for omitted variables by means of fixed effects and time differencing, and propose an instrument to address residual endogeneity.

Omitted variables The importer-product-year and exporter-product-year fixed effects capture the changes in multilateral resistance terms. These control not just for the textbook multilateral resistance forces, such as time-varying importer- or exporter-product-specific demand or supply shocks, but also broad tariff changes by a country across a number of products simultaneously, and any aggregate effects of tariffs, such as trade-induced technology upgrading.

It has been recognized that unobserved bilateral taste or trade cost shifters are important for the variation in trade flows. If these shifters are correlated with tariffs, not accounting for them in estimation leads to omitted variable bias. For instance, if consumers in a particular importing country have idiosyncratically high taste for products from a particular exporter, the policymaker might set lower tariffs on those imports.[7] In their Handbook chapter Head and Mayer (2014, p. 162) recommend including bilateral fixed effects. Indeed, some papers in the literature control for bilateral unobservables via either bilateral fixed effects (see, among others, Lai and Trefler, 2002; Baier and Bergstrand, 2007; Donaldson, 2018), or time-differencing (e.g Feenstra, 1994; Head and Ries, 2001; Romalis, 2007; Imbs and Mejean, 2015). For this reason, our estimating equations are time-differenced, which removes all time-invariant importer-exporter-product determinants of bilateral product-level trade flows. After presenting the main results, Section 5.2 provides a detailed treatment of this point, and shows that controlling for bilateral unobservables is the key reason for the comparatively low elasticity estimates we report.

In addition, our baseline specifications include source-destination-product fixed effects, that absorb trends in product-specific impacts of bilateral resistance forces like distance, as well as trends in bilateral taste shocks for a product, that could be correlated with tariffs applied on the product.

Residual endogeneity in changes Despite fixed effects and differencing, an identification problem can still arise from time-varying, bilateral, non-tariff barriers $\Delta_h \ln \kappa_{i,j,p,t}$, or other time-varying, bilateral product-specific supply or demand shocks. In practice tariffs are set by governments which, in turn, are influenced by lobbyists, and subject to the WTO policy framework. There are three concerns with viewing applied tariff changes as exogenous. First, it is possible that a third factor in the importing country drives both tariff changes and changes in trade flows. A newly elected government, for instance, could change not only tariffs but also other policies that affect import demand. In a similar spirit, business cycle fluctuations could induce governments to change tariffs (Bown and Crowley, 2013; Lake and Linask, 2016). Again, imports would change in part because of the tariff change, and in part due to the changes in economic conditions. Further, a taste shock for a product from a specific source country could trigger both larger imports of the product and lower tariffs on that product due to lobbying. Second, there could be reverse causality, whereby the

[7] In the theoretical and quantitative trade literature, inclusion of taste shifters is ubiquitous; see the Handbook chapter by Costinot and Rodríguez-Clare (2014) for a typical representation of preferences.

importer's government changes tariffs because of observed or anticipated changes in trade patterns (e.g. Trefler, 1993). Third, foreign governments could influence the importer's government to change tariffs, either through the WTO body, or through other channels (Gawande, Krishna, and Robbins, 2006; Antràs and Padró i Miquel, 2011).

An instrument for tariff changes is difficult to find, as changes in trade policy are unlikely to ever be orthogonal to economic activity in general and trade flows in particular. We turn to the WTO's MFN tariff system to construct a plausibly exogenous instrument. All WTO member countries are bound by treaty to apply tariffs uniformly to all other WTO countries. Thus, when a WTO country changes its MFN tariffs, those tariffs change for all of its partners that trade on MFN terms. Of course, when a country changes its MFN rate on a product, it might do so due to concerns about imports from an important partner country, or lobbying by an important partner country. The baseline instrument uses the insight that third countries are also affected by this tariff change if they are MFN partners. From the point of view of these third countries, the tariff change is plausibly exogenous. The response of imports from these third countries can then identify the trade elasticity. As a control group we use countries unaffected by the MFN tariff change because they do not trade on MFN terms. These are countries in preferential trade agreements (PTAs).

Our baseline instrument is:

$$\begin{aligned}
\Delta_0 \ln \tau^{instr}_{i,j,p,t} \quad = \quad & \mathbf{1}\left(\tau_{i,j,p,t} = \tau^{\text{applied MFN}}_{i,j,p,t}\right) \times \mathbf{1}\left(\tau_{i,j,p,t-1} = \tau^{\text{applied MFN}}_{i,j,p,t-1}\right) \\
& \times \mathbf{1}\,(\text{not a major trading partner in } t-1 \text{ in aggregate}) \\
& \times \mathbf{1}\,(\text{not a major trading partner in } t-1 \text{ at product level}) \\
& \times \mathbf{1}\,(\text{not a major trading partner in } t \text{ in aggregate}) \\
& \times \mathbf{1}\,(\text{not a major trading partner in } t \text{ at product level}) \\
& \times \left[\ln \tau^{\text{applied MFN}}_{i,j,p,t} - \ln \tau^{\text{applied MFN}}_{i,j,p,t-1}\right].
\end{aligned} \tag{2.5}$$

We estimate equations (2.2) and (2.3) with $\Delta_0 \ln \tau^{instr}_{i,j,p,t}$ as an instrument for $\Delta_0 \ln \tau_{i,j,p,t}$ and equation (2.4) using $\Delta_0 \ln \tau^{instr}_{i,j,p,t}$ as the instrument for the h-year endogenous tariff change $\Delta_h \ln \tau_{i,j,p,t}$.

The first two indicators simply say that the applied MFN tariff is binding for the countries and product in question both in the pre-period $t-1$ and the impact period t. The next four indicators relate to whether the exporter is a major trading partner in $t-1$ or t, either in terms of aggregate trade, or in terms of trade in product p. At both the aggregate and the product levels, a trading partner is coded as major if its rank is in the top 10.[8] Finally $\ln \tau^{\text{applied MFN}}_{i,j,p,t} - \ln \tau^{\text{applied MFN}}_{i,j,p,t-1}$ is simply the log change in the tariff from $t-1$ to t.

[8] We also carried out the analysis considering the top 5 partners as major. The results were very similar.

Note that the instrument conditions on minor trading partners. We presume that endogeneity concerns that survive the fixed effects will mostly apply to the importer's major MFN trading partners. Thus, major MFN partners are dropped from the sample. We stress that the classification into major and minor trading partners is from the perspective of each individual importer. Section 5.1 shows that this filter does not produce a treated group composed of only small countries. This is because large countries are often minor trading partners from the perspective of individual importing countries.

Discussion To succinctly state the source of the identifying variation: we compare the changes in imports from countries hit by a plausibly exogenous tariff change to the changes in imports from countries to whom those tariff changes did not apply. The "treatment" countries experienced tariff changes because they are part of the MFN system. The "control" countries did not experience the MFN tariff changes because they trade on different terms. Note that the IV strategy is more than simply a sample restriction to minor MFN partners. Importantly, it constrains the identifying variation to MFN tariff changes. By doing so, the instrument sets up a comparison between treated and control observations, and is thus an "instrumented differences-in-differences."

Our approach thus follows the long tradition in the literature of estimating the trade elasticity based on the comparison of trade flows across product-country pairs subjected to differential tariff changes. It is well-understood that this strategy is correct and internally consistent in an environment with sector-level isoelastic gravity (see, e.g., the Handbook chapter by Head and Mayer, 2014), which characterizes the large majority of both empirical and theoretical work in trade. This environment allows for the "non-treated" trade flows to change following a tariff change in the treated group. That is, if an importing country raises its MFN tariffs, imports from non-MFN source countries can increase, as would be predicted by any CES aggregator. Our gravity-based approach will still correctly identify the elasticity, as long as importer-product multilateral resistance terms are used in estimation.

The estimation approach affects the interpretation of the coefficients. As in nearly all of the literature on elasticity estimation in international trade, our elasticities do not capture the general-equilibrium response of cross-border trade to a net fall in tariffs. Computing that response requires a general equilibrium model, for which our estimates can serve as an input. The aggregate response of trade to tariffs will then depend on the mechanisms operating in the model, including those that are absorbed by fixed effects in estimation. For example, if a fall in tariffs leads to technology adoption and higher productivity (Lileeva and Trefler, 2010; Bustos, 2011), exporter-product-time effects will absorb this channel, but a quantitative model can be used to capture its impact on aggregate trade.

Section 5.3 contains further discussion of threats to identification, alternative instruments, as well as extensive robustness checks.

2.4 Institutional background and examples of MFN tariff changes

To understand why countries change MFN tariffs, we provide some institutional background and discuss some examples.[9] When countries join the WTO, their accession treaty sets maximum MFN tariff rates ("bounds") that they can impose on imports from WTO member countries. These MFN bounds are country- and product-specific, and vary from very low rates for developed countries and large economies to much higher rates for developing countries. For instance, the average bound rate is 3.5% in the US, 10.0% in China, and 48.6% in India. The number of products covered by the bounds is also negotiated and varies by country. In many countries, including the US and China, 100% of products are covered by the bounds. By contrast, 74% of products are subject to MFN bounds in India, and 50% in Turkey. The bounds themselves vary substantially across products. In the US in 2015, about 40% of products had a bound of 0, while about one-tenth of products had bounds above 10%. Once these MFN bounds are set, they rarely change, except in subsequent rounds of WTO negotiations. As such, changes in MFN bounds do not provide sufficient variation for an instrument.

In practice, actual applied MFN tariffs are frequently far below the bounds. Thus, countries can and do legally vary their applied tariffs below the bounds. Some motives are business-cycle related. For instance Turkey raised a number of MFN tariffs temporarily around its financial crisis. The tariffs were lowered again post-crisis. Similar patterns were observed in Argentina. Sometimes the rationale for changing the MFN rates is less clear – India raises and lowers tariffs on varied products year-to-year. Finally, MFN rates might also be changed while countries are engaged in a trade war. China lowered MFN rates on 1449 consumer goods and 1585 industrial products while raising tariffs on the US as part of the US-China trade war in 2018. As a result, China's average tariffs on the US were 20.7% in late 2018, while those faced by other exporters to China were only 6.7%, on average. Since the US was the motivation for these MFN tariff changes, they are plausibly exogenous from the perspective of small exporters to China.[10]

This discussion makes clear the endogeneity of many tariff changes, and the rationale for the inclusion of a rich set of fixed effects (to remove business cycles and broad partner-specific variation). Further, the US-China trade war example illustrates the need to eliminate major partners from the instrument, in order to isolate the exogenous component of MFN tariff changes for third countries.

[9] Further details can be found in Bown and Crowley (2016). We are grateful to Chad Bown for useful suggestions and examples.

[10] See the blogpost by Bown, Jung and Zhang in June 2019 for a discussion.

3 Data and Basic Patterns

Our trade dataset is the BACI version of UN-COMTRADE, covering years 1995-2018. The data contain information on the trade partners, years, and product codes at the HS 6-digit level of disaggregation, as well as the value and quantity traded. We link these data to information on tariffs from the TRAINS dataset, also covering 1995-2018. This database reports the applied and the MFN tariff rates. The applied tariffs can differ from MFN tariffs for country pairs that are part of a PTA. Unfortunately, for many countries comprehensive information on tariff rates is often not available before they join the WTO. The sample covers 183 economies and over 5,000 HS6 categories.

We drop observations for which trade is subject to non-ad valorem (specific or compound) tariffs. For these tariffs TRAINS reports ad-valorem equivalents. However, computation of these equivalents requires data on quantities, which are often noisy and could also endogenously respond to changes in tariffs. Since the large majority of MFN tariffs are ad valorem, the impact of dropping these observations for our sample size is small.[11]

The most detailed product classification available in the trade data is at the HS6 level. However, we face the constraint that the data are provided in several different revisions of HS codes. Further, even within the same year, countries sometimes report trade flows in different vintages of HS codes.[12] While some concordances of HS6 codes over time are available, we do not implement these fully as they necessitate splitting values of trade across product codes in different revisions or aggregating product codes. As we do not observe transaction-level trade, any such split will introduce composition effects into our tariff measures. In particular, we could have spurious tariff changes coming from averaging tariffs when product codes are combined over time. Instead, our definition of a product is an HS6 code of a specific revision, tracked over time. We link product codes across revisions only when there is a one-to-one mapping between the codes across revisions. This approach is conservative, but it does reduce the effective sample size – and hence widens the standard errors – for any very long run elasticity estimates, as over a longer horizon there will be fewer product codes that map uniquely across revisions. Hence, the maximum horizon over which we estimate the trade elasticity in the baseline analysis is 10 years, which typically corresponds to only two transitions in HS revisions. Appendix Table A1 provides the fraction of codes that map uniquely across revisions. In a single revision transition, on average 89% of product codes have a unique mapping.[13] In a small number of instances, the meaning of HS4 codes changes across revisions, which would imply that importer-exporter-HS4 fixed effect categories combine substantively different products across time

[11] Among the 148 WTO members in 2013, the median fraction of HS6 products covered by non-ad valorem tariffs is 0.01%, and the mean fraction is 1.76% (World Trade Organization, 2014).

[12] As far as we are aware, there is no double counting of trade flows reported under different HS revisions.

[13] Naturally, alternative specifications that include several lags of tariff changes require longer horizons than ten years, reducing the sample size and increasing the standard errors of the estimates.

periods. We manually identified those instances and eliminated them.[14]

While HS6 product lines are often the most detailed level at which applied tariffs vary, a few countries have tariffs that vary within HS6 product groups (for instance at the HS8 or HS9 level). We do not have trade flows at a more detailed level, so we assess the robustness of our results to excluding series where countries apply different MFN tariffs within an HS6 product group.

The values of trade flows reported in these data are not inclusive of tariffs. Thus, the elasticities estimated by our procedure are tariff-exclusive, and must be appropriately adjusted to obtain the elasticity relevant from the consumer's perspective.[15]

Patterns in tariff changes Figure A1 plots histograms of tariff changes. It shows that while tariff decreases occur more frequently than increases, a substantial share of tariff increases exist in our data. Further, the treatment and the control group both experience a range of tariff changes. Note that our identification strategy does not require the control group to experience no tariff changes. Since our specifications include importer-product-time fixed effects, we exploit differential changes in MFN and non-MFN tariffs for identification. Below we also check the robustness of our estimates by removing from the control group observations in which non-MFN tariffs change. Figure A2 plots the autocorrelation function of tariff changes in our data. It highlights a negative first-order autocorrelation. This pattern motivates the use of time-series methods that explicitly account for the fact that impact tariff changes are not fully permanent, but partially reversed in subsequent periods.

Appendix A presents additional summary statistics about our sample: (i) the average share of imports by destination (Figure A3) and by product (Figure A4); and (ii) the incidence of MFN and non-MFN trade (Figure A5) in the sample.

Examples of the treatment/control assignments Appendix Table A2 provides an illustration of how the instrument is implemented. As our instrument is defined at the product level, we illustrate it for a 4-digit HS code 6403, "Footwear; with outer soles of rubber, plastics, leather or composition leather and uppers of leather." For three large importers (the USA, Japan, and Germany) in 2006, we list partner countries that fall into each of the following three categories: treatment group, control group, and excluded group.

Columns 1-2 list the 10 largest MFN trading partners at $t-1$ and t. Trading on MFN terms is the first criterion for being assigned to the treatment group. (Of course, there are many more than 10

[14]We have similarly implemented manual fixes for the very few HS6 codes that also change meaning over revisions.

[15]Section 6 contains the complete discussion. As an example, if the underlying model Armington, our long-run estimates would correspond to the elasticity in the CES aggregator $-\sigma$, while the trade elasticity inclusive of tariffs would be $1-\sigma$.

countries in this category). Columns 3-4 list the 10 major trade partners in terms of aggregate trade. These countries are disqualified from the treatment group. Columns 5-6 list the 10 major trading partners for the product code HS 6403. These are also disqualified from the treatment group. As expected, there is imperfect overlap between the set of major partners overall and in a specific HS code.

After these countries are dropped, columns 7-9 list the treatment, control, and excluded groups. As the table highlights, for the US NAFTA countries such as Canada and Mexico are important in the control group. The excluded group comprises large trading partners like Germany, China, and France, but also smaller economies such as Vietnam that are important exporters of footwear to the US. The treatment group includes smaller trading partners in footwear who trade at MFN rates, such as Portugal, Poland, Slovakia, and Hungary. While we do not incorporate explicit data on regional trade agreements, the instrument design appropriately assigns countries in customs unions or PTAs to control or excluded groups.[16] For Germany, for instance, EU member countries do not appear in the treatment groups, and are only part of the control groups.

4 Main Results

We begin by estimating the effects of a one-time tariff change on h-periods ahead trade flows and tariffs, as in equations (2.2)-(2.3), using our instrument as described above. For the baseline estimation, the product disaggregation for the fixed effects is at the HS4-level.[17] We also exclude major trading partners at the HS4-level in the baseline estimation. The left panel of Figure 1 reports the time path of tariff changes h periods after the initial 1-unit change. Thus, by construction the $h = 0$ coefficient is 1. A partial mean reversion in tariff changes is evident: following the initial impulse, about 80% of the change remains after 5 years, and approximately 75% after 10 years. At the same time, the pattern shown in the figure is contrary to the hypothesis that our low elasticity estimates may come from using temporary/short-term tariff changes as the source of variation. Figure 1 makes clear that the large majority of the initial tariff change persists for (at least) a decade.[18] These results illustrate the need for an estimation method that takes explicit account of the non-trivial

[16]The instrument might be improved if we could additionally incorporate information on PTAs. This would help in particular in assigning observations to the control group instead of the excluded group in some instances where the PTA rate is the same as the MFN rate and the country is a large trading partner. Currently, these observations have to be excluded. Unfortunately, while aggregate datasets on PTAs are available, these are typically not product-specific. Many free trade agreements exclude certain products, and applying them to all products is problematic for our estimation. Assigning observations to the excluded group increases our standard errors but is the conservative option.

[17]With many fixed effects, standard errors may be biased downward if there are many "singleton" observations that are perfectly absorbed by a fixed effect (Correia, 2015). The routine we use drops singleton observations from the sample prior to estimation, addressing this concern.

[18]In the robustness analysis, we show that the elasticity estimates are similar if we constrain the sample to only the "truly" permanent MFN tariff changes coming from the Uruguay round for estimation.

Figure 1: Local Projections: Tariffs and Trade

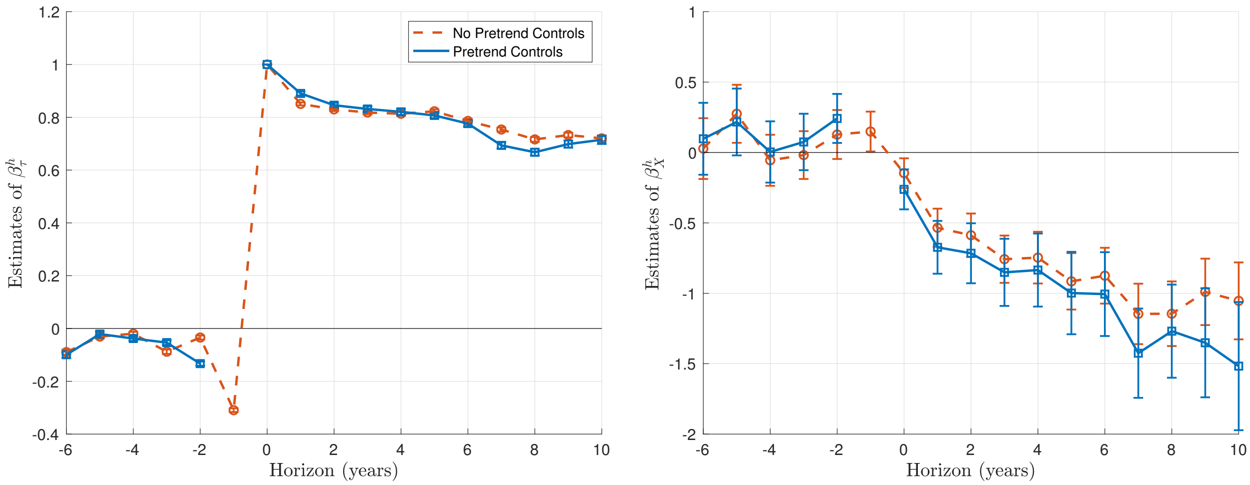

Tariffs　　　　Trade

Notes: This figure displays the results from estimating equations (2.2) and (2.3) – the local projection of h-period tariff growth (left panel) and h-period imports (right panel) on one period tariff growth, instrumented with our baseline instrument (2.5). The bars display 95% confidence intervals. Standard errors are clustered at the country-pair-product level.

time series behavior of tariffs.

The figure suggests a presence of a pre-trend. A tariff increase of one percent is preceded by a reduction of approximately 0.3 in the pre-period, reflecting again the negative first order autocorrelation highlighted above. We control for this pre-trend by including a lagged pre-trend control for both tariffs and trade in our baseline estimates throughout. The blue lines in Figure 1 depict the estimates after including the pre-trend controls. They make little difference to the results. We include additional lags in robustness checks.

The right panel of Figure 1 displays the impact of an initial one percent tariff change on trade flows. Trade falls gradually and stabilizes at 1 to 1.5% after 7 to 10 years. Unlike for tariffs, there is no evident pre-trend in trade flows, regardless of whether we use pre-trend controls, ruling out an important role for anticipation effects. Including the pretrend control modestly amplifies the point estimates of the effect of the tariff shock on trade values at longer horizons, though the difference is not significant. Columns 1 and 4 of Appendix Table A4 report the estimated impulse response coefficients and standard errors for tariffs and trade, respectively.

Figure 2 reports the baseline estimates of the trade elasticity ε^h across horizons. The impact ($h = 0$) elasticity is -0.26. Our data are annual, and it is unlikely that all tariff changes go into effect on January 1. Thus, we do not focus attention on the impact elasticity as it can be low due to partial-year effects. The point estimate in the year following the tariff change is probably a better indicator of the short-run elasticity. At $h = 1$, the elasticity is around -0.76. The 10-year elasticity is -2.12. Over the first 7 years, the elasticity converges smoothly to the long-run value, and then is stable for years 7-10. Appendix Table A5 reports the first stage F statistics, which indicate that the instrument is very strong.

Our short-run elasticities are similar to the low elasticities of trade flows to exchange rates typically found in the literature (e.g. Hooper, Johnson, and Marquez, 2000; Fitzgerald and Haller, 2018; Fontagné, Martin, and Orefice, 2018), and lend some support to the assumption often adopted in international business cycle literature that the Armington elasticity is below 1 (e.g. Heathcote and Perri, 2002).

The red line in Figure 2 reports the "all data/all tariffs 2SLS" estimates. This specification implements (2.4) on all the available data rather than the exogenous subset, instrumenting the horizon-h tariff change $\Delta_h \ln \tau_{i,j,p,t}$ with the initial tariff change $\Delta_0 \ln \tau_{i,j,p,t}$. As discussed in Section 2, using only the initial tariff change variation allows us to estimate the horizon-h trade elasticity. In contrast, estimation in long differences conflates trade elasticities of different horizons (see Appendix Proposition C.1).[19] At horizon 0, this approach amounts to a standard OLS estimation in differences. Note that because this strategy uses all tariff changes rather than the exogenous subset, it is subject to the concern that tariff changes are endogenous. Thus, the economic interpretation of these 2SLS horizon-h estimates should be in the spirit of OLS.

All data/all tariffs 2SLS actually produces a significantly smaller trade elasticity than our baseline IV at all horizons, a fact we revisit in Section 5.2. A substantive explanation for our baseline IV estimates being larger in absolute value than the all data/all tariffs 2SLS is that – conditional on all the fixed effects – tariffs are endogenously higher when imports are also high. One possible rationalization of this pattern is that greater import competition leads to more intense lobbying for protection. Trefler (1993) formalizes this argument, and shows that accounting for this type of endogeneity in US non-tariff barriers increases coefficient estimates of their impact of trade substantially.

Our estimates of β_X^h, β_τ^h, and ε^h should be interpreted as averages in the following sense. For a given size shock to tariffs, the subsequent evolution of tariff changes likely differs across shocks in our sample. Further, the responses of tariffs and trade could depend on the initial state of the world, they could vary by country pair, and/or depend on the product p for which the tariff changes. The estimation approach above will effectively average tariff and trade responses over all shocks, all evolution of tariffs, all initial states of the world, and all country-pairs and products. We now relax this assumption somewhat and report elasticities for broad product groups.

[19] Additionally, relying on higher frequency variation typically reduces confounding.

FIGURE 2: Trade Elasticity: Baseline and All Data/All Tariffs 2SLS

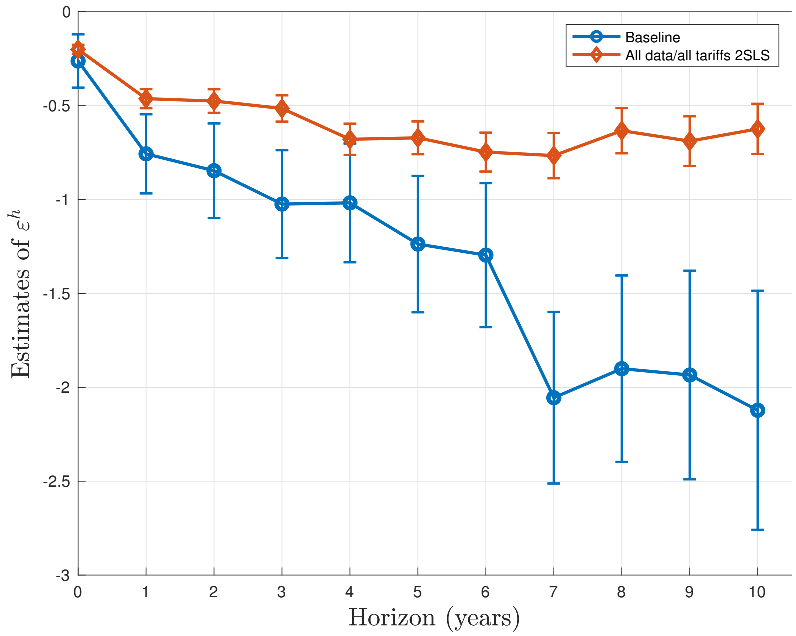

Notes: This figure displays estimates of the trade elasticity based on specification (2.4), and including one lag of the changes in tariffs and trade as pre-trend controls. 2SLS estimates with the baseline instrument (2.5) are in blue, and all data/all tariffs 2SLS estimates are in red. The bars display 95% confidence intervals. Standard errors are clustered at the bilateral country-pair-product level.

4.1 Sectoral heterogeneity

HS codes are organized into 21 sections that are consistent across countries. These sections describe broad categories of goods, such as "Live Animals, Animal Products" (Section 1). In practice, there is insufficient tariff variation in some of these sections to obtain precise estimates of the elasticity at all horizons. Thus, we combine a few of the sections together, leaving us with 11 sections. Appendix Table A3 describes the sections and lists the sections that are aggregated.

Figure 3 plots the point estimates of the trade elasticities over h for the 11 HS product groups. To contain the role of estimation error, we also report the median of the coefficients across horizons 7 to 10 for each product group in the figure. The long-run elasticities range from -0.75 to approximately -5 even in this coarse sectoral breakdown. The highest elasticities are in HS sections 8 (leather articles), 11 (textiles and apparel), whereas the least elastic sections are 18 (optical and precision instruments) and 20 (miscellaneous manufacturing). In addition, the elasticities fan out over time. The range at $h = 1$ is from -0.5 to about -1.5, much narrower than the long-run range.

One might be concerned that the headline elasticity values in the baseline analysis are unrepresenta-

FIGURE 3: Trade Elasticity: Sectoral Heterogeneity

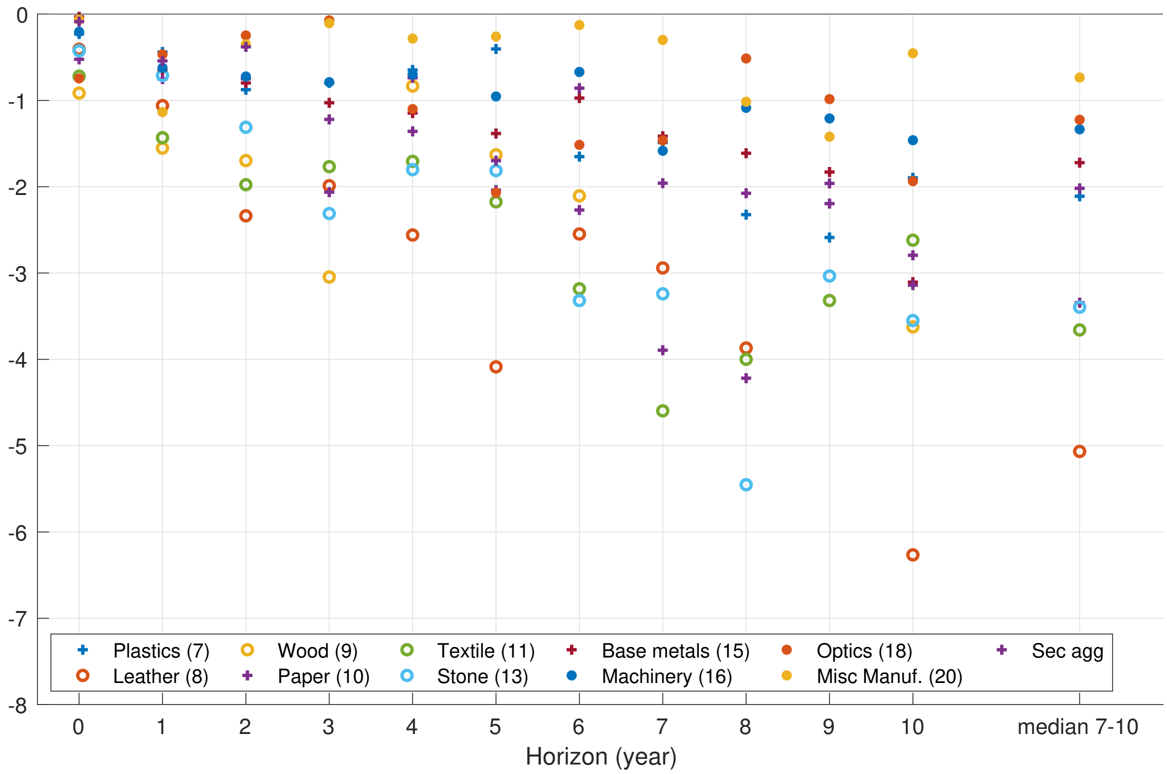

Notes: This figure displays the trade elasticity point estimates by HS Section based on specification (2.4) and using the baseline instrument (2.5). Some HS Sections are grouped into a single aggregate section "Sec agg" as described in the text.

tive of world trade, if product groups with higher or lower elasticities predominate in the data. Appendix Figure A8 plots the baseline horizon-specific elasticity, together with the world-trade-weighted mean and median of the sector-specific elasticities reported in Figure 2. The trade-weighted mean elasticities essentially coincide with the pooled baseline estimates, allaying compositional concerns. The trade-weighted median elasticities exhibit a similar time path but are, if anything, closer to zero.

5 Additional Results and Robustness

We have now presented the main estimation results of the paper. This section explores our findings in greater detail. In particular, it (i) shows that our identification strategy is based on broad variation across countries and sectors; (ii) discusses the relationship between our estimates and the conventional wisdom values in the literature, uncovering the source of the differences; (iii) reports a large battery of robustness checks. Together, these exercises demonstrate that our estimates are both quite stable and are not an artefact of non-representative data or non-standard estimation strategies.

5.1 Identifying variation

One might be concerned that the coefficient estimates are identified from special and/or non-representative segments of world trade. One possibility might be, for instance, that dropping major trading partners leaves a treated group composed of only small developing countries. Another possibility is that tariff changes might occur predominantly in products that account for relatively little of world trade. These potential concerns would be exacerbated by the large number of fixed effects, that further sweep out "singleton"-like observations, for instance cases in which the entirety of an importer-product trade is carried out on MFN basis.

To better understand the identifying variation in the data, we regress the one year (Δ_0) change in log trade flows and tariffs on the full set of fixed effects, and discard observations that are perfectly explained by the fixed effects. In this step we also impose the sample restriction that drops major trading partners. The resulting sample reflects the variation in trade flows and tariffs that is potentially available to identify the coefficients of interest. The observed patterns are reassuring on several fronts. The left panel of Figure 4 plots the (log) counts of instances countries appear as treatment or controls in the residualized data.[20] The relative size of the circle reflects country GDP. It is apparent that the same countries appear in both treatment and control groups, and indeed economies large in absolute size are frequently in the treatment group. The figure rules out the possibility that identifying variation comes from very small or esoteric countries. It also allays the concern that the control group countries are dramatically different from the treatment group.[21] Appendix Figure A6 projects the frequency of country appearance in treatment or control group on per-capita income. It is evident that a broad range of income levels is represented in both treatment and control groups.

The right panel of Figure 4 plots the distribution across HS sections. The green bars plot the shares of observations of all trade data. The orange bars display the shares of observations remaining in the residualized data after the fixed effects are taken out and sample restrictions imposed. The available variation is spread across all broad product groups, and is representative of the unconditional sectoral distribution of trade. The figure thus suggests that we are not identifying our elasticity coefficients from sectorally un-representative trade flows. Appendix Figure A7 plots the frequency of different product groups in our residualized data at a finer level of sectoral disaggregation (HS-2).

[20] The figure reports shares of observations rather than shares of value, because our regressions are unweighted.

[21] For instance, it does not appear to be the case that small countries that are recipients of GSP tariffs are disproportionately in the control group.

FIGURE 4: Country and Product Variation

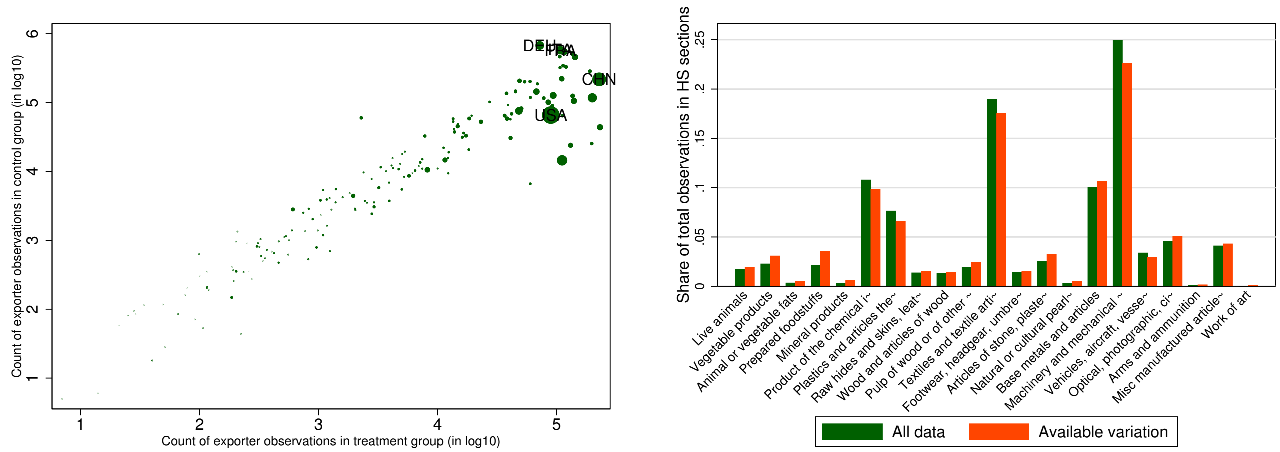

Countries HS Sections

Notes: The left panel displays the scatterplot of base 10 log counts an exporter appears in the control group on the vertical axis against the log count the same country appears in the treatment group on the horizontal axis. The size of the circle is proportional to relative country size as measured by GDP. The plot is based on a residualized sample from which importer-product-time, exporter-product-time, and importer-exporter-product fixed effects have been taken out, and the sample restrictions have been imposed. The right panel displays the sectoral distribution of all trade data in our sample (green bars), and the residualized sample after fixed effects have been taken out and the sample restrictions have been imposed (orange bars).

5.2 Relationship to other estimates

Our preferred IV estimates of the trade elasticity are -0.76 in the short run, falling to about -2 in the long run. These are substantially smaller in absolute value than the conventional wisdom range of -5 to -10 (see for instance the review in Anderson and van Wincoop, 2004). Interestingly, even our all data/all tariffs 2SLS estimates, which treat all tariff variation as exogenous as typical in the literature, are much smaller than the values commonly estimated in other studies.

Table 1 investigates the source of these differences. Columns 1-2 of the table estimate the elasticity using a log-levels OLS specification, assuming all tariff variation is exogenous. This specification, both without fixed effects and with the most commonly used multilateral resistance fixed effects (importer-product-time and exporter-product-time), yields values between -3.7 and -6.7, similar to previous estimates.

However, as argued above, failing to control for bilateral unobservables can bias the OLS estimates if these bilateral unobservables are correlated with tariffs. Columns 3-4 include country-pair-product fixed effects and time-difference the data by 5 years, respectively. In both cases, the elasticity estimates fall sharply to about -1. Thus, as recognized in several important contributions in the lit-

Table 1: Elasticity Estimates: Alternative Approaches

	(1)	(2)	(3)	(4)	(5)	(6)
	Log Levels			5-year Log Differences		
	No FE	Multilateral FE	Multilateral + Bilateral FE	OLS	Baseline IV	
$\ln \tau_{i,j,p,t}$	-3.70***	-6.96***	-1.04***	-0.66***	-1.17***	-1.11***
	(0.02)	(0.05)	(0.02)	(0.02)	(0.11)	(0.12)
R^2	0.013	0.383	0.571	0.180		
obs	107.1	106.2	104.9	38.20	21.4	21.13
First stage F					39917	36815
Fixed effects						
Imp×HS4×year, Exp×HS4×year		Yes	Yes	Yes	Yes	Yes
Imp×Exp×HS4			Yes			Yes

Notes: This table presents the results of estimating the trade elasticity in a traditional gravity specification. The dependent variable is log of trade values, and the independent variable is log of tariffs (columns 1-3), or 5-year differences (columns 4-6). Column 1 reports the results with no fixed effects. Column 2 adds importer-product-year and exporter-product-year fixed effects. Column 3 further adds importer-exporter-product fixed effects. Column 4 estimates the coefficient by OLS, while columns 5 and 6 present the results using our baseline IV. Column 6 includes our baseline bilateral-HS4 fixed effects. Standard errors clustered by country-pair-product are in parentheses. *** denotes significance at the 99% level. Numbers of observations are reported in millions. The differenced specifications do not have pretrend controls.

erature, controlling for unobserved bilateral determinants of trade is important for obtaining reliable estimates.

Columns 5 and 6 report IV estimates using our baseline instrument. The elasticity estimates of -1.17 and -1.11 at the five year horizon are very similar to one another, demonstrating that the importer-exporter-HS4 fixed effects do not play an important role in differenced specifications. Relative to the OLS estimates in column 4, and consistent with the main results above, instrumental variables push estimates away from 0.

Sample composition The relatively low headline elasticity values we report are not due to any potential lack of representativeness of our baseline sample. The log-levels specifications in columns 1-3 have quite similar sample sizes, all three are estimated essentially on all of the world's trade. Thus, the large drop in the elasticity in column 3 is not due to changes in sample composition. The log-differenced specification in column 4 has fewer observations, as it requires non-zero trade flows in both beginning and end periods. The IV estimates reported in columns 5 and 6 have even fewer observations, as the sample is constrained to minor exporters. The patterns in the coefficients are nonetheless not driven by changes in the sample. Appendix Table A6 replicates Table 1 on a sample that is constant across columns. If anything, the difference between the cross-sectional estimates in columns 1-2 and the fixed effects/differenced specifications in columns 3-4 is even starker, as the cross-sectional variation implies even larger elasticity estimates in this subsample (as high as $8-11$).[22] Section 5.3 addresses in detail a related phenomenon, namely the extensive margin of trade.

Fixed effects and omitted variables bias Another concern might be that "overcontrolling" for bilateral product-level determinants of trade by means of either high-dimensional fixed effects or differencing may remove "too much" of the variation available for the purposes of elasticity estimation. Table 2 explores whether the estimated coefficients fall because the fixed effects are absorbing an excessive amount of variation in the data. Column 1 reproduces the traditional gravity specification without any bilateral fixed effects from Table 1. Columns 2-5 add bilateral fixed effects in increasing order of resolution, starting from the coarsest possible (country-pair, as suggested by Head and Mayer, 2014), up through the country-pair-HS4 fixed effects.

The takeaway from the table is that by far the largest fall in the estimated elasticity comes from adding the very coarsest bilateral fixed effect at the country-pair level (column 1 versus column 2). There are comparatively few of these (around 30 thousand relative to a sample of 106 million). Thus the concern about oversaturating the data with too many fixed effects is the least applicable for these fixed effects. And yet, it is these country-pair effects that lower the elasticity estimates the most. While by no means a formal proof, Table 2 suggests that controlling for omitted variables via either fixed effects or differencing is very much worthwhile in spite of overcontrolling concerns.[23]

Finally, our baseline estimation also includes country-pair-HS4 fixed effects, that in a differenced specification absorb *trends* in these bilateral shifters, rather than levels. We favor including these because we found that the results were slightly more stable across specifications under this approach. These fixed effects are not the reason for the low elasticity estimates. Appendix Figure A9 plots the time path of trade elasticities with multiple versions of bilateral fixed effects including without any bilateral fixed effects. There are at most modest, and not statistically significant differences across these specifications. Omitting bilateral fixed effects actually produces a somewhat lower trade elasticity.

Measurement error A related concern is that differencing may exacerbate measurement error on the right-hand side, biasing OLS coefficients towards zero. While differencing removes a country-pair-HS6 fixed effect, amplification of residual measurement error due to differencing is not the reason behind the low elasticity estimates. As Table 2 makes clear, the coarse country-pair effects in a levels specification is sufficient to sharply lower the coefficient estimates.

[22]Table A6 also reports the estimates when differencing over 3 and 7 years, and shows that there is nothing special about 5-year differences.

[23]Appendix Table A7 replicates Table 2 with multilateral resistance terms at the HS6 rather than HS4 level. The results are virtually identical.

Table 2: Log-Level Elasticity Estimates Varying the Fineness of the Importer-Exporter Effects

	(1) No Bilateral	(2) Country-pair	(3) Country-pair ×HS2	(4) Country-pair ×HS3	(5) Country-pair ×HS4
$\ln \tau_{i,j,p,t}$	-6.96***	-1.39***	-1.21***	-1.17***	-1.04***
	(0.05)	(0.02)	(0.02)	(0.02)	(0.02)
R^2	0.383	0.476	0.507	0.517	0.571
obs	106.2	106.2	106.1	105.9	104.9
Fixed effects					
Imp×HS4×year, Exp×HS4×year	Yes	Yes	Yes	Yes	Yes
Imp×Exp		Yes			
Imp×Exp×HS2			Yes		
Imp×Exp×HS3				Yes	
Imp×Exp×HS4					Yes

Notes: This table presents the results of estimating the trade elasticity in log levels in a traditional gravity specification. The dependent variable is log of trade values, and the independent variable is log of tariffs. Column 1 reports the results with importer-product-year and exporter-product-year fixed effects. Columns 2-5 additionally include progressively finer bilateral fixed effects, from importer×exporter to importer×exporter×HS4. Standard errors clustered by country-pair-product are in parentheses. *** denotes significance at the 99% level. Numbers of observations are reported in millions.

More broadly, we argue that right-hand side measurement error should not be an overriding concern here for three reasons. First, our right-hand side variable is tariffs, which are statutory policy instruments less likely to be measured with error. Second, we have done extensive checks on the tariff data, and eliminated known issues such as specific or compound tariffs, and product reclassifications. A few countries set tariffs at the HS8 or HS9 level, rather than the HS6 level of our data. By constraining the sample to instances of zero standard deviation in tariffs within an importer-HS6-year, we can eliminate cases in which tariffs are set at the 8-digit or 9-digit level, as well as any other cases in which the tariff is not exactly ad valorem at the HS6 level. Doing so barely changes the estimates (see Table 4 below). Third, the solution to measurement error on the right hand side is to use an instrument, which we employ in our baseline approach. This will help if any measurement error in the instrument is not correlated with any residual measurement error in the tariff data overall. The one-year initial MFN tariff changes that form the basis of the instrument are broad, published, changes affecting the applied tariffs to several countries, and least likely to be measured with error. Finally, even if measurement error remains an important worry, it must still be traded off against an equally important and well-recognized concern about omitted variables, as discussed in detail above.

In summary, country-pair specific variation is key for the estimation of trade elasticities. Estimates are large if this variation is used, and small if this variation is taken out with importer-exporter(-product) fixed effects or through differencing. Relative to the all data/all tariffs 2SLS estimation that encompasses endogenous variation in tariffs, our baseline instrument raises estimates of the trade elasticity.

5.3 Robustness

Pre-trends and anticipation effects Tariff decreases often follow tariff increases (tariff changes are negatively autocorrelated), as shown above. Indeed, the left panel of Figure 1 reveals some evidence of a pre-trend in tariffs. We account for differential pre-trends in tariffs using the standard approach of controlling for lagged tariff and trade changes. Our baseline estimates use a single lag of both as pretrend controls. Columns 2-3 of Table 3 report results with no lags and 5 lags, respectively, to compare the results to the baseline in column 1. The substantive conclusions change little when adding or subtracting lags, although with more lags the sample size drops substantially and the standard errors increase. Columns 1-3 and 4-6 of Table A4 reports the results of local projections of tariffs and trade flows directly on the initial tariff change, as in (2.2)-(2.3), while allowing for 1, 0 and 5 lags. Once again, the point estimates change little when adding lags.

A distinct concern is anticipation effects. Even if pre-treatment tariffs are constant, countries might begin to adjust their exports in response to an expected future MFN tariff change by the importer. We check for the presence of such anticipation effects by examining pre-trends in the trade volume equation estimates. Figure 1 shows no evidence of pre-trends in trade values even without controlling for lagged changes in tariffs and trade.

Alternative samples and standard errors Column 4 of Table 3 restricts the sample so that each fixed effect is estimated from at least 50 observations. Column 5 two-way clusters the standard errors by importer-exporter-HS4 and year. In both cases the estimates and their precision change little. Column 6 reports estimates on a constant sample. While the point estimates are slightly lower in absolute value, the standard errors widen substantially. Overall, the difference from the other specifications is typically not statistically significant. This is reassuring as the constant sample conditions on positive trade flows for all time horizons. This sample likely has different characteristics than the full sample, but the stability of the estimates suggests that sample selection is not a big concern. Finally column 7 reports the results from an estimation where we drop observations from the control group that experience tariff changes. The estimates are slightly lower than the baseline, but not significantly so at most horizons.

Our estimated tariff impulse responses stabilize fast and are very persistent, with about 75% of

Table 3: Trade Elasticity, Robustness: Pre-Trends, Alternative Clustering, Alternative Samples

	Baseline	Zero Lag	Five Lags	FE50	Two-way Clustering	Constant Sample	Alternative Control Group
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
t	-0.262***	-0.147***	0.166	-0.232**	-0.262***	-0.592**	-0.187**
	(0.072)	(0.054)	(0.138)	(0.094)	(0.096)	(0.290)	(0.077)
obs	31.66	41.46	14.58	17.62	31.66	5.04	27.33
$t+1$	-0.756***	-0.628***	-0.129	-0.599***	-0.756***	-0.098	-0.490***
	(0.108)	(0.081)	(0.208)	(0.133)	(0.141)	(0.380)	(0.117)
obs	26.18	32.85	12.52	15.19	26.18	5.04	22.63
$t+3$	-1.024***	-0.926***	-0.625**	-0.865***	-1.024***	-0.895**	-0.743***
	(0.146)	(0.105)	(0.315)	(0.173)	(0.195)	(0.472)	(0.160)
obs	20.8	26.16	9.76	12.48	20.8	5.04	17.86
$t+5$	-1.237***	-1.112***	-1.146***	-1.012***	-1.237***	-0.916**	-0.792***
	(0.185)	(0.124)	(0.429)	(0.215)	(0.253)	(0.437)	(0.201)
obs	16.69	21.13	7.3	10.22	16.69	5.04	14.27
$t+7$	-2.055***	-1.521***	-2.330***	-1.853***	-2.055***	-0.990**	-1.383***
	(0.233)	(0.145)	(0.595)	(0.270)	(0.357)	(0.489)	(0.251)
obs	13.22	16.95	5.25	8.2	13.22	5.04	11.12
$t+10$	-2.122***	-1.463***	-2.550**	-1.760***	-2.122***	-1.818***	-1.600***
	(0.325)	(0.194)	(1.016)	(0.374)	(0.332)	(0.544)	(0.379)
obs	8.31	11.25	3.21	5.25	8.31	5.04	6.84

Notes: This table presents robustness exercises for the results from estimating equation (2.4). All specifications include importer-HS4-year, exporter-HS4-year and importer-exporter-HS4 fixed effects, and the baseline pretrend controls (one lag) unless otherwise specified. Columns 2 and 3 vary the pretrend controls (including alternatively zero lags or five lags of import growth and tariff changes). Column 4 reports the results when the sample is restricted to fixed-effects clusters with a minimum of 50 observations. Standard errors are clustered at the importer-exporter-HS4 level, except in Column 5 where they are additionally clustered by year. Column 6 restricts the sample to a constant sample. Column 7 reports results where the control group only contains observations with zero tariff changes. ***, ** and * indicate significance at the 99, 95 and 90 percent level respectively. Observations are reported in millions.

the initial shock surviving 10 years.[24] This alleviates concerns that our estimates are driven by very short-run temporary MFN tariff changes. To further explore the impact of potentially more permanent tariff changes, we estimate elasticities using only the tariff changes of the Uruguay Round GATT/WTO negotiations. It is likely that firms viewed these as persistent or permanent—at least until the next successful multilateral negotiation. In practice, we constrain the sample to only MFN tariff changes during 1995-1997, which corresponds to the staggered phasing in of the Uruguay round MFN bounds. Reassuringly, we find all data/all tariffs 2SLS estimates that are not significantly different from our baseline IV coefficients (Appendix Table A10). This may suggest that the Uruguay Round tariff changes were more "exogenous" than typical tariff changes, since they resulted from protracted multilateral negotiations. Estimates using our baseline IV on the 1995-1997 sample are imprecise and not informative, as the sample size is drastically reduced.

Alternative instruments, outcome variables, and fixed effects The baseline instrument excludes large trading partners from both treatment and the control groups. Column 2 of Table 4 reports the results when we include all trading partners but restrict the variation to MFN tariff changes ("All data/MFN tariffs 2SLS"). In this case the instrument is simply the change in the MFN tariff rate for all countries subject to the MFN tariff rate. The point estimates fall to about -0.9 for the long-run elasticity. Columns 3 and 4 report results for quantities and unit values, respectively. It turns out that the impact in the long run is mostly on quantities. The response of unit values is noisy and in general insignificant. For interpreting the unit values coefficients, it is important to keep in mind that these are unit values exclusive of tariffs. Thus, a zero estimated coefficient on unit values indicates complete pass-through of tariff changes to the buyers in the importing country. Appendix Table A10 contains results for the elasticity estimated with the multilateral resistance terms at the HS6 level. The estimates are somewhat smaller than the baseline, though the sample shrinks and the standard errors widen.[25]

Extensive margin Our baseline specifications are in log differences and our data are at the country-pair product level. Thus, our sample consists of instances where country-pair product flows are positive in both the initial and end periods. Many trade models emphasize exit and entry of firms into export markets (see, e.g. Melitz, 2003; Ruhl, 2008; Alessandria and Choi, 2014; Alessandria, Choi, and Ruhl, 2014; Ruhl and Willis, 2017). The firm-level entry and exit in country-pair-product markets with positive trade is already reflected in our baseline elasticity estimates.[26]

[24] Consistent with our estimates, Bown and Crowley (2014) document that most MFN tariff changes below bounds are permanent or very persistent.

[25] Note that in our baseline estimation, time differencing already eliminates importer-exporter-HS6 fixed effect in levels.

[26] While we cannot examine the firm-level extensive margin using our data, available empirical evidence often suggests that it is not large quantitatively. For example, Buono and Lalanne (2012) analyze the response of French firm-level exports to the Uruguay round tariff reductions, and conclude that extensive margin responses did not materially contribute to the overall trade volume changes. Fitzgerald and Haller (2018) estimate that in Ireland, the contribution

Our baseline estimation abstracts from the possibility that tariff changes lead to (dis)appearance of trade flows at the country-pair-product level. As a benchmark for how important the product-level extensive margin can be, Kehoe and Ruhl (2013) report that it contributed only 10% of the overall growth in North-American trade following NAFTA implementation. While instances of rapid economic growth and structural change – such as South Korea – can be associated with a contribution of the product-level extensive margin as high as 25%, the extensive margin plays a negligible role in trade growth under more conventional circumstances (such as US-UK).

To implement specifications with the product-level extensive margin, we use the differenced inverse hyperbolic sine transformation instead of log differences as suggested by Burbidge, Magee, and Robb (1988). This transformation allows us to include zero or missing trade flows, while approximating logs for larger values of the data.[27,28]

We stress that including zero trade observations in the sample need not increase the trade elasticity point estimates. How the point estimates change relative to the baseline depends on the relative importance of observations where trade switches from, say, zero to positive, compared to observations where trade goes from zero to zero. If a tariff falls and many zero trade observations turn positive, the elasticity will be pushed up. However, if following a tariff reduction many zero observations stay at zero, the elasticity estimate will be pushed down, since, on average, trade changes become less responsive to tariff changes.

As a result, elasticity estimates that incorporate the extensive margin are sensitive to *which* zeros are added to the sample. We report two sets of estimates. In the first, we include all available zero trade observations for exporter-HS section to any importer in instances where some exports are ever observed.[29] In the second, we only include observations where trade goes from zero to positive, or from positive to zero. This approach gives the extensive margin maximum chance to increase the absolute values of elasticity estimates, in the sense that it only admits observations for

of the firm-level extensive margin to the long-run elasticity of trade to tariffs is less than 10%.

[27] Tariff data are typically not missing and we can always construct $\ln \tau_{i,j,p,t}$, so we do not need the inverse hyperbolic transformation for tariffs. Bellemare and Wichman (2020) highlight that caution must be used in interpreting the estimated coefficient as an elasticity, but in our case the estimated β^h can be interpreted as an elasticity. The estimated coefficient converges to an elasticity as the underlying variable being transformed (trade values in our case) takes on large enough values on average. This is the case in the trade data.

[28] A common approach to include variation along the extensive margin is to use Poisson pseudo maximum likelihood (PPML) estimation. This method can neither easily be run in differences, nor with an instrument. Further, available software for PPML estimation (in levels) with high-dimensional fixed effects fails to converge even on subsets of our sample that include 5% of the observations.

[29] That is, if country A ever exports any product in HS 1-digit section Z to importer B in any year, all the zero exports of products belonging to section Z from A to B in every year are added to the sample. This leaves out of the estimation sample export flows between pairs of countries in broad sectors that never occurred, and thus are unlikely to respond to tariff changes. A more extreme approach is to just include all the possible zeros. Predictably, this leads to even lower elasticity point estimates, as it increases the fraction of the sample in which trade flows go from zero to zero. Note that fixed effects will automatically absorb instances in which there is never any trade within a fixed effect category, and those observations will not contribute to elasticity estimates.

which extensive margin changes actually occur. This sample restriction corresponds more closely to quantitative models and firm-level analyses where the extensive margin is active. However, it should interpreted as an upper bound on the sensitivity of trade flows to tariffs as it effectively selects the sample based on outcomes. All extensive margin estimates do not include pretrend controls. Therefore the results in this exercise must be compared to the baseline estimates without pretrend controls (Column 2 of Table 3).[30]

The resulting estimates in columns 5 and 6 of Table 4 can be interpreted as the total elasticity, inclusive of both the intensive and product-level extensive margins. When including more zeros (column 5), the point estimates are similar to the baseline initially, smaller in the long run. We conjecture that this is because the estimation sample now includes many instances of trade being zero at both $t-1$ and $t+h$. Since these appear as zero changes in the sample, they drive down the point estimate. Column 6 reports the extensive margin response when we only include zeros in instances where trade goes from zero to positive, or from positive to zero. As expected, the 10 year elasticity including the extensive margin is slightly higher (-1.64) than the corresponding intensive margin instrumented specification without pretrend controls (-1.46).

Additional results, diagnostics, and robustness Column 7 of Table 4 estimates a distributed-lag model as an alternative to the local projection specification. This approach has two disadvantages relative to the baseline: (i) it requires a panel of non-missing log growth rates for trade, tariffs, and the instrument for every lag, reducing the estimation sample greatly; and (ii) it imposes linearity on the estimates. Caveats aside, the distributed lag specification with 10 lags yields a long-run trade elasticity of 3.17 with a standard error of 1.25, while the number of observations falls to just around 6.08 million. This point estimate is statistically indistinguishable from our baseline estimates.[31] Finally, the last column of Table 4 estimates the elasticity on a sample where tariffs do not vary within an importer-HS6. This drops importer-product instances where tariffs are set at finer levels

[30]Including pre-trend controls leads to elasticity estimates much lower in absolute value, and below the baseline (intensive margin) estimates. This appears to be due to the fact that adding zero observations adds to the sample many instances of occasional exporting, where entry is followed by exit and vice versa. As a result, the pre-trend control for lagged log change in trade has a negative sign and is a very powerful predictor of the subsequent change in trade (t-statistic of about 2000). If this part of the sample is dominated by idiosyncratic shocks that manifest themselves in occasional exporting behavior, there would be less for tariff changes to explain. Reporting extensive margin estimates without pre-trend controls thus gives the extensive margin maximum chance to produce larger elasticities relative to the baseline.

[31]Formally, we estimate the equation $\Delta_0 \ln X_{i,j,p,t} = \sum_{k=0}^{10} \gamma^k \Delta_0 \ln \tau_{i,j,p,t-k} + \delta_{i,p,t} + \delta_{j,p,t} + \delta_{i,j,p} + u_{i,j,p,t}$ instrumenting $\Delta_0 \ln \tau_{i,j,p,t-k}$ with $\Delta_0 \ln \tau^{instr}_{i,j,p,t-k}$, for all k. The trade elasticity at horizon h reported in Table 4 is then the estimate of $\sum_{k=0}^{h} \gamma^k$. As this estimation requires 11 instruments for 11 endogenous variables, we report the Sanderson-Windmeijer F-statistic for weak instruments in Appendix Table A5. Conceptually, there is a subtle difference between the object estimated by local projections and the distributed lag approach. Whereas the local projections take into account the time series behavior of the tariff variable, the distributed lag coefficients cumulated up to horizon h are estimates of the response of trade to a permanent once-and-for-all change in tariffs that happened at horizon 0. This distinction does not matter for the long-run limit, but is relevant for finite h.

of disaggregation, such as HS8 or HS10. Again, the results are very similar to the baseline at all horizons.[32]

Appendix B presents the results for all the specifications at every horizon. This appendix also reports the first stage F-statistics for the baseline specification for every horizon. In all cases, the first stage F-statistics are much higher than 10.

Other candidates for instruments There are other candidate instruments that could in principle be considered under the WTO framework. Here, we discuss these potential instruments and issues with each of them.

A natural candidate instrument is WTO accession. When a country such as China joins the WTO, the negotiations are protracted, and there are substantial anticipation effects (see for instance Pierce and Schott, 2016). However, once China joins the WTO and sets its MFN tariffs, small third countries in the WTO are also affected by these MFN tariffs. These countries are plausibly facing an exogenous change, conditional on the anticipation effects, as they were likely not key players in the negotiations. While there are a few WTO accessions in our data, a key problem with implementing this instrument is that product-level tariff data are typically not available in standard datasets for countries before they join the WTO. It is therefore not possible to construct tariff changes (the change from the pre-WTO rate to the MFN rate) for estimation.

A second instrument would be a change in the MFN bound, which is the maximum tariff a country in the WTO can apply against other countries. While these are likely less discretionary, the MFN bounds are set in the WTO accession treaty and very hard to change ex-post. The lack of instances of changes in the bounds implies there is insufficient variation in this potential instrument to estimate the elasticity.

6 Theory and Applications

We stress that equations (2.2), (2.3), and (2.4) are "model free", and under our identification assumptions will produce estimates of ε^h by definition. The mapping between these estimates and parameters in theoretical models then depends on model structure. This section provides a mapping to dynamic and static trade models. We first develop a simple dynamic model of sluggish adjustment to trade cost shocks. The recent literature on trade dynamics is rich in both substantive mechanisms and quantification (see, among many others, Costantini and Melitz, 2007; Ruhl, 2008; Burstein and

[32] We have also checked whether the trade response depends on the size of the tariff shock. To do so, we estimated separate elasticities depending on whether the absolute value of the initial (nonzero) tariff change is below or above the median nonzero absolute value tariff change. The estimated elasticities for both size categories are very similar and we do not report them here.

Table 4: Trade Elasticity, Robustness: Alternative Instruments, Outcomes, Fixed Effects, Samples and Models

	Baseline (1)	All data/MFN tariffs (2)	Quantities (3)	Unit Values (4)	Extensive (5)	Extensive Sel (6)	DL (7)	SD1 (8)
t	-0.262***	-0.275***	-0.179*	-0.049	0.025	-0.082	-0.413	-0.313***
	(0.072)	(0.030)	(0.093)	(0.057)	(0.039)	(0.063)	(0.338)	(0.096)
obs	31.66	57.14	31.66	31.66	131.03	56.35	6.08	28.7
$t+1$	-0.756***	-0.624***	-0.664***	-0.027	-0.484***	-0.805***	-0.523	-0.889***
	(0.108)	(0.047)	(0.136)	(0.080)	(0.063)	(0.089)	(0.466)	(0.142)
obs	26.18	47.19	26.18	26.18	108.13	49.11	6.08	23.76
$t+3$	-1.024***	-0.648***	-0.810***	-0.134	-0.590***	-0.955***	-1.578**	-1.228***
	(0.146)	(0.061)	(0.184)	(0.104)	(0.084)	(0.107)	(0.685)	(0.200)
obs	20.8	38.16	20.8	20.8	87.22	41.5	6.08	18.86
$t+5$	-1.237***	-0.718***	-1.422***	0.291**	-0.730***	-1.183***	-2.097**	-1.180***
	(0.185)	(0.073)	(0.232)	(0.128)	(0.100)	(0.122)	(0.861)	(0.253)
obs	16.69	30.89	16.69	16.69	69.73	34.61	6.08	15.12
$t+7$	-2.055***	-0.940***	-2.166***	0.162	-0.902***	-1.496***	-2.710***	-1.990***
	(0.233)	(0.091)	(0.292)	(0.161)	(0.114)	(0.138)	(1.016)	(0.318)
obs	13.22	24.64	13.22	13.22	55.24	28.4	6.08	11.96
$t+10$	-2.122***	-0.866***	-1.765***	-0.078	-0.941***	-1.638***	-3.166**	-2.360***
	(0.325)	(0.122)	(0.406)	(0.221)	(0.153)	(0.181)	(1.252)	(0.444)
obs	8.31	15.92	8.31	8.31	35.09	19.24	6.08	7.52

Notes: This table presents alternative estimates for the results from estimating equation (2.4), varying the instrument or outcome variable. Column 2 uses an alternative definition of the instrument where all trade partners subject to the MFN regime are included. Column 3 reports results for quantities, and Column 4 the results for unit values. Column 5 presents results including the extensive margin using the inverse hyperbolic sine transformation for trade flows, and all zero trade observations for importer-exporter-section pair with ever positive trade. Column 6 presents results including the extensive margin using the inverse hyperbolic sine transformation for trade flows, and only zero trade observations when trade switches from zero to positive, or vice versa. Column 7 presents results from a distributed lag model. Column 8 reports the results based on a sample where tariffs do not vary within an importer-HS6. All specifications include importer-HS4-year, exporter-HS4-year and importer-exporter-HS4 fixed effects and the baseline pretrend controls (one lag). Standard errors are clustered at the importer-exporter-HS4 level. ***, ** and * indicate significance at the 99, 95 and 90 percent level respectively. Observations are reported in millions.

Melitz, 2013; Drozd and Nosal, 2012; Alessandria and Choi, 2014; Alessandria, Choi, and Ruhl, 2014; Ruhl and Willis, 2017; Blaum, 2019; Fitzgerald, Haller, and Yedid-Levi, 2019; Alessandria, Arkolakis, and Ruhl, 2020). The goal of this section is not to revisit all of the proposed mechanisms for gradual adjustment of trade. Rather, we focus on the minimal common structure that characterizes these models. Appendix C lays out the model details and proves the propositions in this section.

An attractive feature of our model is that it delivers analytical expressions for trade elasticities at all horizons that clarify the determinants of the adjustment dynamics. In this setting, we state the short- and long-run model-implied elasticities and the properties of their time path. We also show that this framework delivers the estimating equations used above up to a first order approximation. We then explore a simple quantification exercise that matches the time path of the trade elasticities following a tariff shock. Finally, turning to the mapping from our estimates to the parameter relevant for static trade models, we explore the quantitative implications of our estimates for the long run gains from trade.

6.1 Dynamics of Trade Elasticities

Setup The minimalist model that can capture differing trade elasticities in the short vs. the long run has to feature a variable that determines trade flows but cannot instantaneously and fully adjust upon a change in trade costs. In addition, a long and smooth path of increasing trade elasticities requires some curvature in the costs of adjustments, such that the long run is not reached in the first period after the shock. Following a long tradition in the literature, we assume that foreign markets are served by monopolistically-competitive firms that face CES demand. We focus on the partial equilibrium decisions of firms from one market selling to another, and thus suppress importer, exporter, and product subscripts. Consistent with the gravity tradition, general equilibrium objects such as domestic unit costs or foreign demand shifts are absorbed by country-product-time fixed effects, and thus we ignore general equilibrium forces in this section. Throughout, we assume that marginal costs are constant at the firm level and thus exporting decisions are separable across locations. The setup below nests versions of the Krugman (1980), Melitz (2003), and Arkolakis (2010) models, as well as extensions with pricing to market (e.g. Burstein, Neves, and Rebelo, 2003; Atkeson and Burstein, 2008).

Trade between the two countries can be expressed as

$$X_t = p_t^x q_t n_t,$$

where n_t is a generic mass, p_t^x is the exporters' price exclusive of tariffs, and q_t is the quantity exported per unit mass. Crucially for the short vs. long-run distinction, we assume that p_t^x and q_t adjust instantaneously to tariff changes, whereas n_t is pre-determined by one period, and can only change from the next period onwards. Quantity and price are a function of tariffs, and quantity must be consistent with market clearing at the price: $p_t^x = p^x(\tau_t)$ and $q_t = q(p_t^x, \tau_t)$. Exporting generates flow profits $\pi(\tau_t) = q_t(p_t^x - c)$, where c is the unit cost. Define the following elasticities:

$$\eta_{q,p} := \frac{\partial \ln q}{\partial \ln p^x}, \quad \eta_{q,\tau} := \frac{\partial \ln q}{\partial \ln \tau}, \quad \eta_{p,\tau} := \frac{\partial \ln p^x}{\partial \ln \tau}, \quad \eta_{\pi,\tau} := \frac{\partial \ln \pi}{\partial \ln \tau}, \tag{6.1}$$

where we assume that $\eta_{q,p} < 0$, $\eta_{q,\tau} < 0$, and $\eta_{\pi,\tau} < 0$.

The measure n_t comes from profit-maximizing agents serving the export market. Let r denote the real interest rate at which firms discount future profits, and G a positive and increasing function. Dynamics in this model are governed by two equations:

$$v_t = \frac{1}{1+r}\mathbb{E}_t\left[\pi_{t+1} + (1-\delta)\, v_{t+1}\right], \tag{6.2}$$

$$n_t = n_{t-1}(1-\delta) + G(v_{t-1}), \tag{6.3}$$

subject to the transversality condition $\lim_{t\to\infty}\left(\frac{1-\delta}{1+r}\right)^t v_t = 0$. The forward-looking equation (6.2) states that the value of exporting v_t is the expected present value of future flow profits from exporting. The backward-looking equation (6.3) describes how the mass n_t evolves. The increment to the mass n_t today $G(v_{t-1})$ is a function of the value of exporting last period, when the entry or investment decision was made. Parameter δ is a rate of depreciation or an exogenous exit rate.

The model's tractability stems from the fact that equations (6.2) and (6.3) can be solved sequentially. For any stochastic process for tariffs $\{\tau_t\}_{t=0}^{\infty}$, equation (6.2) can be solved forward to obtain

$$v_t = \frac{1}{1+r}\mathbb{E}_t\left[\sum_{k=0}^{\infty}\left(\frac{1-\delta}{1+r}\right)^k \pi(\tau_{t+k+1})\right]. \tag{6.4}$$

Importantly the value v_t does not depend on the evolution of n_t. The resulting sequence $\{v_t\}_{t=0}^{\infty}$, can then be used to obtain n_t after solving equation (6.3) backwards,

$$n_t = \sum_{\ell=0}^{t-1}(1-\delta)^{\ell}\, G(v_{t-1-\ell}) + (1-\delta)^t\, n_0. \tag{6.5}$$

For a given initial value of n_0 and a stochastic process for tariffs $\{\tau_t\}_{t=0}^{\infty}$, equations (6.4)-(6.5) and elasticities (6.1) characterize the path of the mass of exporters n_t. The evolution of n_t together with the static price and quantity decisions then fully determines exports $X_t = p_t^x q_t n_t$. We treat the elasticities (6.1) as constant throughout, which amounts to solving the model to first order.

Examples In the Krugman (1980) model or the Arkolakis (2010) model with a representative firm, $\eta_{p,\tau} = 0$ (recall this is the tariff-exclusive price elasticity), and $\eta_{q,p} = \eta_{q,\tau} = \eta_{\pi,\tau} = -\sigma$, where σ is the demand elasticity. In the Melitz (2003) model, if the exporting cutoff can change instantaneously conditional on the constant mass of firms n_t, $\eta_{p,\tau} = -\partial \ln \widetilde{\varphi}/\partial \ln \tau$, where $\widetilde{\varphi}$ is an aggregate productivity measure of firms serving the export market, and $\eta_{q,p} = \eta_{q,\tau} = -\sigma$. In the Krugman (1980) and Melitz (2003) models, n_t is the mass of entering firms and $G(\cdot)$ is the cumulative distribution function of the sunk costs of entry into exporting. To ensure smooth adjustment of the mass of firms following a change in trade costs, we assume that this distribution is nontrivial. In the Arkolakis (2010) model with a representative firm, n_t is the fraction of the foreign market penetrated by the firm, and the function G is a transformation of the convex cost of acquiring new customers. Appendix C.2 provides a detailed discussion of specific microfoundations of this model.

The short-run trade elasticity Let t_0 denote the date of the tariff change. The short run trade elasticity is:

$$\varepsilon^0 := \frac{d \ln X_{t_0}}{d \ln \tau_{t_0}} = (1 + \eta_{q,p})\, \eta_{p,\tau} + \eta_{q,\tau}. \tag{6.6}$$

Recall that the mass n_t is predetermined within the period, and hence the derivative of n_{t_0} with respect to τ_{t_0} is zero. The short-run trade elasticity is determined by the exporters' price response ($\eta_{p,\tau}$), the quantity response to tariff changes ($\eta_{q,\tau}$), and the quantity response to price changes ($\eta_{q,p}$). Because p_t^x and q_t are static decisions, they are fully determined by period-t tariffs. Thus, the short-run elasticity is not a function of future tariffs. As an example, in the Krugman (1980) model the short-run trade elasticity is $\varepsilon^0 = -\sigma$.

The long-run trade elasticity The long-run trade elasticity is the steady state change in trade following a steady state change in tariffs. The long-run trade elasticity differs from the short-run elasticity because n_t adjusts. If tariffs are constant ($\tau_t = \tau\ \forall t$) equation (6.4) becomes $v = \frac{\pi(\tau)}{\delta + r}$. Equation (6.5) then implies that n_t monotonically converges to $n = \frac{G(v)}{\delta}$.[33] It follows that $\frac{d \ln n}{d \ln \tau} = \chi \eta_{\pi,\tau}$, where $\chi := \frac{g(v)v}{G(v)}$. These two expressions characterize the non-stochastic steady state of the model. Hence, the long-run trade elasticity is

$$\varepsilon := \frac{d \ln X}{d \ln \tau} = \varepsilon^0 + \frac{d \ln n}{d \ln \tau} = \varepsilon^0 + \chi \eta_{\pi,\tau}. \tag{6.7}$$

In the long run, the response of trade to tariff changes depends on $\chi > 0$ and $\eta_{\pi,\tau} < 0$, the elasticity of flow profits with respect to tariffs. Consistent with intuition, the more sensitive are profits to tariffs, the greater the absolute value of the long-run trade elasticity.

The long-run trade elasticity increases in the elasticity χ of mass n with respect to value v. The

[33] The convergence of n_t to its steady state value is geometric and monotone. The rate of convergence is δ. We provide details in Appendix C.3.

precise meaning of χ depends on the underlying microfoundation. In the dynamic Krugman (1980) model, χ captures the mass of firms at the margin of entry. The greater the mass of firms at the margin, the more n changes in response to a change in per-firm profits and hence value v. In the dynamic Arkolakis (2010) model, firms face a convex cost function $f(a)$ of adding a mass of a new customers. In that case, $\chi = \left(\frac{f''(a)a}{f'(a)}\right)^{-1}$. Greater curvature of this cost function leads to a lower value of χ, implying a smaller trade response to tariff shocks.

Transitional dynamics and horizon-h elasticities To derive a horizon-specific elasticity, we must specify further details of the time path of tariffs. This is because unlike in the short run or the steady state calculations, the entire path of (expected) tariffs matters for the entry decision in each period. To make progress, we consider an unexpected change to tariffs at time t_0. This shock is followed by a subsequent evolution of tariffs (an impulse response), denoted by $\left\{\frac{d\ln\tau_{t_0+h}}{d\ln\tau_{t_0}}\right\}_{h=0}^{\infty}$. This sequence is the model counterpart of our estimated impulse response function of tariff changes as depicted in the left panel of Figure 1. Since the tariff shock at time t_0 may be followed by further shocks thereafter, agents cannot perfectly predict future tariffs or profits and therefore form expectations as in equation (6.4).

The horizon-h impulse response function of trade to the tariff shock at t_0 is:

$$\frac{d\ln X_{t_0+h}}{d\ln\tau_{t_0}} = \varepsilon^0 \frac{d\ln\tau_{t_0+h}}{d\ln\tau_{t_0}} + \frac{d\ln n_{t_0+h}}{d\ln\tau_{t_0}}. \tag{6.8}$$

The horizon-h trade elasticity is then computed as the ratio of the two impulse response functions:

$$\varepsilon^h := \frac{\frac{d\ln X_{t_0+h}}{d\ln\tau_{t_0}}}{\frac{d\ln\tau_{t_0+h}}{d\ln\tau_{t_0}}} = \varepsilon^0 + \frac{\frac{d\ln n_{t_0+h}}{d\ln\tau_{t_0}}}{\frac{d\ln\tau_{t_0+h}}{d\ln\tau_{t_0}}}, \tag{6.9}$$

as long as this object is finite (i.e. $\frac{d\ln\tau_{t_0+h}}{d\ln\tau_{t_0}} \neq 0$). Note that this definition of the horizon-h trade elasticity coincides with equation (2.1) for a tariff change of one marginal unit, when we replace the infinitesimal difference with the difference operator Δ.

To fully characterize the horizon-h trade elasticity, we must characterize the last term in (6.9), the adjustment of n_t to the tariff shock.

Proposition 1. *Consider an arbitrary evolution of tariffs* $\left\{\frac{d\ln\tau_{t_0+\ell}}{d\ln\tau_{t_0}}\right\}_{\ell=1}^{\infty}$ *after the shock at* t_0*. The impulse response function of* $\ln n_t$ *at horizon* $h = 0, 1, 2, ...$ *is*

$$\frac{d\ln n_{t_0+h}}{d\ln\tau_{t_0}} = \chi\eta_{\pi,\tau}\frac{\delta+r}{1+r}\delta\sum_{k=0}^{h-1}(1-\delta)^{h-1-k}\,\mathbb{E}_{t_0+k}\left[\sum_{\ell=0}^{\infty}\left(\frac{1-\delta}{1+r}\right)^{\ell}\frac{d\ln\tau_{t_0+k+\ell+1}}{d\ln\tau_{t_0}}\right]. \tag{6.10}$$

Proof. See Appendix C.4. □

Plugging (6.10) into (6.9) delivers the horizon-h trade elasticity. As is clear from equations (6.8) and (6.9), the sluggish adjustment of trade to tariff shocks is entirely driven by the sluggish adjustment of n_t. While this adjustment is somewhat complicated (equation 6.10), it delivers a useful insight: in general, all tariff changes from time t_0 into the infinite future affect the trade response to tariff shocks. Proposition 1 captures these tariff changes as the elasticities of time $t_0 + \ell$ tariffs with respect to the tariff shock at time t_0, for $\ell = 1, 2, ...$ For a given time horizon h, elasticities for $0 \leq \ell < h$ reflect changes to past tariffs, the elasticity for $\ell = h$ reflects a change to current tariffs, and elasticities for $\ell > h$ reflect expected changes to future tariffs.

As the following proposition shows, ε^h converges to the long-run trade elasticity, unless the tariff change induced by the shock in period t_0 returns to zero in the limit.

Proposition 2. *If* $\lim_{h\to\infty} \frac{d \ln \tau_{t_0+h}}{d \ln \tau_{t_0}} \neq 0$ *and is finite, then* $\lim_{h\to\infty} \varepsilon^h = \varepsilon$.

Proof. See Appendix C.4. □

Although not surprising, this result is important because it validates our interpretation of horizon-h trade elasticities for large h as estimates of the long-run elasticity.

For concreteness, we next consider two simple examples.

Example 1: tariff constant after 1 period Let there be a surprise change in the tariff sequence of the form $\left\{ \frac{d \ln \tau_{t_0+h}}{d \ln \tau_{t_0}} \right\}_{h=0}^{\infty} = \{1, \Delta \ln \tau_{>t_0}, \Delta \ln \tau_{>t_0}, \Delta \ln \tau_{>t_0}, ...\}$. That is, the tariff change takes the value one in the impact period, and is subsequently constant at $\Delta \ln \tau_{>t_0}$. Note that this example nests a one-time permanent change in tariffs (if $\Delta \ln \tau_{>t_0} = 1$), and is a good approximation of our estimated impulse response function in Figure 1.

At horizon $h \geq 1$ the trade elasticity is

$$\varepsilon^h = \varepsilon^0 + \chi \eta_{\pi,\tau} \left(1 - (1 - \delta)^h \right), \tag{6.11}$$

with ε^0 given by (6.6). The trade elasticity converges geometrically to the long-run trade elasticity at the rate δ. Convergence occurs in one period if $\delta = 1$.

Example 2: AR(1) Second, let the tariffs follow a first order autoregressive process following an initial shock, so that $\Delta \ln \tau_{t+1} = \rho \cdot \Delta \ln \tau_t$ for $t > t_0$ and $0 < \rho < 1$. Since this process is mean-reverting, the tariff change approaches zero as h tends to infinity. It follows that the premise of Proposition 2 does not hold and that the long-run trade elasticity is not defined in this case. However, we can still compute the elasticity at a finite horizon.

First, consider the case $1-\delta<\rho$. Intuitively, this condition requires that the rate of depreciation is higher than the rate of mean reversion of tariffs. In this case the horizon-h trade elasticity is

$$\varepsilon^{h}=\varepsilon^{0}+\chi\eta_{\pi,\tau}\frac{(\delta+r)\,\delta}{[1+r-(1-\delta)\,\rho]\left(1-\frac{1-\delta}{\rho}\right)}\left(1-\left(\frac{1-\delta}{\rho}\right)^{h}\right). \tag{6.12}$$

As in Example 1, the trade elasticity increases with time horizon h in absolute value. Further, with $1-\delta<\rho$ the horizon-h trade elasticity does converge, although not generally to the long-run trade elasticity. While convergence is still geometric, the rate of convergence now depends on the persistence of the tariff process. Convergence is faster for more persistent tariff processes, i.e. greater values of ρ. If tariffs mean-revert sufficiently quickly, $\rho\leq 1-\delta$, the horizon-h trade elasticity does not converge.

Notice that as ρ approaches 1, the horizon-h trade elasticity in the AR(1) case (6.12) converges pointwise to the horizon-h trade elasticity under a permanent tariff change (6.11). This property is important for our empirical application. Although tariffs changes in our sample retain 75% of their initial impulse 10 years later, in short samples it is not possible to statistically distinguish between tariff processes featuring truly permanent or highly persistent tariff changes (Hamilton, 1994, p. 445). Since Proposition 2 does not apply under mean-reverting tariffs ($\rho<1$), one may be concerned that the horizon-h trade elasticity is not informative about the long-run trade elasticity. This property alleviates this concern. For ρ sufficiently close to one, the horizon-h trade elasticity essentially converges to the long-run trade elasticity, even though tariffs mean-revert in the very long run.

Estimating equations While we led off the paper with an atheoretical estimating equation, we now show that this estimating equation can be microfounded by means of the model above.

Proposition 3. *The model delivers estimating equation* (2.2), *where*

$$\beta_{X}^{h}=\chi\eta_{\pi,\tau}\frac{\delta+r}{1+r}\delta\sum_{k=0}^{h-1}(1-\delta)^{h-1-k}\sum_{\ell=0}^{\infty}\left(\frac{1-\delta}{1+r}\right)^{\ell}\beta_{\tau}^{k+\ell+1}+\varepsilon^{0}\beta_{\tau}^{h}.$$

β_{τ}^{h} *is defined as the regression coefficient of* $\Delta_{h}\ln\tau_{i,j,p,t}$ *on* $\Delta_{0}\ln\tau_{i,j,p,t}$ *in the population, and can be estimated from equation* (2.3).

After augmenting the model with additional shocks, the fixed effects $\delta_{j,p,t}^{X,h}$ *and* $\delta_{i,p,t}^{X,h}$ *capture a weighted sum of past, present, and expected future supply and demand shocks, respectively. The error term* *includes past, present, and expected future time-varying bilateral and product-specific demand shocks and non-tariff trade barriers, as well as the initial state.*

Proof. See Appendix C.4. □

Quantification Next, we explore the time path of elasticities. To do this, we calibrate the dynamic model and subject it to the two tariff shocks in the examples above.

We choose a demand elasticity σ of 1.1. This parameter immediately determines the short-run elasticity, since in the CES-monopolistic competition model $\varepsilon^0 = -\sigma$. Based on equation (6.7), and using the fact that $\eta_{\pi,\tau} = -\sigma$ in the CES-monopolistic competition model, we set $\chi = 0.82$ to match our estimated long-run elasticity of $\varepsilon \approx -2$. We further set the depreciation rate to $\delta = 0.25$ to roughly match the rate of convergence to the long run. Calibration of these parameters is sufficient to compute the transition path of exports in Example 1. For Example 2, we also need the interest rate and the AR(1) coefficient. We set these to $r = 0.03$ and $\rho = 0.955$. The latter parameter is chosen to roughly match the impulse response function of tariffs.

The left panel of Figure 5 plots the paths of tariffs. The red line depicts the tariff response of Example 1, where tariffs increase by one unit in the impact period, and then stay constant at 0.75 starting in period 1 onwards. The blue line is the AR(1) path of tariffs following an impulse of unit size (Example 2). The green line plots the impulse response of tariffs estimated in the data, which is quite similar to the two model experiments.

The right panel of Figure 5 displays the trade elasticities. The green line depicts the econometric estimates. Because the data are annual, and it is unlikely that all tariff changes went into effect on January 1, the year-zero trade elasticity is most likely subject to partial-year effects. Thus, for the purposes of comparing to the model, we consider the $h = 1$ empirical estimate to be the impact elasticity ε^0. The red and blue lines depict the model trade elasticity in the two experiments. They are nearly indistinguishable from one another.

The model succeeds in delivering a smooth path of adjustment that takes approximately a decade. The key parameter for the speed of adjustment is the depreciation rate δ. The slow adjustment observed in the data implies that δ is substantially below 1. The main shortcoming of the model is that it cannot match our short-run elasticity point estimate of -0.76, since the CES-monopolistic competition assumption requires that $\sigma > 1$.[34]

[34] A natural conjecture is that flexible markups may help push the short-run trade elasticity below 1. We experimented with versions of the model with local distribution costs à la Burstein, Neves, and Rebelo (2003). With local distribution costs, the net-of-tariff price received by the exporter p_t^x falls when a tariff increases, helping push down the trade elasticity all else equal. However, the flip side of a fall in p_t^x is a *ceteris paribus* increase in the quantity imported. It

FIGURE 5: Time Path of Elasticities in the Dynamic Model

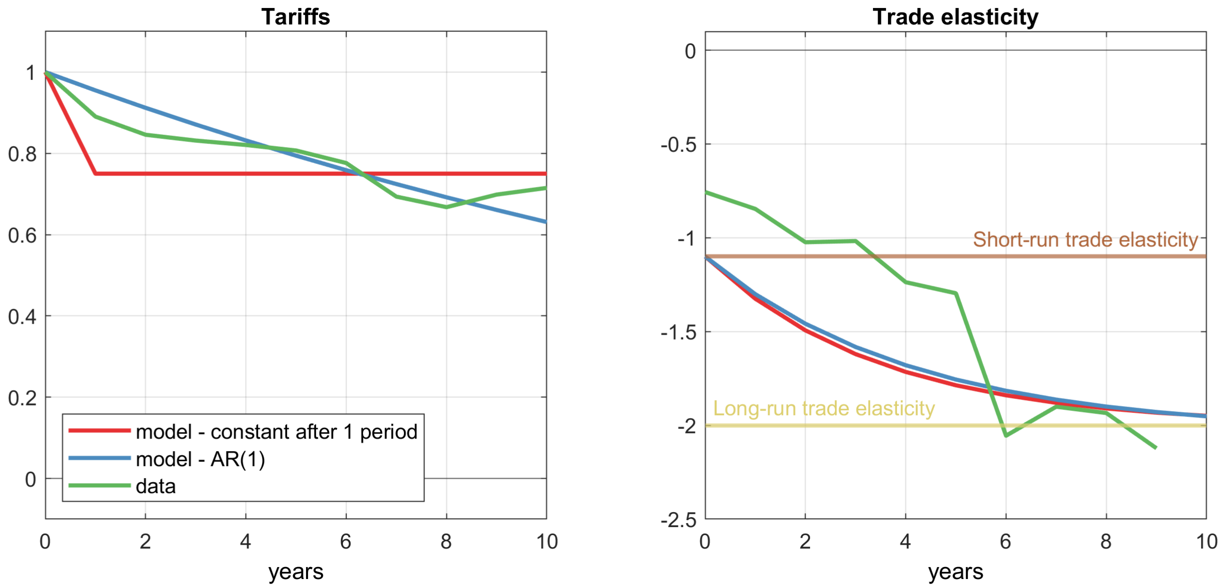

Notes: This figure illustrates the trade elasticities as implied by the model.

6.2 The Long Run Welfare Gains from Trade

As is well known from Arkolakis, Costinot, and Rodríguez-Clare (2012, henceforth ACR), the gains from trade relative to autarky in many quantitative trade models can be expressed as a function of the trade elasticity and the domestic absorption share: $1 - \lambda_{jj}^{1/\theta}$, with λ_{jj} the share of spending on domestically-produced goods in total spending. The static models in which this formula applies are metaphors for the long run, and the gains from trade should be interpreted as steady state comparisons between autarky and trade. Thus, we use the longest horizon elasticity estimated above, $h = 10$, as the long-run value.

To translate our estimates to the welfare-relevant elasticity θ, it is important to note that in our data the outcome variable $X_{i,j,t}$ does not include tariff payments. On the other hand, most theoretical gravity relationships relate tariff-inclusive spending by domestic agents to trade costs, with the trade elasticity defined correspondingly. To go from our estimated coefficient to the elasticity relevant for

turns out there is no combination of $\sigma > 1$ and local distribution share between 0 and 1 that delivers a less than unitary trade elasticity as we measure it (of $p_t^x q_t$ with respect to τ_t). In addition, Table 4 shows a virtually nil response of p_t^x to tariffs, a finding consistent with recent estimates using the US-China trade war (Fajgelbaum et al., 2020; Cavallo et al., 2019). Both of these points suggest that imperfect pass-through into net-of-tariff prices is unlikely to produce a short-run elasticity below 1. Developing a framework that can successfully reproduce a short-run elasticity below 1 remains a fruitful avenue for future research. One possibility is variable distribution margins. Indeed, Cavallo et al. (2019) document a fall in retail margins for US imports affected by the trade war.

welfare, we must add 1. Thus, our estimates imply that the welfare-relevant elasticity θ is around -1.[35]

Figure 6 displays the gains from trade as a function of λ_{jj}, under our value of θ and under an elasticity of -5 considered by ACR.[36] As expected, the gains from trade are substantially larger with our elasticity. For the US, gains from trade are 5.27% for $\theta = -1$, compared to 1.0% for $\theta = -5$. The median welfare gain is 22.9% in a sample of 64 countries, compared to 4.2% implied by $\theta = -5$. Table A11 reports the gains from trade under $\theta = -1$, -5, and -10 for selected countries in the sample.

The blue bars in Appendix Figure A10 report the gains from trade using the multi-sector ACR formula and our sector-specific elasticity values (section 4.1). We benchmark these to the sector-specific trade elasticity estimates from Ossa (2015), which explores the properties of multi-sector ACR formulas. To do this, we concord the sectoral elasticity estimates in that paper to the 11 HS sections for which we estimate elasticities. Once again, the gains from trade implied by our estimates are considerably larger than previously suggested in the literature. Our estimates applied to the ACR multi-sector formula imply mean gains from trade of 26.7%, compared to 12.8% using the elasticities in Ossa (2015).

We caveat these results in two respects. First, we acknowledge that ACR formulas are not known to apply in explicitly dynamic models (for some results bridging ACR with dynamics, see Arkolakis, Eaton, and Kortum, 2011; Alessandria, Choi, and Ruhl, 2014). This is a general critique of all applications of the ACR formulas in static environments. Nonetheless, the widespread use of ACR formulas makes them a natural setting for benchmarking the implications of our elasticity estimates relative to the conventional values. Exploring the implications of our estimates for gains from trade in explicitly dynamic settings is a fruitful avenue for future research.

Second, care must be taken when going from the micro elasticity estimated in our empirical work to the macro elasticity that enters the ACR formula. The calculations above make the implicit assumption that the two coincide. While there are many models in which this is not true, some of this concern can be allayed by using the multi-sector variant of the formula, that aligns more closely

[35] Let $\widetilde{X}_{i,j} = \tau_{i,j} X_{i,j}$ be the steady state spending by consumers in economy i on goods from j inclusive of tariffs. The elasticity in the ACR formula is $\theta = d\ln \widetilde{X}_{i,j}/d\ln \tau_{i,j}$. The object we estimate is the elasticity of trade flows exclusive of tariff payments $X_{i,j}$ to tariffs: $d\ln X_{i,j}/d\ln \tau_{i,j} = \theta - 1$. As an example, in an Armington/Krugman setting θ corresponds to $1-\sigma$, where σ is the elasticity of substitution between goods coming from different origins. In that case, our long-run coefficient estimate has the interpretation $\varepsilon = -\sigma$. In an Eaton-Kortum setting, θ is the Frechet dispersion parameter. In that case, it can be recovered by adding 1 to our ε estimates. This calculation assumes that $X_{i,j}$ is recorded as c.i.f. The relationship between ε and θ is the same if $X_{i,j}$ is f.o.b., since the iceberg costs $\kappa_{i,j}$ are log-additive and make up the error term in our estimation.

[36] We use data from the OECD IO tables for 64 countries for the year 2006, the midpoint of our trade and tariff sample. We compute import penetration by dividing imports by gross output as in ACR.

FIGURE 6: Gains from Trade

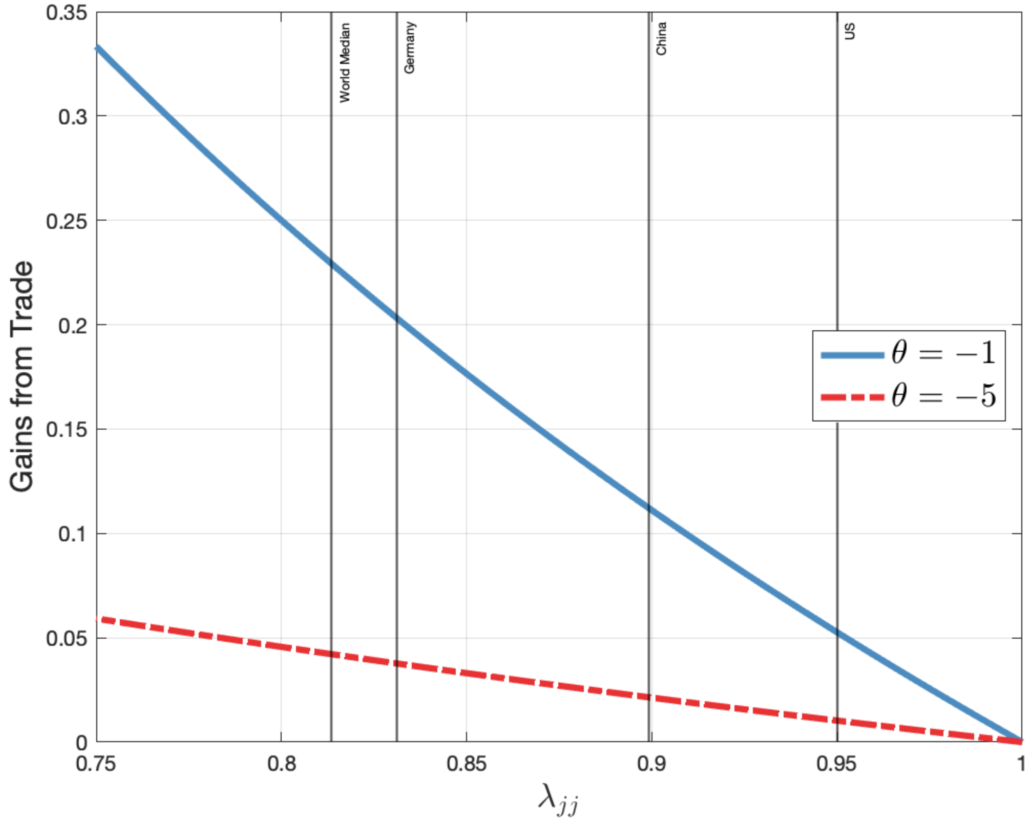

Notes: This figure displays the gains from trade as a function of the domestic absorption ratio λ_{jj} under our baseline welfare-relevant elasticity of -1 (solid blue line) and a comparison elasticity of -5 (red dashed line). "World Median" denotes the median domestic absorption ratio of the 64 countries in the OECD world input-output tables in 2006.

the levels of disaggregation at which the coefficients are estimated and the theory. Using our micro elasticity values in place of the macro elasticity is conservative in the sense that we would expect the elasticities of substitution to be higher at finer levels of product disaggregation.

7 Conclusion

We develop a novel method to estimate the trade elasticity, a key parameter in virtually all models in international economics. To tackle the endogeneity problem that tariffs and trade flows are jointly determined, we propose an instrument that relies on the WTO's MFN principle. We estimate trade elasticities at different horizons, and find short-run values of about -0.76, and long-run values close to -2. The estimates are robust to alternative specifications of the instrument and controls, and uniformly larger compared to estimates that use endogenous tariff variation. Our empirical strategy is not specific to a particular theoretical framework, and applies to all models that have a gravity structure.

The long-run estimates imply the welfare-relevant trade elasticity is around 1 in absolute value. This is significantly smaller than conventional wisdom in the literature, suggesting the welfare gains from trade are larger than previously thought. Our finding that the trade elasticity differs by horizon and converges to the "long-run" after about 7-10 years implies substantial adjustment costs to changing trade volumes.

References

Adão, Rodrigo, Costas Arkolakis, and Sharat Ganapati. 2020. "Not-Parametric Gravity: Measuring the Aggregate Implications of Firm Heterogeneity." Mimeo, Chicago Booth, Yale, and Georgetown.

Alessandria, George, Costas Arkolakis, and Kim J. Ruhl. 2020. "Firm dynamics and trade." Forthcoming, *Annual Review of Economics.*

Alessandria, George and Horag Choi. 2014. "Establishment heterogeneity, exporter dynamics, and the effects of trade liberalization." *Journal of International Economics* 94 (2):207 – 223.

———. 2019. "Entry, Trade, and Exporting over the Cycle." *Journal of Money, Credit and Banking* 51 (S1):83–126.

Alessandria, George, Horag Choi, and Kim Ruhl. 2014. "Trade Adjustment Dynamics and the Welfare Gains from Trade." NBER Working Paper 20663.

Anderson, James and Eric van Wincoop. 2004. "Trade Costs." *Journal of Economic Literature* 42 (3):691–751.

Anderson, James E. 1979. "A Theoretical Foundation for the Gravity Equation." *American Economic Review* 69 (1):106–116.

Anderson, James E. and Eric van Wincoop. 2003. "Gravity with Gravitas: A Solution to the Border Puzzle." *American Economic Review* 93 (1):170–192.

Antràs, Pol and Gerard Padró i Miquel. 2011. "Foreign Influence and Welfare." *Journal of International Economics* 84 (2):135–148.

Arkolakis, Costas. 2010. "Market Penetration Costs and the New Consumers Margin in International Trade." *Journal of Political Economy* 118 (6):1151–1199.

Arkolakis, Costas, Arnaud Costinot, and Andrés Rodríguez-Clare. 2012. "New Trade Models, Same Old Gains?" *American Economic Review* 102 (1):94–130.

Arkolakis, Costas, Jonathan Eaton, and Samuel S. Kortum. 2011. "Staggered Adjustment and Trade Dynamics." Mimeo, Yale and Penn State.

Atkeson, Andrew and Ariel Burstein. 2008. "Trade Costs, Pricing-to-Market, and International Relative Prices." *American Economic Review* 98 (5):1998–2031.

Bagwell, Kyle and Robert W. Staiger, editors. 2016. *Handbook of Commercial Policy*, vol. 1. Elsevier-North Holland.

Baier, Scott L. and Jeffrey H. Bergstrand. 2007. “Do free trade agreements actually increase members’ international trade?” *Journal of International Economics* 71 (1):72 – 95.

Baldwin, Richard and Daria Taglioni. 2006. “Gravity for Dummies and Dummies for Gravity Equations.” NBER Working Paper 12516.

Bas, Maria, Thierry Mayer, and Mathias Thoenig. 2017. “From micro to macro: Demand, supply, and heterogeneity in the trade elasticity.” *Journal of International Economics* 108:1 – 19.

Bellemare, Marc F. and Casey J. Wichman. 2020. “Elasticities and the Inverse Hyperbolic Sine Transformation.” *Oxford Bulletin of Economics and Statistics* 82 (1):50–61.

Blaum, Joaquin. 2019. “Global Firms in Large Devaluations.” Mimeo, Brown University.

Bown, Chad P. and Meredith A Crowley. 2013. “Import Protection, Business Cycles, and Exchange Rates: Evidence from the Great Recession.” *Journal of International Economics* 90 (1):50–64.

Bown, Chad P. and Meredith A. Crowley. 2014. “Emerging economies, trade policy, and macroeconomic shocks.” *Journal of Development Economics* 111:261 – 273.

Bown, Chad P. and Meredith A Crowley. 2016. “The Empirical Landscape of Trade Policy.” In *The Handbook of Commercial Policy*, vol. 1, edited by Kyle Bagwell and Robert W. Staiger, chap. 1. Elsevier-North Holland, 3–108.

Broda, Christian and David Weinstein. 2006. “Globalization and the Gains from Variety.” *Quarterly Journal of Economics* 121 (2):541–85.

Buono, Ines and Guy Lalanne. 2012. “The effect of the Uruguay round on the intensive and extensive margins of trade.” *Journal of International Economics* 86 (2):269 – 283.

Burbidge, John B., Lonnie Magee, and A. Leslie Robb. 1988. “Alternative Transformations to Handle Extreme Values of the Dependent Variable.” *Journal of the American Statistical Association* 83 (401):123–127.

Burstein, Ariel and Marc Melitz. 2013. “Trade Liberalization and Firm Dynamics.” In *Advances in Economics and Econometrics Tenth World Congress. Applied Economics*, vol. 2. Cambridge, UK: Cambridge University Press.

Burstein, Ariel T., Joao C. Neves, and Sergio Rebelo. 2003. “Distribution costs and real exchange rate dynamics during exchange-rate-based stabilizations.” *Journal of Monetary Economics* 50 (6):1189–1214.

Bustos, Paula. 2011. “Trade Liberalization, Exports, and Technology Upgrading: Evidence on the Impact of MERCOSUR on Argentinian Firms.” *American Economic Review* 101 (1):304–40.

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

Canova, Fabio and Luca Sala. 2009. "Back to square one: Identification issues in DSGE models." *Journal of Monetary Economics* 56 (4):431 – 449.

Cavallo, Alberto, Gita Gopinath, Brent Neiman, and Jenny Tang. 2019. "Tariff Passthrough at the Border and at the Store: Evidence from US Trade Policy." Forthcoming, *American Economic Review: Insights*.

Correia, Sergio. 2015. "Singletons, Cluster-Robust Standard Errors and Fixed Effects: A Bad Mix." Mimeo, Duke University.

Costantini, James and Marc Melitz. 2007. "The Dynamics of Firm-Level Adjustment to Trade Liberalization." In *The Organization of Firms in a Global Economy*, edited by E. Helpman, D. Marin, and T. Verdier. Cambridge, MA: Harvard University Press.

Costinot, Arnaud and Andrés Rodríguez-Clare. 2014. "Trade Theory with Numbers: Quantifying the Consequences of Globalization." *Handbook of International Economics* 4:197–261.

Donaldson, Dave. 2018. "Railroads of the Raj: Estimating the Impact of Transportation Infrastructure." *American Economic Review* 108 (4-5):899–934.

Drozd, Lukasz A. and Jaromir B. Nosal. 2012. "Understanding International Prices: Customers as Capital." *American Economic Review* 102 (1):364–95.

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

Fajgelbaum, Pablo D., Pinelopi K. Goldberg, Patrick J. Kennedy, and Amit K. Khandelwal. 2020. "The Return to Protectionism." *Quarterly Journal of Economics* 135 (1):1–55.

Feenstra, Robert. 1994. "New Product Varieties and the Measurement of International Prices." *American Economic Review* 84 (1):157–77.

Feenstra, Robert C., Philip Luck, Maurice Obstfeld, and Katheryn N. Russ. 2018. "In Search of the Armington Elasticity." *Review of Economics and Statistics* 100 (1):135–150.

Fitzgerald, Doireann and Stefanie Haller. 2018. "Exporters and shocks." *Journal of International Economics* 113 (C):154–171.

Fitzgerald, Doireann, Stefanie Haller, and Yaniv Yedid-Levi. 2019. "How Exporters Grow." Mimeo, Federal Reserve Bank of Minneapolis and Trinity College Dublin and Tiomkin School of Economics.

Fontagné, Lionel, Philippe Martin, and Gianluca Orefice. 2018. "The international elasticity puzzle is worse than you think." *Journal of International Economics* 115:115 – 129.

Gallaway, Michael P, Christine A McDaniel, and Sandra A Rivera. 2003. "Short-run and long-run industry-level estimates of U.S. Armington elasticities." *The North American Journal of Economics and Finance* 14 (1):49 – 68.

Gawande, Kishore, Pravin Krishna, and Michael J. Robbins. 2006. "Foreign Lobbies and U.S. Trade Policy." *The Review of Economics and Statistics* 88 (3):563–571.

Giri, Rahul, Kei-Mu Yi, and Hakan Yilmazkuday. 2020. "Gains from Trade: Does Sectoral Heterogeneity Matter?" Forthcoming, *Journal of International Economics*.

Hamilton, James D. 1994. *Time Series Analysis*. Princeton, NJ: Princeton University Press.

Head, Keith and Thierry Mayer. 2014. "Gravity Equations: Workhorse, Toolkit, and Cookbook." In *Handbook of International Economics*, vol. 4, edited by Elhanan Helpman, Kenneth Rogoff, and Gita Gopinath. Elsevier B.V., 131–195.

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

Heathcote, Jonathan and Fabrizio Perri. 2002. "Financial autarky and international business cycles." *Journal of Monetary Economics* 49 (3):601–627.

Hooper, Peter, Karen Johnson, and Jaime Marquez. 2000. "Trade Elasticities for the G-7 Countries." *Princeton Studies in International Economics* 87.

Imbs, Jean and Isabelle Mejean. 2015. "Elasticity Optimism." *American Economic Journal: Macroeconomics* 7 (3):43–83.

———. 2017. "Trade Elasticities." *Review of International Economics* 25 (2):383–402.

Jordà, Òscar. 2005. "Estimation and Inference of Impulse Responses by Local Projections." *American Economic Review* 95 (1):161–182.

Kehoe, Timothy J. and Kim J. Ruhl. 2013. "How Important Is the New Goods Margin in International Trade?" *Journal of Political Economy* 121 (2):358–392.

Khan, Shafaat Y. and Armen Khederlarian. 2020. "How Does Trade Respond to Anticipated Tariff Changes? Evidence from NAFTA." Mimeo, University of Rochester.

Krugman, Paul. 1980. "Scale Economies, Product Differentiation, and the Pattern of Trade." *American Economic Review* 70 (5):950–59.

Lai, Huiwen and Daniel Trefler. 2002. "The Gains from Trade with Monopolistic Competition: Specification, Estimation, and Mis-Specification." NBER Working Paper 9169.

Lake, James and Maia K. Linask. 2016. "Could tariffs be pro-cyclical?" *Journal of International Economics* 103:124 – 146.

Lileeva, Alla and Daniel Trefler. 2010. "Improved Access to Foreign Markets Raises Plant-level Productivity…For Some Plants*." *The Quarterly Journal of Economics* 125 (3):1051–1099.

Lind, Nelson and Natalia Ramondo. 2020. "Trade with Correlation." Mimeo, Emory University and Boston University.

Melitz, Marc J. 2003. "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity." *Econometrica* 71 (6):1695–1725.

Novy, Dennis. 2013. "International trade without CES: Estimating translog gravity." *Journal of International Economics* 89 (2):271 – 282.

Ossa, Ralph. 2015. "Why trade matters after all." *Journal of International Economics* 97 (2):266 – 277.

Pierce, Justin R. and Peter K. Schott. 2016. "The Surprisingly Swift Decline of US Manufacturing Employment." *American Economic Review* 106 (7):1632–62.

Redding, Stephen and Anthony J. Venables. 2004. "Economic geography and international inequality." *Journal of International Economics* 62 (1):53–82.

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

Ruhl, Kim J. 2008. "The International Elasticity Puzzle." Working Paper.

Ruhl, Kim J. and Jonathan L. Willis. 2017. "New Exporter Dynamics." *International Economic Review* 58 (3):703–726.

Sequeira, Sandra. 2016. "Corruption, Trade Costs, and Gains from Tariff Liberalization: Evidence from Southern Africa." *American Economic Review* 106 (10):3029–63.

Simonovska, Ina and Michael E. Waugh. 2014. "The Elasticity of Trade: Estimates and Evidence." *Journal of International Economics* 92 (1):34–50.

Sims, Christopher A. 1980. "Macroeconomics and Reality." *Econometrica* 48 (1):1–48.

Soderbery, Anson. 2015. "Estimating import supply and demand elasticities: Analysis and implications." *Journal of International Economics* 96 (1):1–17.

———. 2018. “Trade elasticities, heterogeneity, and optimal tariffs.” *Journal of International Economics* 114:44 – 62.

Trefler, Daniel. 1993. “Trade Liberalization and the Theory of Endogenous Protection: An Econometric Study of U.S. Import Policy.” *Journal of Political Economy* 101 (1):138–160.

World Trade Organization. 2014. “World Tariff Profiles 2014.” Geneva: WTO, UNCTAD, and ITC.

Yilmazkuday, Hakan. 2019. “Estimating the trade elasticity over time.” *Economics Letters* 183:108579.

Appendix A Data

Figure A1: Patterns in Tariffs: Frequency of Changes

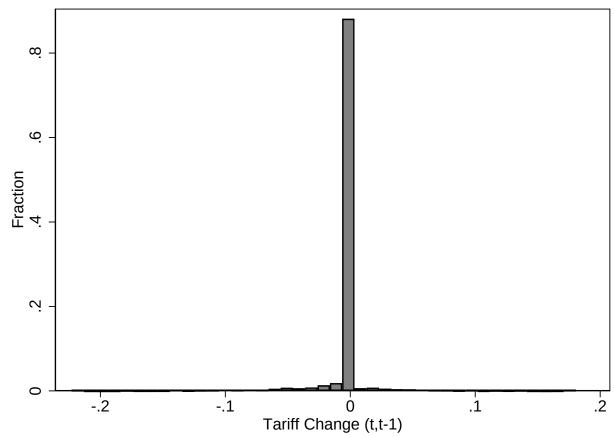

Unconditional

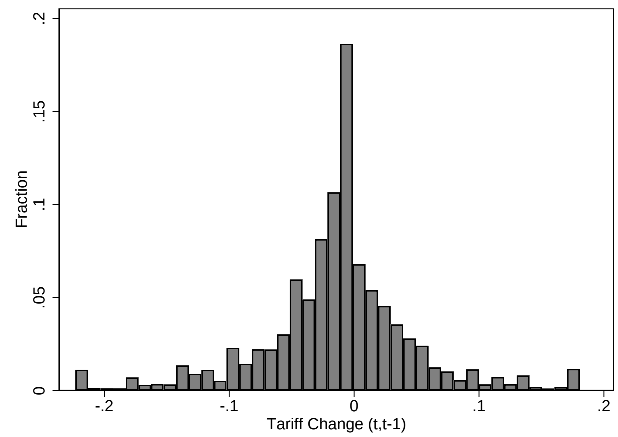

Unconditional, Excluding Zeros

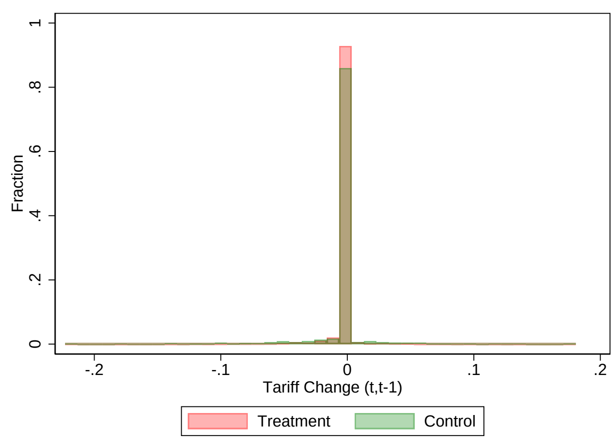

Treatment and Control

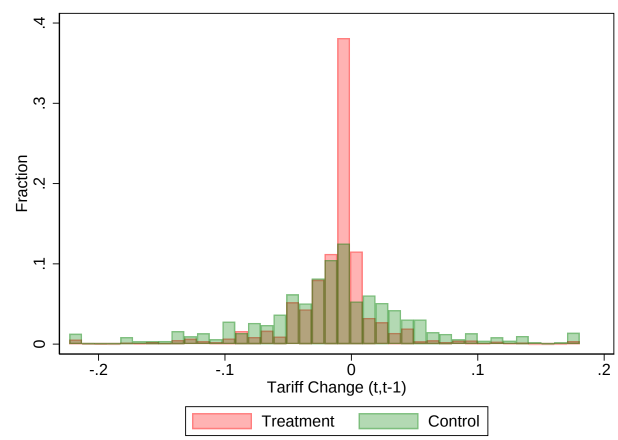

Treatment and Control, Excluding Zeros

Notes: These figures display the frequency of tariff changes in our data. The top two panels display the unconditional frequency of all tariff changes (top left) and frequency excluding zeros (top right). The bottom panel displays the overlap in the frequency of changes in the treatment and control groups, including zero changes (left panel) and removing zero changes (right panel).

Figure A2: Patterns in Tariffs: Autocorrelation

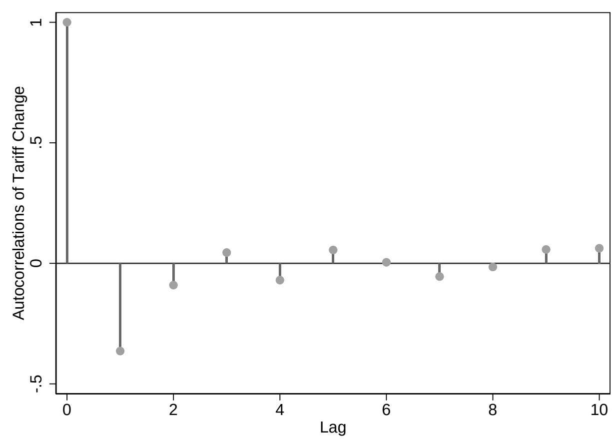

Notes: This figure displays the unconditional autocorrelation of tariff changes in the sample.

Figure A3: Fraction of World Imports (Average, %)

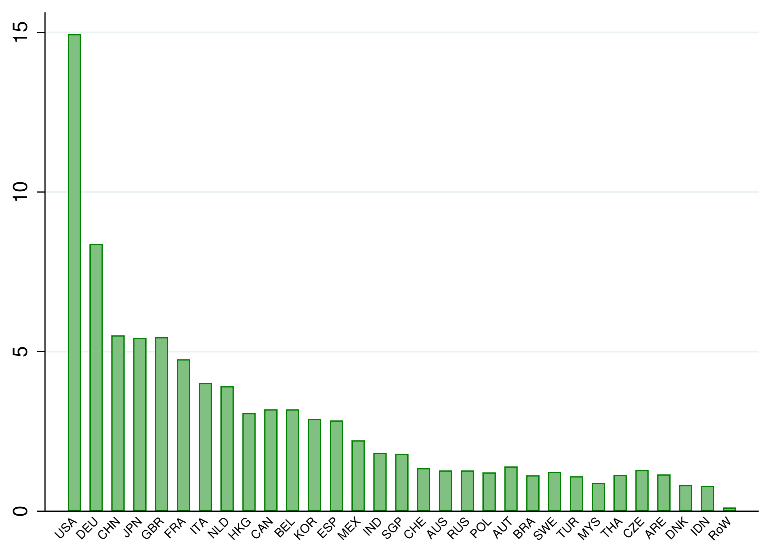

Notes: This figure shows the average fraction of world trade flows by importer in our sample. "ROW" is the mean share of world trade among countries outside of the top 20 importers.

Figure A4: Fraction of World Imports by HS Section (Average, %)

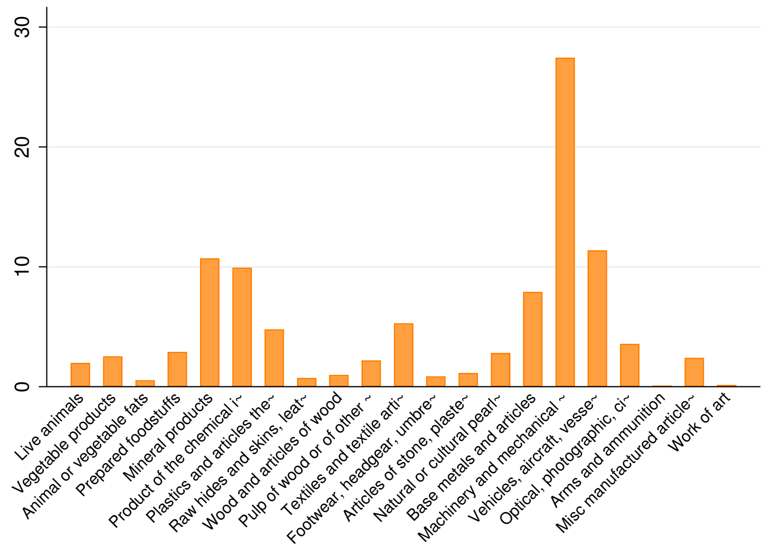

Notes: This figure shows the average fraction of trade that is in each HS Section in our sample.

Figure A5: Fraction of World Imports: MFN vs. non-MFN (%)

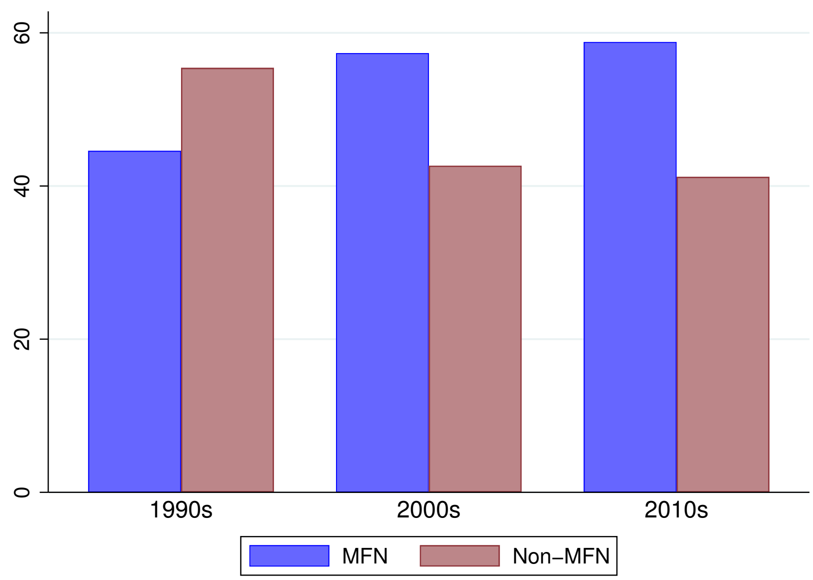

Notes: This figure shows the average fraction of world trade that is subject to MFN tariffs and non-MFN tariffs by decade in our sample.

Figure A6: Explaining Country Variation

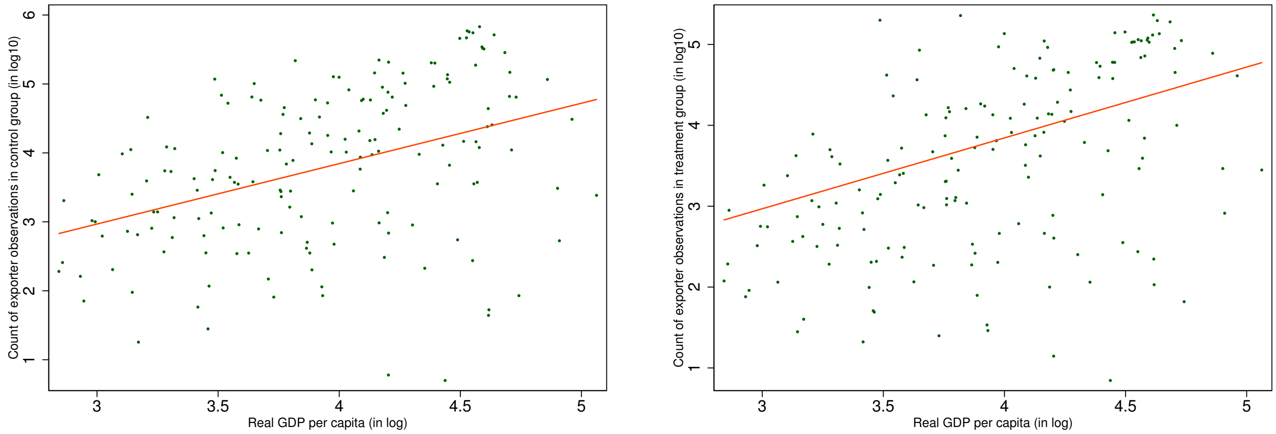

Control Treatment

Notes: This figure plots the (log) counts a country appears in the control group (left panel) and in the treatment group (right panel) against log real PPP-adjusted per capita income from the Penn World Tables, after taking out the variation absorbed by the fixed effects. The line depicts the OLS fit.

Figure A7: Product Variation

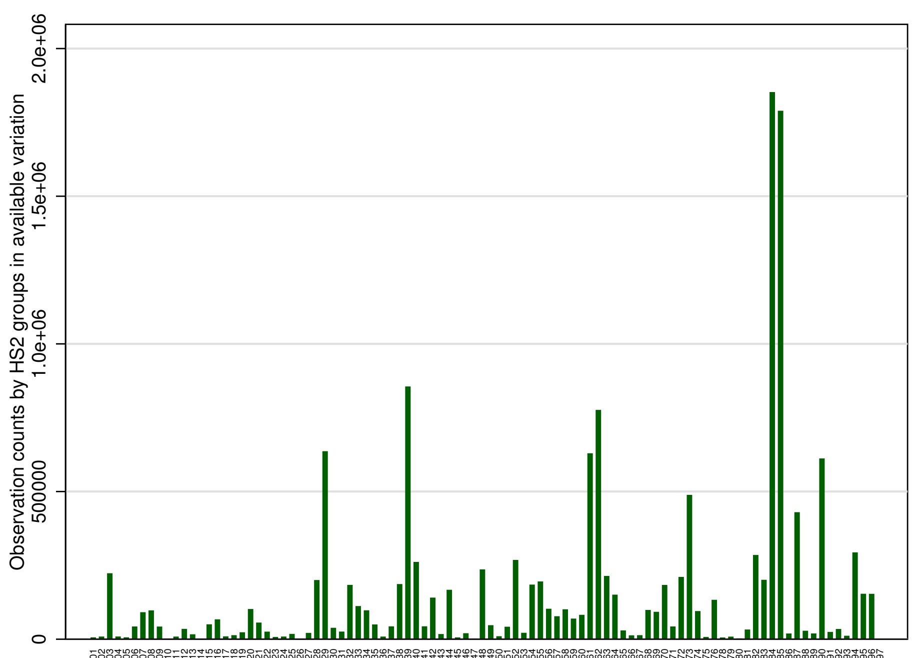

Notes: This figure plot the frequency of observations belonging to each HS-2 category, after taking out the variation absorbed by the fixed effects and imposing the sample restrictions.

Table A1: Fraction of Unique Mappings Across HS Revisions (percent)

		Mapped to:				
		HS-92	HS-96	HS-02	HS-07	HS-12
Mapped from:	HS-96	89.38				
	HS-02	81.55	90.81			
	HS-07	73.34	80.74	88.48		
	HS-12	68.17	74.91	81.81	91.93	
	HS-17	61.85	67.92	73.62	81.99	88.05

Notes: This table presents the fraction of HS codes that can be mapped uniquely from one HS revision (in the "Mapped from" row) to another HS revision (in a "Mapped to" column). All numbers are in percent.

Table A2: Instrument – Illustration

Importer	MFN Trade Partners		Major Trade Partners Aggregate		Major Trade Partners HS 6403		Treatment	Control	Excluded
	2005	2006	2005	2006	2005	2006			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A: Germany									
	USA	USA	FRA	FRA	ITA	ITA	HKG	ITA	CHN
	CHN	CHN	CHN	CHN	CHN	CHN	KOR	VNM	IND
	JPN	JPN	NLD	NLD	VNM	VNM	SGP	PRT	USA
	KOR	KOR	ITA	ITA	PRT	PRT	NZL	AUT	JPN
	IND	IND	USA	USA	AUT	AUT	CAN	NLD	
	CAN	CAN	GBR	GBR	IND	IND	AUS	SVK	
	HKG	HKG	BEL	BEL	NLD	NLD	PRK	IDN	
	SGP	RUS	AUT	AUT	SVK	SVK		ESP	
	BRA	SGP	CHE	CHE	ESP	IDN		GBR	
	RUS	BRA	JPN	JPN	ROU	ROU		FRA	
Panel B: Japan									
	CHN	CHN	CHN	CHN	CHN	CHN	GBR	KHM	CHN
	USA	USA	USA	USA	ITA	ITA	PRT	MMR	ITA
	KOR	KOR	AUS	SAU	KHM	KHM	BRA	BGD	VNM
	AUS	AUS	IDN	ARE	VNM	VNM	MAR	MEX	IDN
	ITA	ITA	KOR	AUS	IDN	IDN	IND	LAO	ESP
	CAN	FRA	DEU	IDN	MMR	MMR	CHE	NPL	FRA
	DEU	CAN	THA	KOR	BGD	BGD	HUN	LBN	DEU
	FRA	DEU	MYS	QAT	ESP	ESP	SVK		THA
	VNM	VNM	ARE	DEU	FRA	FRA	LKA		USA
	DNK	DNK	SAU	THA	DEU	DEU	AUT		KOR
Panel C: USA									
	CHN	CHN	CAN	CAN	CHN	CHN	PRT	MEX	CHN
	JPN	JPN	CHN	CHN	ITA	ITA	SVK	CAN	ITA
	DEU	DEU	MEX	MEX	BRA	BRA	POL	DOM	BRA
	KOR	KOR	JPN	JPN	VNM	VNM	HKG	ISR	VNM
	ITA	ITA	DEU	DEU	IDN	IDN	HUN	MAR	IDN
	GBR	GBR	KOR	KOR	THA	THA	CHE	COL	THA
	FRA	FRA	GBR	VEN	MEX	MEX	ALB	SLV	ESP
	IND	IND	FRA	GBR	ESP	ESP	BGR	AUS	IND
	HKG	HKG	ITA	FRA	IND	IND	DNK	ZAF	FRA
	VNM	VNM	MYS	MYS	DOM	DOM	AUT	PER	DEU

Notes: This table illustrates the construction of our instrument, using as an example product code 6403 "Footwear; with outer soles of rubber, plastics, leather or composition leather and uppers of leather" in 2006. Columns 1-2 list the top exporters to three importing countries – USA, Germany and Japan – exporting under the MFN regime in periods $t = 2006$ and $t - 1 = 2005$. Columns 3-4 list the importing countries' major aggregate trading partners in these periods. Columns 5-6 list the major trading partners in product 6403. Columns 7-9 then list the main countries in the treatment, control and excluded group for imports of product 6403 to the three importing countries.

Table A3: HS Sections

Code	Name
7	PLASTICS AND ARTICLES THEREOF; RUBBER AND ARTICLES THEREOF
8	RAW HIDES AND SKINS, LEATHER, FURSKINS AND ARTICLES THEREOF; SADDLERY AND HARNESS; TRAVEL GOODS,
9	WOOD AND ARTICLES OF WOOD; WOOD CHARCOAL; CORK AND ARTICLES OF CORK; MANUFACTURES OF STRAW, OF ESPARTO OR OF OTHER PLAITING MATERIALS; BASKETWARE AND WICKERWORK
10	PULP OF WOOD OR OF OTHER FIBROUS CELLULOSIC MATERIAL; RECOVERED (WASTE AND SCRAP) PAPER OR PAPERBOARD; PAPER AND PAPERBOARD AND ARTICLES THEREOF
11	TEXTILES AND TEXTILE ARTICLES
13	ARTICLES OF STONE, PLASTER, CEMENT, ASBESTOS, MICA OR SIMILAR MATERIALS;
15	BASE METALS AND ARTICLES OF BASE METAL CERAMIC PRODUCTS; GLASS AND GLASSWARE
16	MACHINERY AND MECHANICAL APPLIANCES; ELECTRICAL EQUIPMENT; PARTS THEREOF; SOUND RECORDERS AND REPRODUCERS, TELEVISION IMAGE AND SOUND RECORDERS AND REPRODUCERS, AND PARTS AND ACCESSORIES OF SUCH ARTICLES ARTIFICIAL FLOWERS; ARTICLES OF HUMAN HAIR
18	OPTICAL, PHOTOGRAPHIC, CINEMATOGRAPHIC, MEASURING, CHECKING, PRECISION, MEDICAL OR SURGICAL INSTRUMENTS AND APPARATUS; CLOCKS AND WATCHES; MUSICAL INSTRUMENTS; PARTS AND ACCESSORIES THEREOF
20	MISCELLANEOUS MANUFACTURED ARTICLES
	Aggregated
1	LIVE ANIMALS; ANIMAL PRODUCTS
2	VEGETABLE PRODUCTS
3	ANIMAL OR VEGETABLE FATS AND OILS AND THEIR CLEAVAGE PRODUCTS HANDBAGS AND SIMILAR CONTAINERS; ARTICLES OF ANIMAL GUT (OTHER THAN SILK-WORM GUT)
4	PREPARED EDIBLE FATS;ANIMAL OR VEGETABLE WAXES PREPARED FOODSTUFFS; BEVERAGES, SPIRITS AND VINEGAR; TOBACCO AND MANUFACTURED TOBACCO SUBSTITUTES
5	MINERAL PRODUCTS
6	PRODUCTS OF THE CHEMICAL OR ALLIED INDUSTRIES
12	FOOTWEAR, HEADGEAR, UMBRELLAS, SUN UMBRELLAS, WALKING-STICKS, SEAT-STICKS, WHIPS, RIDING-CROPS AND PARTS THEREOF; PREPARED FEATHERS AND ARTICLES MADE THEREWITH;
14	NATURAL OR CULTURED PEARLS, PRECIOUS OR SEMI-PRECIOUS STONES, PRECIOUS METALS, METALS CLAD WITH PRECIOUS METAL AND ARTICLES THEREOF; IMITATION JEWELLERY; COIN
17	VEHICLES, AIRCRAFT, VESSELS AND ASSOCIATED TRANSPORT EQUIPMENT
19	ARMS AND AMMUNITION; PARTS AND ACCESSORIES THEREOF
21	WORKS OF ART, COLLECTORS' PIECES AND ANTIQUES

Notes: This table describes the 21 internationally compatible HS "Sections", which are groupings of HS product codes. We also list the 9 HS Sections that we aggregate in the main text into a Section 'aggregate', as there is insufficient variation in tariffs in these sections to estimate the elasticity. Figures 3 reports the elasticity estimates by section, and Figure A8 reports trade-weighted means and medians of section-specific elasticities.

Appendix B Robustness

Figure A8: Trade elasticities: Full Sample Pooled vs. Trade-Weighted Sectoral Averages

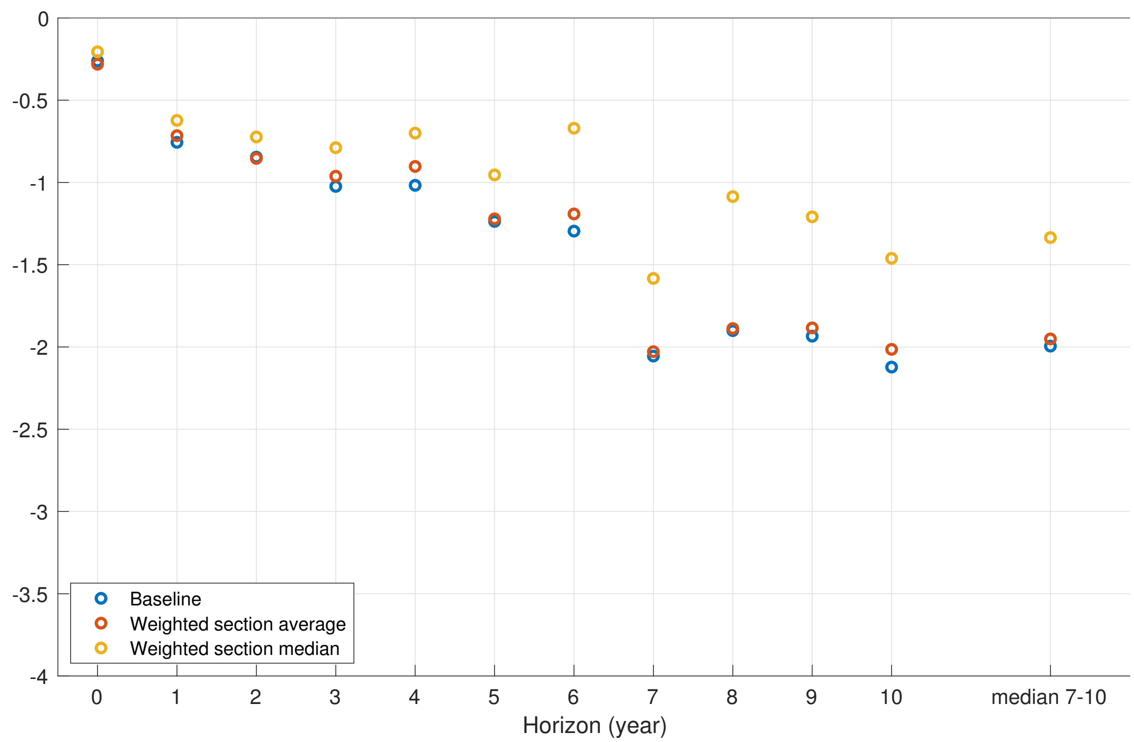

Notes: The blue circles reproduce the baseline elasticity point estimates depicted in Figure 2. The red circles display world trade-weighted means of the HS section-specific elasticities reported in Figure 3. The yellow circles display world trade-weighted medians of the HS section-specific elasticities reported in Figure 3. Weighting uses the 2006 shares of world trade, and excludes the estimates of the combined HS aggregate section as described in the text.

Figure A9: Trade Elasticity: Comparing Bilateral Fixed Effects

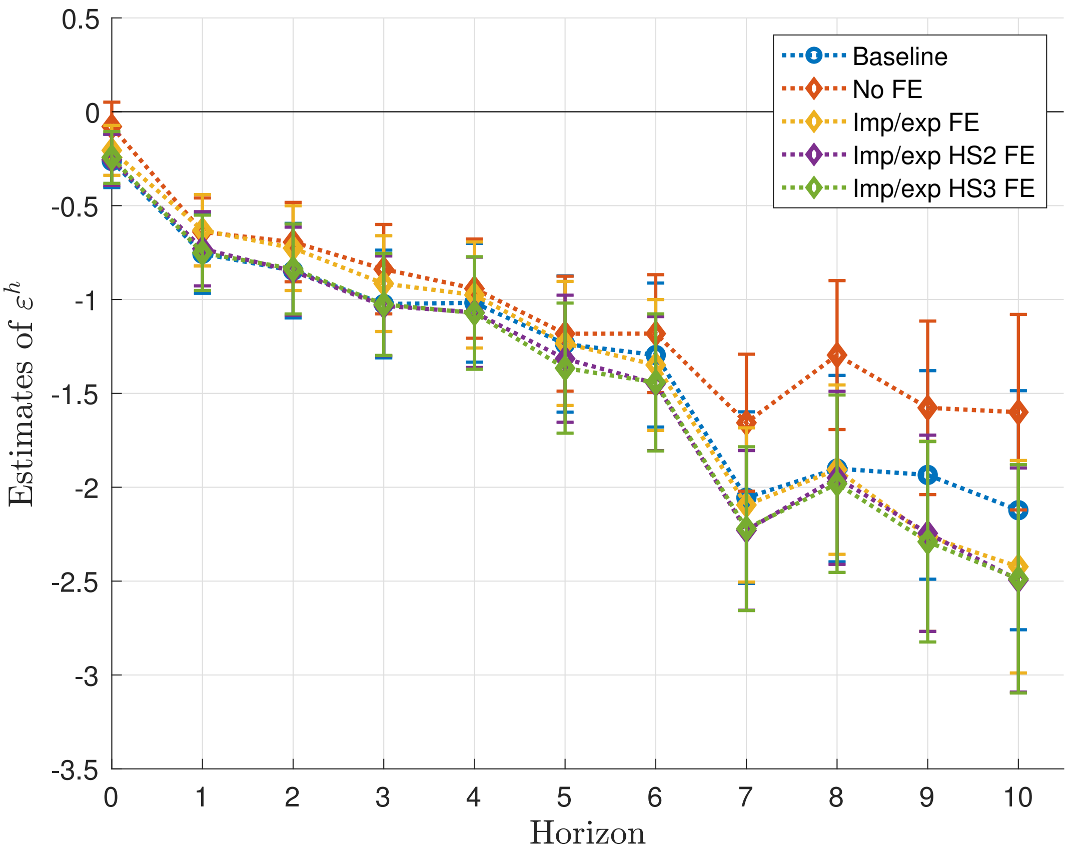

Notes: This figure displays estimates of the trade elasticity based on specification (2.4), with the baseline instrument (2.5) and including one lag of the changes in tariffs and trade as pre-trend controls. The bilateral fixed effects are either importer-exporter-HS4 (the baseline), importer-exporter-HS3, importer-exporter-HS2, importer-exporter or no bilateral fixed effects. The bars display 95% confidence intervals. Standard errors are clustered at the bilateral country-pair-product level.

Figure A10: Gains From Trade: Multiple Sectors

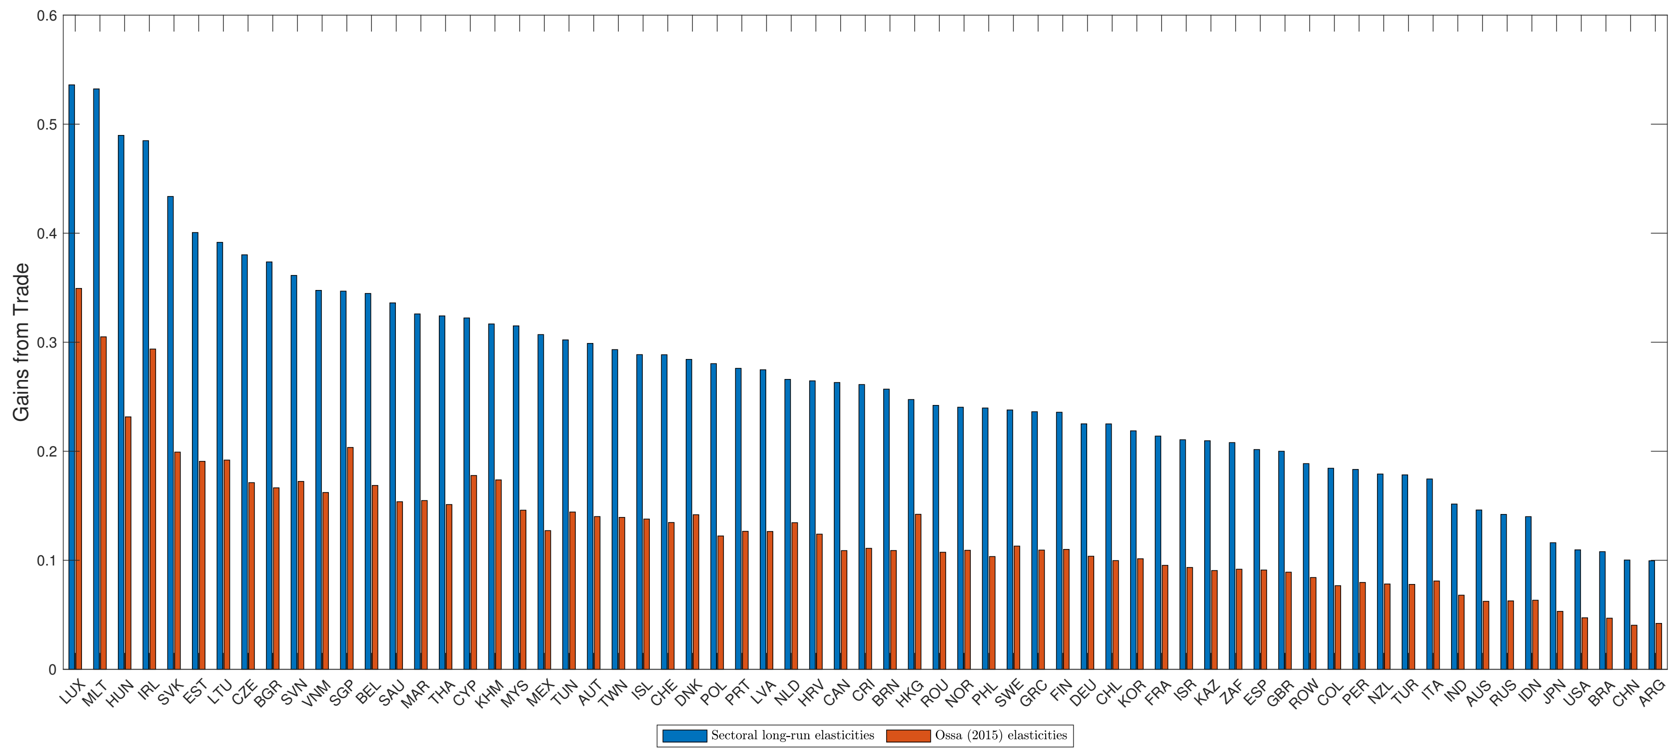

Notes: Gains from trade relative to autarky are computed using the formula $1 - \sum_s \lambda_{jj,s}^{-\beta_{j,s}/\theta_s}$, where $\beta_{j,s}$ is the share of sector s in country j's total absorption and $\lambda_{jj,s}$ is 1 minus the import share in sector s. The numbers for China and Mexico include export-processing activities (China) and global manufacturing activities (Mexico). "Sectoral long-run elasticities" refer to the HS-section level elasticities estimated in Section 4.1. We use the median estimate between years 7-10 for each section as the long-run value. For a comparison, the red bars use elasticities obtained from Ossa (2015). Data are from the OECD IO tables for 64 countries in year 2006. The OECD IO tables are converted to HS classification using an OECD's concordance between ISIC and HS. The GTAP sector estimates from Ossa (2015) are converted to the HS classification using GTAP's concordance table between GTAP sectors and HS classifications. The number of HS-6 categories in each GTAP-HS section pair is used as a weight.

Table A4: Robustness: Local Projections

	Panel A: Tariffs			Panel B: Trade		
	Baseline	Zero Lag	Five Lags	Baseline	Zero Lag	Five Lags
	(1)	(2)	(3)	(4)	(5)	(6)
$t-6$	-0.099***	-0.090***	-0.102***	0.097	0.028	0.305*
	(0.006)	(0.005)	(0.007)	(0.130)	(0.110)	(0.165)
$t-5$	-0.021***	-0.031***	.	0.217	0.275***	.
	(0.005)	(0.004)	.	(0.121)	(0.105)	.
$t-4$	-0.038***	-0.020***	.	0.004	-0.055	.
	(0.005)	(0.004)	.	(0.111)	(0.093)	.
$t-3$	-0.053***	-0.088***	.	0.075	-0.018	.
	(0.005)	(0.004)	.	(0.102)	(0.087)	.
$t-2$	-0.133***	-0.034***	.	0.242***	0.128	.
	(0.004)	(0.004)	.	(0.089)	(0.089)	.
$t-1$	.	-0.309***	.	.	0.149**	.
	.	(0.004	.	.	(0.072)	.
t	.	.	.	-0.262***	-0.147***	0.166
	.	.	.	(0.072)	(0.054)	(0.138)
$t+1$	0.890***	0.851***	0.842***	-0.673***	-0.535***	-0.109
	(0.004)	(0.003)	(0.006)	(0.096)	(0.069)	(0.175)
$t+2$	0.846***	0.830***	0.788***	-0.716***	-0.588***	-0.129
	(0.005)	(0.003)	(0.007)	(0.109)	(0.079)	(0.206)
$t+3$	0.831***	0.818***	0.769***	-0.851***	-0.758***	-0.481**
	(0.005)	(0.004)	(0.008)	(0.122)	(0.086)	(0.242)
$t+4$	0.821***	0.813***	0.747***	-0.835***	-0.747***	-0.214
	(0.006)	(0.004)	(0.009)	(0.132)	(0.094)	(0.271)
$t+5$	0.807***	0.823***	0.718***	-0.998***	-0.915***	-0.823***
	(0.005)	(0.004)	(0.010)	(0.149)	(0.102)	(0.308)
$t+6$	0.776***	0.786***	0.663***	-1.006***	-0.875***	-0.496**
	(0.006)	(0.004)	(0.011)	(0.152)	(0.101)	(0.320)
$t+7$	0.694***	0.754***	0.594***	-1.425***	-1.146***	-1.383***
	(0.006)	(0.004)	(0.011)	(0.162)	(0.110)	(0.352)
$t+8$	0.668***	0.716***	0.552***	-1.269***	-1.145***	-0.932**
	(0.007)	(0.005)	(0.012)	(0.169)	(0.117)	(0.391)
$t+9$	0.699***	0.733***	0.634***	-1.351***	-0.990***	-1.366***
	(0.007)	(0.005)	(0.015)	(0.198)	(0.120)	(0.513)
$t+10$	0.715***	0.720***	0.641***	-1.517***	-1.054***	-1.635**
	(0.008)	(0.005)	(0.018)	(0.232)	(0.139)	(0.649)

Notes: This table presents the results from estimating the local projections equations (2.3) (Panel A) and (2.2) (Panel B). The first column in each panel presents the baseline local projects results, while the second and third columns in each panel present results with 2 and 5 lags of tariffs and trade as pre-trend controls respectively. Standard errors clustered by country-pair-product are in parentheses. ***, ** and * denote significance at the 99, 95 and 90% levels.

Table A5: Trade Elasticity: Estimates and First Stage F-Statistics

	Baseline IV	F-stat	Distributed Lag	SW F-stat
t	-0.262***	91422	-0.413	20536
	(0.072)		(0.338)	
$t+1$	-0.756***	42231	-0.523	18831
	(0.108)		(0.466)	
$t+2$	-0.846***	36469	-0.920	22499
	(0.129)		(0.586)	
$t+3$	-1.024***	28537	-1.578**	22570
	(0.146)		(0.685)	
$t+4$	-1.017***	23771	-1.598**	16227
	(0.161)		(0.773)	
$t+5$	-1.237***	22697	-2.097**	12463
	(0.185)		(0.861)	
$t+6$	-1.296***	19439	-2.185**	14722
	(0.196)		(0.931)	
$t+7$	-2.055***	15481	-2.710***	13473
	(0.233)		(1.016)	
$t+8$	-1.901***	13933	-2.798**	13475
	(0.253)		(1.109)	
$t+9$	-1.934***	10201	-3.084***	14278
	(0.283)		(1.180)	
$t+10$	-2.122***	8252	-3.166**	10962
	(0.325)		(1.252)	

Notes: This table presents the first-stage F-statistics for the main estimates. For the Distributed Lag model we report the Sanderson-Windmeijer F-statistic to test for weak instruments as we have 11 instruments and 11 endogenous variables.

Table A6: "Traditional Gravity" Elasticity Estimates in Differences

	No FE	Multilateral FE	Multilateral + Bilateral FE	OLS	Baseline IV	
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Baseline Sample		Log Levels Gravity			5-year Differences	
Elasticity	-7.01***	-11.26***	-0.95***	-0.77***	-1.18***	-1.11***
	(0.05)	(0.11)	(0.05)	(0.03)	(0.11)	(0.12)
R^2	0.02	0.42	0.59	0.19		
obs	21.13	21.13	21.13	21.13	21.13	21.13
First stage F					38220	36815
Panel B: Shorter horizon					3-year Differences	
Elasticity				-0.48***	-1.01***	-0.93***
				(0.02)	(0.09)	(0.10)
R^2				0.15		
obs				46.60	26.43	26.16
First stage F					56290	47082
Panel C: Longer horizon					7-year Differences	
Elasticity				-0.86***	-1.54***	-1.52***
				(0.02)	(0.13)	(0.15)
R^2				0.21		
obs				30.8	17.19	16.95
First stage F					34952	31202
Fixed effects						
Imp×HS4×year, Exp×HS4×year		Yes	Yes	Yes	Yes	Yes
Imp×Exp×HS4			Yes			Yes

Notes: This table presents the results from estimating the trade elasticity in log-levels and in differences. The dependent variable in columns 1-3 is the log of trade values, and in columns 4-5 is the log-difference in trade value, over 5 years on the constant baseline IV sample for $h = 5$ (Panel A), 3 years (Panel B), and 7 years (Panel C). Columns 1-3 report the OLS "conventional gravity" results, column 4 a differenced specification, and columns 5 and 6 instrumenting with the baseline instrument, where the fixed effects vary. Standard errors, clustered at the importer-exporter-HS4 level, are in parentheses. ***, ** and * denote significance at the 99, 95 and 90% levels. Number of observations reported in millions. Differenced specifications do not have pretrend controls for comparability with the log-levels specifications.

Table A7: "Traditional Gravity" Elasticity Estimates in Levels, HS6×year Multilateral Resistance Terms

	No Bilateral	Country-pair	Country-pair×HS2	Country-pair×HS3	Country-pair×HS4
	(1)	(2)	(3)	(4)	(5))
$\tau_{i,j,p,t}$	-12.62***	-1.70***	-1.36***	-1.28***	-1.04***
SE	(0.08)	(0.03)	(0.02)	(0.02)	(0.02)
R^2	0.55	0.67	0.70	0.71	0.76
N	103.1	103.1	102.9	102.7	101.7
Bilateral Fixed Effects					
	None	Imp×Exp	Imp×Exp×HS2	Imp×Exp×HS3	Imp×Exp×HS4

Notes: This table presents the results from estimating the trade elasticity in log-levels where the multilateral resistance terms are HS6×year. The dependent variable is the log of trade value. Column 1 reports the results with no fixed effects. Column 2 adds country-pair effects, Column 3 includes country-pair-HS2 bilateral fixed effects , column 4 includes country-pair-HS3 effects and Column 4 uses country-pair-HS4 fixed effects. Standard errors, clustered at the importer-exporter-HS4 level, are in parentheses. ***, ** and * denote significance at the 99, 95 and 90% levels. Number of observations are reported in millions.

Table A8: Trade Elasticity, Every Horizon:, Robustness: Pre-Trends, Alternative Clustering, Alternative Samples

	Baseline	Zero Lag	Five Lags	FE50	Two-way Clustering	Balanced Panel	Alternative Control Group
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
t	-0.262***	-0.147***	0.166	-0.232**	-0.262***	-0.592**	-0.187**
	(0.072)	(0.054)	(0.138)	(0.094)	(0.096)	(0.290)	(0.077)
obs	31.66	41.46	14.58	17.62	31.66	5.04	27.33
$t+1$	-0.756***	-0.628***	-0.129	-0.599***	-0.756***	-0.098	-0.490***
	(0.108)	(0.081)	(0.208)	(0.133)	(0.141)	(0.380)	(0.117)
obs	26.18	32.85	12.52	15.19	26.18	5.04	22.63
$t+2$	-0.846***	-0.708***	-0.164	-0.724***	-0.846***	-0.755*	-0.449***
	(0.129)	(0.095)	(0.262)	(0.156)	(0.189)	(0.444)	(0.141)
obs	23.27	29.21	11.09	13.77	23.27	5.04	20.05
$t+3$	-1.024***	-0.926***	-0.625**	-0.865***	-1.024***	-0.895*	-0.743***
	(0.146)	(0.105)	(0.315)	(0.173)	(0.195)	(0.472)	(0.160)
obs	20.8	26.16	9.76	12.48	20.8	5.04	17.86
$t+4$	-1.017***	-0.918***	-0.287	-0.784***	-1.017***	-0.769	-0.613***
	(0.161)	(0.115)	(0.362)	(0.187)	(0.232)	(0.480)	(0.171)
obs	18.67	23.49	8.53	11.33	18.67	5.04	16
$t+5$	-1.237***	-1.112***	-1.146***	-1.012***	-1.237***	-0.916**	-0.792***
	(0.185)	(0.124)	(0.429)	(0.215)	(0.253)	(0.437)	(0.201)
obs	16.69	21.13	7.3	10.22	16.69	5.04	14.27
$t+6$	-1.296***	-1.113***	-0.747	-1.051***	-1.296***	-0.532	-0.570***
	(0.196)	(0.129)	(0.482)	(0.231)	(0.274)	(0.463)	(0.216)
obs	14.92	18.9	6.17	9.2	14.92	5.04	12.62
$t+7$	-2.055***	-1.521***	-2.330***	-1.853***	-2.055***	-0.990**	-1.383***
	(0.233)	(0.145)	(0.595)	(0.270)	(0.357)	(0.489)	(0.251)
obs	13.22	16.95	5.25	8.2	13.22	5.04	11.12
$t+8$	-1.901***	-1.599***	-1.690**	-1.888***	-1.901***	-1.094**	-1.079***
	(0.253)	(0.163)	(0.709)	(0.289)	(0.465)	(0.510)	(0.275)
obs	11.53	15.02	4.42	7.19	11.53	5.04	9.63
$t+9$	-1.934***	-1.351***	-2.155***	-1.781***	-1.934***	-1.600***	-1.087***
	(0.283)	(0.164)	(0.812)	(0.323)	(0.500)	(0.551)	(0.306)
obs	9.85	13.11	3.79	6.19	9.85	5.04	8.2
$t+10$	-2.122***	-1.463***	-2.550**	-1.760	-2.122***	-1.818***	-1.600***
	(0.325)	(0.194)	(1.016)	(0.374)	(0.332)	(0.544)	(0.379)
obs	8.31	11.25	3.21	5.25	8.31	5.04	6.84

Notes: This table presents robustness exercises for the results from estimating equation (2.4). All specifications include importer-HS4-year, exporter-HS4-year and importer-exporter-HS4 fixed effects. Columns 2 and 3 vary the pretrend controls (including alternatively two lags or five lags of import growth and tariff changes). Column 4 reports the results when the sample is restricted to fixed-effects clusters with a minimum of 50 observations. Standard errors are clustered at the importer-exporter-HS4 level, except in Column 5 where they are additionally clustered by year. Column 6 restricts the sample to have positive trade flows for all time horizons. Column 7 reports results where the control group only contains observations with zero tariff changes. All columns except 2 and 3 include the baseline pretrend controls (one lag). ***, ** and * indicate significance at the 99, 95 and 90 percent level respectively. Observations are reported in millions.

Table A9: Trade Elasticity, Every Horizon, Robustness: Alternative Instruments, Outcomes, Fixed Effects, and Samples

	Baseline (1)	All Partners (2)	Quantities (3)	Unit Values (4)	Extensive (5)	Extensive Sel (6)	DL (7)	SD1 (8)
t	-0.262***	-0.275***	-0.434***	-0.179*	-0.049	0.025	-0.082***	-0.313***
	(0.072)	(0.030)	(0.123)	(0.093)	(0.057)	(0.039)	(0.063)	(0.096)
obs	31.66	57.14	29.3	31.66	31.66	131.03	56.35	28.7
$t+1$	-0.756***	-0.624***	-1.228***	-0.664***	-0.027	-0.484***	-0.805***	-0.889***
	(0.108)	(0.047)	(0.174)	(0.136)	(0.080)	(0.063)	(0.089)	(0.142)
obs	26.18	47.19	24.26	26.18	26.18	108.13	49.11	23.76
$t+2$	-0.846***	-0.647***	-1.550***	-0.663***	-0.139	-0.432***	-0.731***	-0.861***
	(0.129)	(0.053)	(0.213)	(0.162)	(0.094)	(0.073)	(0.098)	(0.174)
obs	23.27	42.29	21.42	23.27	23.27	96.99	45.23	21.1
$t+3$	-1.024***	-0.648***	-1.669***	-0.81***	-0.134	-0.590***	-0.955***	-1.228***
	(0.146)	(0.061)	(0.241)	(0.184)	(0.104)	(0.084)	(0.107)	(0.200)
obs	20.8	38.16	19.14	20.8	20.8	87.22	41.5	18.86
t+4	-1.017***	-0.657***	-1.637***	-0.801***	-0.09	-0.449***	-0.846***	–1.171***
	(0.161)	(0.068)	(0.264)	(0.203)	(0.115)	(0.095)	(0.117)	(0.219)
obs	18.67	34.39	17.14	18.67	18.67	77.82	37.92	16.92
$t+5$	-1.237***	-0.718***	-1.686***	-1.422***	0.291**	-0.730***	-1.183***	-1.180***
	(0.185)	(0.073)	(0.285)	(0.232)	(0.128)	(0.100)	(0.122)	(0.253)
obs	16.69	30.89	15.32	16.69	16.69	69.73	34.61	15.12
$t+6$	-1.296***	-0.781***	-2.477***	-1.374***	0.125	-0.864***	-1.212***	-1.267***
	(0.196)	(0.078)	(0.296)	(0.244)	(0.135)	(0.101)	(0.124)	(0.261)
obs	14.92	27.65	13.64	14.92	14.92	62.07	31.34	13.51
$t+7$	-2.055***	-0.940***	-3.756***	-2.166***	0.162	-0.902***	-1.496***	-1.990***
	(0.233)	(0.091)	(0.401)	(0.292)	(0.161)	(0.114)	(0.138)	(0.318)
obs	13.22	24.64	12.06	13.22	13.22	55.24	28.4	11.96
$t+8$	-1.901***	-1.000***	-3.635***	-2.076***	0.205	-0.688***	-1.429***	-1.955***
	(0.253)	(0.099)	(0.430)	(0.318)	(0.173)	(0.130)	(0.157)	(0.352)
obs	11.53	21.67	10.48	11.53	11.53	48.51	25.4	10.44
$t+9$	-1.934***	-0.970***	-3.149***	-1.656***	-0.180	-1.000***	-1.650***	-2.185***
	(0.283)	(0.109)	(0.451)	(0.353)	(0.193)	(0.137)	(0.161)	(0.388)
obs	9.85	18.69	8.8	9.85	9.85	41.82	22.36	8.93
$t+10$	-2.122***	-0.866***	-1.865***	-1.765***	-0.0780	-0.941***	-1.638***	-2.360***
	(0.325)	(0.122)	(0.453)	(0.406)	(0.221)	(0.153)	(0.181)	(0.444)
obs	8.31	15.92	7.54	8.31	8.31	35.09	19.24	7.52

Notes: This table presents alternative estimates for the results from estimating equation (2.4), varying the instrument or outcome variable. Column 2 uses an alternative definition of the instrument where all trade partners subject to the MFN regime are included. Column 3 reports results for quantities, and Column 4 the results for unit values. Column 5 presents results including the extensive margin using the inverse hyperbolic sine transformation for trade flows, and all zero trade observations for importer-exporter-section pair with ever positive trade. Column 6 presents results including the extensive margin using the inverse hyperbolic sine transformation for trade flows, and only zero trade observations when trade switches from zero to positive, or vice versa. Column 7 presents results from a distributed lag model. Column 8 reports the results based on a sample where tariffs do not vary within an importer-HS6. All specifications include importer-HS4-year, exporter-HS4-year and importer-exporter-HS4 fixed effects. All columns include the baseline pretrend controls (one lag). Standard errors are clustered at the importer-exporter-HS4 level. ***, ** and * indicate significance at the 99, 95 and 90 percent level respectively. Observations are reported in millions.

Table A10: Trade Elasticity: Further Robustness

	Uruguay Round		HS6 Multilateral Effects	
	All data/all tariffs 2SLS	Baseline IV	All data/all tariffs 2SLS	Baseline IV
	(1)	(2)	(3)	(4)
t	-0.285	-0.181	-0.186***	0.077
	(0.179)	(0.851)	(0.014)	(0.106)
obs	0.9	0.55	54.3	30.2
$t+1$	-1.157**	-2.011	-0.261***	-0.682***
	(0.581)	(2.392)	(0.022)	(0.134)
obs	0.84	0.51	54.3	24.9
$t+2$	-0.771	-1.493	-0.278***	-0.599***
	(0.663)	(2.662)	(0.027)	(0.028)
obs	0.78	0.5	48.5	22.2
$t+3$	-1.489***	-5.094	-0.288***	-0.589***
	(0.563)	(2.469)	(0.030)	(0.173)
obs	0.78	0.47	43.7	19.8
$t+4$	-1.608***	-1.434***	-0.373***	-0.767***
	(0.592)	(1.960)	(0.032)	(0.190)
obs	0.77	0.46	39.5	17.8
$t+5$	-1.287*	-2.077	-0.406***	-1.041***
	(0.683)	(2.462)	(0.034)	(0.219)
obs	0.71	0.43	35.6	15.9
$t+6$	-1.428**	-1.152	-0.446***	-1.090***
	(0.644)	(2.306)	(0.036)	(0.219)
obs	0.71	0.43	31.9	14.2
$t+7$	-1.533*	0.260	-0.429***	-1.371***
	(0.831)	(2.905)	(0.041)	(0.266)
obs	0.7	0.42	28.6	12.6
$t+8$	-1.848**	-4.740	-0.353***	-0.993***
	(0.814)	(2.928)	(0.045)	(0.287)
obs	0.74	0.46	24.4	11.0
$t+9$	-1.041	-3.610	-0.405***	-0.981***
	(0.801)	(2.515)	(0.044)	(0.330)
obs	0.74	0.46	22.2	9.45
$t+10$	-0.296	-2.967	-0.495***	-0.468
	(0.991)	(3.278)	(0.043)	(0.355)
obs	0.65	0.4	19.0	7.95

Notes: This table presents the results from estimating the trade elasticity using both all data/all tariffs 2SLS (column 1) and the baseline instrument (column 2) for tariff changes only in years 1995-1997 ("Uruguay round"). Columns (4) and (5) present the all data/all tariffs 2SLS and baseline IV specifications when the multilateral resistance terms are country-HS6-year level. In these columns we drop the bilateral fixed effect. Standard errors are clustered at the importer-exporter-HS4 level. All columns include the baseline pretrend controls (one lag). ***, ** and * indicate significance at the 99, 95 and 90 percent level respectively. Observations are reported in millions.

Table A11: Gains from Trade

Country	$\theta = -1$	$\theta = -5$	$\theta = -10$
G7			
Canada	21.0%	3.9%	1.9%
France	15.6%	3.0%	1.5%
Germany	21.7%	3.8%	1.9%
Italy	13.9%	2.7%	1.3%
Japan	7.7%	1.5%	0.7%
UK	13.9%	2.6%	1.3%
US	5.3%	1.0%	0.5%
Major Emerging Markets			
Brazil	7.4%	1.4%	0.7%
China	11.9%	2.1%	1.1%
India	10.5%	2.0%	1.0%
Mexico	6.1%	1.2%	0.6%
Russia	19.5%	3.6%	1.8%
South Africa	14.1%	2.7%	1.3%
Median, 64 Countries	22.9%	4.2%	2.1%

Notes: Data are from the OECD IO tables for 64 countries in year 2006. Gains from trade relative to autarky are computed using the formula $\lambda_{jj}^{1/\theta}$, where λ_{jj} is 1 minus the import share. The import share is calculated as imports divided by gross output. The numbers for China and Mexico include export-processing activities (China) and global manufacturing activities (Mexico).

Appendix C Model

Notation Throughout this appendix, we let tildes denote percent deviations from steady state, e.g. $\tilde{v}_t = \ln v_t - \ln v = d\ln v_t = \frac{v_t - v}{v}$. Variables without subscripts denote steady state values.

For most of this appendix we suppress source and destination country as well as product subscripts for convenience. For clarity we provide an overview on the notation here:

- D_t is a destination and product-specific demand shock, i.e. $D_t = D_{i,p,t}$
- ω_t is a source, destination, and product-specific demand shock, i.e. $\omega_t = \omega_{i,j,p,t}$
- c_t denotes marginal cost and is specific to the source country and product, i.e. $c_t = c_{j,p,t}$
- κ_t denotes non-tariff trade barriers and varies with the source country, the destination country, and by product, i.e. $\kappa_t = \kappa_{i,j,p,t}$
- τ_t is the country-pair and product-specific iceberg tariff, i.e. $\tau_t = \tau_{i,j,p,t}$

C.1 Model summary

The following system of equations characterizes the trade response to tariff shocks. The first set of equations is

$$p_t^x = p^x\left(c_t, \kappa_t\tau_t, \omega_t D_t\right), \tag{C.1}$$

$$q_t = q\left(p_t^x, \kappa_t\tau_t, \omega_t D_t\right), \tag{C.2}$$

$$\pi_t = \pi\left(c_t, \kappa_t\tau_t, \omega_t D_t\right), \tag{C.3}$$

$$X_t = q_t p_t^x n_t, \tag{C.4}$$

where p_t^x is the price of exported goods, q_t a quantity measure, π_t a measure of flow profits, X_t is exports exclusive of tariffs, and n_t a generic mass. Let further v_t denote a generic value. The following dynamic system determines the evolution of v_t and n_t,

$$v_t = \frac{1}{1+r}\mathbb{E}_t\left[\pi\left(c_{t+1}, \kappa_{t+1}\tau_{t+1}, \omega_{t+1}D_{t+1}\right) + (1-\delta)\, v_{t+1}\right], \tag{C.5}$$

$$n_t = n_{t-1}(1-\delta) + G\left(v_{t-1}\right), \tag{C.6}$$

together with $\lim_{t\to\infty}\left(\frac{1-\delta}{1+r}\right)^t v_t = 0$, a given initial value for n_0, and stochastic processes for c_t, κ_t, τ_t, ω_t and D_t, which are exogeneous in our partial equilibrium model.

We define the following constants

$$\eta_{q,p} := \frac{\partial \ln q}{\partial \ln p^x},\ \eta_{q,\tau} := \frac{\partial \ln q}{\partial \ln \tau},\ \eta_{p,\tau} := \frac{\partial \ln p^x}{\partial \ln \tau},\ \eta_{\pi,\tau} := \frac{\partial \ln \pi}{\partial \ln \tau}, \tag{C.7}$$

and assume that $\eta_{q,p} < 0$, $\eta_{q,\tau} < 0$, and $\eta_{\pi,\tau} < 0$. We also define $\chi := \frac{g(v)v}{G(v)}$ where G introduced below, and $g = G'$.

C.2 Microfoundations

We next show that three different frameworks generate the above system of equations.

C.2.1 A dynamic Arkolakis (2010) model

This model is a dynamic extension of the Arkolakis (2010) market penetration framework, where the number of customers adjusts gradually. The model also shares features with Fitzgerald, Haller, and Yedid-Levi (2019), and others.

A single representative firm sells its good in the foreign location, earning profits $\Pi_t = n_t \pi\left(c_t, \kappa_t \tau_t, \omega_t D_t\right)$. Here, n_t denotes the mass of foreign consumers that the firm reaches in the foreign location. Further, $\pi\left(c_t, \kappa_t \tau_t, \omega_t D_t\right)$ denotes flow profits per unit mass of foreign consumers reached, and is a function of the exporter's costs c_t, tariffs τ_t, nontariff trade costs κ_t, and the demand shifters $\omega_t D_t$.

The mass of foreign consumers available for the firm to sell to evolves according to the accumulation equation

$$n_{t+1} = n_t \left(1-\delta\right) + a_t, \tag{C.8}$$

where a_t is the mass of newly added customers in the foreign country. Note that mass n_t is predetermined in the current period, so that adding new consumers this period only affects next period's mass of consumers n_{t+1}. We assume that adding a_t new customers requires a payment of $f(a_t)$, where $f' > 0$, $f'' > 0$, $\lim_{a \to 0} f'\left(a\right) = 0$, $\lim_{a \to \infty} f'\left(a\right) = \infty$, and that the existing mass of consumers already reached by the firm depreciates at rate δ.

The firm discounts at interest rate r and maximizes the present discounted value of future profits,

$$\max_{\{a_t\}_{t=0}^{\infty}} \mathbb{E}_0 \sum_{t=0}^{\infty} \left(\frac{1}{1+r}\right)^t \left[n_t \pi\left(c_t, \kappa_t \tau_t, \omega_t D_t\right) - f\left(a_t\right)\right].$$

Denoting by v_t the multiplier on constraint (accumulation equation), the current value Lagrangian is

$$\mathcal{L} = \mathbb{E}_0 \sum_{t=0}^{\infty} \left(\frac{1}{1+r}\right)^t \left[n_t \pi\left(c_t, \kappa_t \tau_t, \omega_t D_t\right) - f\left(a_t\right) + v_t\left(n_t\left(1-\delta\right) + a_t - n_{t+1}\right)\right].$$

The first order necessary conditions are

$$\begin{aligned} f'\left(a_t\right) &= v_t, \\ v_t &= \frac{1}{1+r} \mathbb{E}_t\left[\pi\left(c_{t+1}, \kappa_{t+1} \tau_{t+1}, \omega_{t+1} D_{t+1}\right) + \left(1-\delta\right) v_{t+1}\right], \end{aligned}$$

and the transversality condition $\lim_{t\to\infty}\left(\frac{1}{1+r}\right)^t v_t n_t = 0$, which implies that $\lim_{t\to\infty}\left(\frac{1-\delta}{1+r}\right)^t v_t = 0$. The firm chooses its investment into accumulating new consumers such that the marginal benefit v_t equals the marginal cost $f'(a_t)$. The shadow value v_t, in turn, is the expected present value of profits generated by each consumer reached in the foreign market.

Note that the above problem is reminiscent of a standard investment problem with convex adjustment costs, except that flow profits are a linear function of n_t, the analogue of the capital stock. This linearity greatly improves the tractability of the problem and permits analytical solutions.

Letting $q_t = q\left(p_t^x, \kappa_t\tau_t, \omega_t D_t\right)$ denote foreign demand per unit mass of consumers, and letting $p_t^x = p^x\left(c_t, \kappa_t\tau_t\right)$ denote the price set by the representative firm, exports are $X_t = q_t p_t^x n_t$. After substituting out a_t, the accumulation equation (C.8) becomes

$$n_t = n_{t-1}\left(1-\delta\right) + \left(f'\right)^{-1}\left(v_{t-1}\right).$$

For $G \equiv \left(f'\right)^{-1}$, the model is described by the set of equations in Section C.1.

C.2.2 A dynamic Krugman (1980) model

We next present a dynamic partial equilibrium version of the Krugman (1980) model. The model also shares features with Costantini and Melitz (2007), Ruhl (2008), and many others.

There is a continuum of firms, and each exporting firm receives flow profits $\pi\left(c_t, \kappa_t\tau_t, \omega_t D_t\right)$ from exporting. Further, exporters exit the bilateral trade relationship with probability δ per period. The value of an exporting firm at the end of period t is

$$v_t = \frac{1}{1+r}E_t\left[\pi\left(c_{t+1}, \kappa_{t+1}\tau_{t+1}, \omega_{t+1}D_{t+1}\right) + \left(1-\delta\right)v_{t+1}\right],$$

where we assume that the value of a non-exporting firm is zero. We also require that $\lim_{t\to\infty}\left(\frac{1-\delta}{1+r}\right)^t v_t = 0$, which follows from the transversality condition of the firms' owner(s).

In every period, a unit mass of firms receives the opportunity to begin exporting to the foreign location. Each of these firms receive i.i.d sunk cost draw ξ_t^s, drawn from distribution G, and then decide whether to start exporting. Each firm solves

$$\max\left\{v_t - \xi_t^s, 0\right\},$$

so a firm enters if and only if $\xi_t^s \leq v_t$. Note that a firm entering this period begins to receive profits from exporting only in the next period. The mass of firms entering into exporting in period t is thus $G\left(v_t\right)$. The mass of exporting firms at the end of period t is denoted by n_t, and it evolves according to

$$n_{t+1} = n_t\left(1-\delta\right) + G\left(v_t\right).$$

Letting $q_t = q\left(p_t^x, \kappa_t\tau_t, \omega_t D_t\right)$ denote foreign demand per unit mass of firms, and letting $p_t^x = p^x\left(c_t, \kappa_t\tau_t\right)$ denote the price set by each firm, exports are $X_t = q_t p_t^x n_t$. It is clear that this model is nested by the set of equations in Section C.1.

C.2.3 A dynamic Melitz (2003) model

Consider a version of the Melitz (2003) model, with a two-stage entry problem. In the first stage of the entry problem, firms do not know their productivity of producing the exported good. Further, they pay a sunk cost to obtain the *right to export on a per-period basis.* Having paid this sunk cost, they learn their productivity and face the following static decision problem going forward: As long as the firm maintains its right to export on a per-period basis, it can pay a fixed cost to obtain the profit of exporing for one period.

First stage Let $\pi\left(c_t, \kappa_t\tau_t, \omega_t D_t\right)$ denote *expected* flow profits from exporting in stage one of the entry problem. The remainder of this stage is isomorphic to the dynamic Krugman (1980) model described above. Firms lose their right to export on a per-period basis with probability δ per period. The expected value of exporting at the end of period t is

$$v_t = \frac{1}{1+r} E_t\left[\pi\left(c_{t+1}, \kappa_{t+1}\tau_{t+1}, \omega_{t+1}D_{t+1}\right) + (1-\delta)\, v_{t+1}\right],$$

where we assume that the value of a non-exporting firm is zero. We also require that $\lim_{t\to\infty}\left(\frac{1-\delta}{1+r}\right)^t v_t = 0$, which follows from the transversality condition of the firms' owner(s).

In every period, a unit mass of firms faces the first stage of the entry problem. Each of these firms receives an i.i.d sunk cost draw ξ_t^s, drawn from distribution G, and then decides whether to enter into the second stage. Each firm solves

$$\max\left\{v_t - \xi_t^s, 0\right\},$$

so a firm enters if and only if $\xi_t^s \leq v_t$. Note that a firm entering this period faces the second stage of the entry problem only in the next period. The mass of firms entering into the second stage in period t is $G\left(v_t\right)$. The mass of firms with the right to export on a per-period basis is denoted by n_t, and evolves according to

$$n_{t+1} = n_t\left(1-\delta\right) + G\left(v_t\right).$$

Foreign consumer We assume that foreign demand takes the form $Q_t = \left(P_t^c\right)^{-\sigma}\omega_t D_t$, where $P_t^c = \kappa_t\tau_t P_t^x$ is the ideal price index of consumer prices paid for the country's exports, so that $Q_t = \left(P_t^x\right)^{-\sigma}\left(\kappa_t\tau_t\right)^{-\sigma}\omega_t D_t$ is the quantity aggregate of the firm-level exports. Q_t takes the CES form

$$Q_t = \left(\int_{j\in J_t} q_t\left(j\right)^{\frac{\sigma-1}{\sigma}} dj\right)^{\frac{\sigma}{\sigma-1}}, \tag{C.9}$$

where j indexes exporting firms and J_t is the set of exporting firms. Profit maximization implies that

$$q_t\left(j\right) = Q_t\left(\frac{p_t^x\left(j\right)}{P_t^x}\right)^{-\sigma}, \tag{C.10}$$

where

$$P_t^x = \left(\int_{j\in J_t}\left(p_t^x\left(j\right)\right)^{1-\sigma} dj\right)^{\frac{1}{1-\sigma}}. \tag{C.11}$$

Measured exports are $X_t = Q_t P_t^x$.

Second stage Once a firm has paid the sunk entry cost, it draws its productivity φ from distribution F, which we assume to be independent of the sunk cost draw ζ_t^s. A firm's marginal costs are

$\frac{c_t}{\varphi}$. Each firm faces demand function (C.10). Profit maximization implies that

$$p_t^x(j) = \frac{\sigma}{\sigma-1}\frac{c_t}{\varphi(j)},$$

and yields flow profits from exporting

$$\begin{aligned}\pi_t(j) &= q_t(j)\left(p_t^x(j) - \frac{c_t}{\varphi(j)}\right) - \xi \\ &= \frac{1}{\sigma}\left(\frac{\sigma}{\sigma-1}\frac{c_t}{\varphi(j)}\right)^{1-\sigma} Q_t\left(P_t^x\right)^{\sigma} - \xi,\end{aligned}$$

where ξ denotes the per-period fixed cost of exporting, which is common across firms.

A firm exports in period t if $\pi_t(j) \geq 0$, and the marginal firm has productivity

$$\varphi_t^m = \frac{\sigma}{\sigma-1}c_t\left(\frac{\sigma\xi}{Q_t\left(P_t^x\right)^{\sigma}}\right)^{\frac{1}{\sigma-1}}.$$

Note that Q_t and P_t^c depend on τ_t and hence changes in tariffs will affect the composition of firms that exports in a given period.

Following Melitz (2003), we write the price index (C.11) as

$$\begin{aligned}P_t^x &= \left(\int_{\varphi_t^m}^{\infty}\left(p_t^x(\varphi)\right)^{1-\sigma} n_t dF(\varphi)\right)^{\frac{1}{1-\sigma}} \\ &= n_t^{\frac{1}{1-\sigma}}\frac{\sigma}{\sigma-1}c_t\left(\int_{\varphi_t^m}^{\infty}\varphi^{\sigma-1}dF(\varphi)\right)^{\frac{1}{1-\sigma}} \\ &= n_t^{\frac{1}{1-\sigma}}\frac{\sigma}{\sigma-1}\frac{c_t}{\tilde{\varphi}_t}\end{aligned}$$

where

$$\tilde{\varphi}_t = \left(\int_{\varphi_t^m}^{\infty}\varphi^{\sigma-1}dF(\varphi)\right)^{\frac{1}{\sigma-1}}. \tag{C.12}$$

Note that $\tilde{\varphi}_t$ denotes an *aggregate* productivity measure of exporting firms, and not an *average*.

Now letting

$$p_t^x(\tilde{\varphi}_t) = \frac{\sigma}{\sigma-1}\frac{c_t}{\tilde{\varphi}_t}, \tag{C.13}$$

we have

$$P_t^x = n_t^{\frac{1}{1-\sigma}} p_t^x(\tilde{\varphi}_t).$$

Again following Melitz (2003), and noting that $q_t(\varphi) = Q_t\left(\frac{p_t^x(\varphi)}{P_t^x}\right)^{-\sigma}$ and

$$q_t(\tilde{\varphi}_t) = Q_t\left(\frac{p_t^x(\tilde{\varphi}_t)}{P_t^x}\right)^{-\sigma}, \tag{C.14}$$

we have that $q_t(\varphi) = \left(\frac{\varphi_t}{\tilde{\varphi}_t}\right)^{\sigma} q(\tilde{\varphi}_t)$. We can then write the quantity index (C.9) as

$$\begin{aligned} Q_t &= \left(\int_{\varphi_t^m}^{\infty} q_t(\varphi)^{\frac{\sigma-1}{\sigma}} n_t dF(\varphi)\right)^{\frac{\sigma}{\sigma-1}} \\ &= n_t^{\frac{\sigma}{\sigma-1}} q_t(\tilde{\varphi}_t). \end{aligned}$$

Now the total value of exports is

$$\begin{aligned} X_t &= Q_t P_t^x \\ &= n_t^{\frac{\sigma}{\sigma-1}} q(\tilde{\varphi}_t)\, n_t^{\frac{1}{1-\sigma}} p^x(\tilde{\varphi}_t) \\ &= n_t q_t(\tilde{\varphi}_t)\, p_t^x(\tilde{\varphi}_t), \end{aligned}$$

where $\tilde{\varphi}_t$, $p_t^x(\tilde{\varphi}_t)$, and $q_t(\tilde{\varphi}_t)$ are defined in equations (C.12), (C.13), and (C.14).

Lastly, expected profits can be written as

$$\pi_t = \frac{1}{\sigma} Q_t (P_t^x)^{\sigma} \left(\frac{\sigma}{\sigma-1}\frac{c_t}{\tilde{\varphi}_t}\right)^{1-\sigma} - \xi(1 - F(\varphi_t^m)).$$

Since our assumptions on foreign demand imply that $Q_t (P_t^x)^{\sigma} = (\kappa_t \tau_t)^{-\sigma} \omega_t D_t$, we can write

$$\begin{aligned} p_t^x(\tilde{\varphi}_t) &= \frac{\sigma}{\sigma-1}\frac{c_t}{\tilde{\varphi}_t} \\ q_t(\tilde{\varphi}_t) &= (p_t^x(\tilde{\varphi}_t))^{-\sigma} (\kappa_t\tau_t)^{-\sigma} \omega_t D_t \\ \pi_t &= \frac{1}{\sigma}(\kappa_t\tau_t)^{-\sigma}\omega_t D_t \left(\frac{\sigma}{\sigma-1}\frac{c_t}{\tilde{\varphi}_t}\right)^{1-\sigma} - \xi(1 - F(\varphi_t^m)) \end{aligned}$$

where $\tilde{\varphi}_t$ is given by equation (C.12) and

$$\varphi_t^m = \frac{\sigma}{\sigma-1} c_t \left(\frac{\sigma\xi}{(\kappa_t\tau_t)^{-\sigma}\omega_t D_t}\right)^{\frac{1}{\sigma-1}}.$$

It is now easy to see that the above functions take the forms assumed in equations (C.1)-(C.4).

While the exact values of elasticities (C.7) depend on the distribution F, it is always true that $\frac{\partial \ln \varphi_t^m}{\partial \ln \tau_t} = \frac{\sigma}{\sigma-1} > 0$, $\frac{\partial \ln \tilde{\varphi}_t}{\partial \ln \tau_t} < 0$, and hence $\frac{\partial \ln p_t^x}{\partial \ln \tau_t} = -\frac{\partial \ln \tilde{\varphi}_t}{\partial \ln \tau_t} > 0$. Further, the partial derivative $\frac{\partial \ln q_t}{\partial \ln \tau_t} = -\sigma$.

C.3 Model solution

Global solution Solving equation (C.5) forward gives, after imposing the transversality condition,

$$v_t = \frac{1}{1+r}\mathbb{E}_t\left[\sum_{\ell=0}^{\infty}\left(\frac{1-\delta}{1+r}\right)^{\ell}\pi_{t+\ell+1}\right].$$

Further, solving equation (C.6) backwards gives

$$n_t = \sum_{k=0}^{t-1}(1-\delta)^k G\left(v_{t-1-k}\right) + (1-\delta)^t n_0.$$

The model solution is unique: for any sequence of $\pi_{t+\ell+1}$'s, the first equation yields a unique v_t, and for any sequence of v_t's, the second equation yields a unique n_t.

Nonstochastic steady state Suppose all exogenous driving forces are constant so that $c_t = c$, $\kappa_t = \kappa$, $\tau_t = \tau$, $\omega_t = \omega$ and $D_t = D$. Then $\pi_t = \pi$, and v_t immediately collapses to

$$v = \frac{\pi}{r+\delta}.$$

Further, n_t converges to

$$n = \frac{G\left(v\right)}{\delta}.$$

These two equations characterize the non-stochastic steady state.

Long-run trade elasticity The long-run trade elasticity is

$$\begin{aligned}\frac{d\ln X}{d\ln\tau} &= \frac{d\ln q}{d\ln\tau} + \frac{d\ln p^x}{d\ln\tau} + \frac{d\ln n}{d\ln\tau}\\ &= \varepsilon^0 + \frac{d\ln n}{d\ln\tau},\end{aligned}$$

where

$$\frac{d\ln n}{d\ln\tau} = \frac{d\ln n}{d\ln v}\frac{d\ln v}{d\ln\tau} = \chi\frac{d\ln\pi}{d\ln\tau} = \chi\eta_{\pi,\tau},$$

and

$$\chi := \frac{d\ln n}{d\ln v} = \frac{d\ln G\left(v\right)}{d\ln v} = \frac{g\left(v\right)v}{G\left(v\right)}.$$

Monotone convergence If $c_t = c$, $\kappa_t = \kappa$, $\tau_t = \tau$, $\omega_t = \omega$ and $D_t = D$, then $v_t = v = \frac{\pi}{r+\delta}$. It then follows from equation (C.6) above that

$$\begin{aligned}n_t - n &= \left(1-\delta\right)\left(n_{t-1}-n\right) + G\left(v\right) - \delta n\\ &= \left(1-\delta\right)\left(n_{t-1}-n\right),\end{aligned}$$

so convergence is monotone.

Linearized economy We characterize all impulse response functions and trade elasticities up to a first order approximation. Letting tildes denote percent deviations from steady state, e.g. $\tilde{v}_t = \ln v_t - \ln v = d \ln v_t = \frac{v_t - v}{v}$, these are

$$\begin{aligned} \tilde{v}_t &= \mathbb{E}_t \left[\frac{\delta + r}{1+r} \tilde{\pi}_{t+1} + \frac{1-\delta}{1+r} \tilde{v}_{t+1} \right], & \text{(C.15)} \\ \tilde{n}_t &= \tilde{n}_{t-1} (1-\delta) + \delta \chi \tilde{v}_{t-1}, & \text{(C.16)} \end{aligned}$$

in recursive form and

$$\begin{aligned} \tilde{v}_t &= \frac{\delta + r}{1+r} \mathbb{E}_t \left[\sum_{\ell=0}^{\infty} \left(\frac{1-\delta}{1+r} \right)^{\ell} \tilde{\pi}_{t+\ell+1} \right], & \text{(C.17)} \\ \tilde{n}_t &= \delta \chi \sum_{k=0}^{t-1} (1-\delta)^{t-1-k} \tilde{v}_k + (1-\delta)^t \tilde{n}_0, & \text{(C.18)} \end{aligned}$$

when solved forwards and backwards, respectively.

Further, the static model block (C.1)-(C.4) takes the form

$$\begin{aligned} \tilde{p}_t^x &= \eta_{p,c} \tilde{c}_t + \eta_{p,\tau} (\tilde{\kappa}_t + \tilde{\tau}_t) + \eta_{p,D} \left(\tilde{\omega}_t + \tilde{D}_t \right), & \text{(C.19)} \\ \tilde{q}_t &= \eta_{q,p} \tilde{p}_t^x + \eta_{q,\tau} (\tilde{\kappa}_t + \tilde{\tau}_t) + \eta_{q,D} \left(\tilde{\omega}_t + \tilde{D}_t \right), & \text{(C.20)} \\ \tilde{\pi}_t &= \eta_{\pi,c} \tilde{c}_t + \eta_{\pi,\tau} (\tilde{\kappa}_t + \tilde{\tau}_t) + \eta_{\pi,D} \left(\tilde{\omega}_t + \tilde{D}_t \right), & \text{(C.21)} \\ \tilde{X}_t &= \tilde{q}_t + \tilde{p}_t^x + \tilde{n}_t. & \text{(C.22)} \end{aligned}$$

C.4 Proofs of propositions and examples

C.4.1 Proof of Proposition 1

Proposition 1. *Consider an arbitrary evolution of tariffs* $\left\{ \frac{d \ln \tau_{t_0+\ell}}{d \ln \tau_{t_0}} \right\}_{\ell=1}^{\infty}$ *after the shock at* t_0*. The impulse response function of* $\ln n_t$ *at horizon* $h = 0, 1, 2, \ldots$ *is*

$$\frac{d \ln n_{t_0+h}}{d \ln \tau_{t_0}} = \chi \eta_{\pi,\tau} \frac{\delta + r}{1+r} \delta \sum_{k=0}^{h-1} (1-\delta)^{h-1-k} \mathbb{E}_{t_0+k} \left[\sum_{\ell=0}^{\infty} \left(\frac{1-\delta}{1+r} \right)^{\ell} \frac{d \ln \tau_{t_0+k+\ell+1}}{d \ln \tau_{t_0}} \right].$$

Proof. Combining equation (C.17) as of time $t_0 + k$ with the fact that $\tilde{\pi}_t = \eta_{\pi,\tau} \tilde{\tau}_t$ in the version of the model with tariff shocks only (see C.21) gives

$$\tilde{v}_{t_0+k} = \eta_{\pi,\tau} \frac{\delta + r}{1+r} \mathbb{E}_{t_0+k} \left[\sum_{\ell=0}^{\infty} \left(\frac{1-\delta}{1+r} \right)^{\ell} \tilde{\tau}_{t_0+k+\ell+1} \right]. \tag{C.23}$$

Next take equation (C.16) at time $t_0 + h$ and solve it backwards until period t_0. This gives

$$\tilde{n}_{t_0+h} = \delta\chi \sum_{k=0}^{h-1} (1-\delta)^{h-1-k} \tilde{v}_{t_0+k} + (1-\delta)^h \tilde{n}_{t_0} \tag{C.24}$$

Now plugging (C.23) into (C.24) gives

$$\tilde{n}_{t_0+h} = \eta_{\pi,\tau}\chi \frac{\delta+r}{1+r}\delta \sum_{k=0}^{h-1} (1-\delta)^{h-1-k} \mathbb{E}_{t_0+k}\left[\sum_{\ell=0}^{\infty}\left(\frac{1-\delta}{1+r}\right)^{\ell} \tilde{\tau}_{t_0+k+\ell+1}\right] + (1-\delta)^h \tilde{n}_{t_0}.$$

Lastly, replace $\tilde{n}_{t_0+h}$ with $d\ln n_{t_0+h}$, etc., differentiate with respect to $d\ln\tau_{t_0}$, and note that $\frac{d\ln n_{t_0}}{d\ln\tau_{t_0}} = 0$.

□

C.4.2 Proof of Proposition 2

Proposition 2. *If* $\lim_{h\to\infty} \frac{d\ln\tau_{t_0+h}}{d\ln\tau_{t_0}} \neq 0$ *and is finite, then* $\lim_{h\to\infty} \varepsilon^h = \varepsilon$.

Proof. We first show that $\{\tilde{v}_{t_0+h}\}_{h=0}^{\infty}$ converges to $\eta_{\pi,\tau}\tilde{\tau}$. Fix an arbitrary $\psi > 0$. Since $\{\tilde{\tau}_{t_0+h}\}_{h=0}^{\infty}$ converges to $\tilde{\tau}$, there exists a h_ψ such that for $\forall h \geq h_\psi : |\tilde{\tau}_{t_0+h} - \tilde{\tau}| < \frac{\psi}{|\eta_{\pi,\tau}|}$. Next note that

$$\tilde{v}_{t+h} - \eta_{\pi,\tau}\tilde{\tau} = \frac{\delta+r}{1+r}\mathbb{E}_{t+h}\left[\sum_{\ell=0}^{\infty}\left(\frac{1-\delta}{1+r}\right)^{\ell} \eta_{\pi,\tau}\left(\tilde{\tau}_{t+h+\ell+1} - \tilde{\tau}\right)\right].$$

Then, for $h \geq h_\psi$, and using Jensen's and the triangle inequality,

$$\begin{aligned} |\tilde{v}_{t+h} - \eta_{\pi,\tau}\tilde{\tau}| &\leq \frac{\delta+r}{1+r}\mathbb{E}_{t+h}\left[\sum_{\ell=0}^{\infty}\left(\frac{1-\delta}{1+r}\right)^{\ell} |\eta_{\pi,\tau}\left(\tilde{\tau}_{t+h+\ell+1} - \tilde{\tau}\right)|\right] \\ &< \frac{\delta+r}{1+r}\mathbb{E}_{t+h}\left[\sum_{\ell=0}^{\infty}\left(\frac{1-\delta}{1+r}\right)^{\ell} \psi\right] = \psi, \end{aligned}$$

and hence $\{\tilde{v}_{t_0+h}\}_{h=0}^{\infty}$ converges to $\eta_{\pi,\tau}\tilde{\tau}$.

We next show that $\{\tilde{n}_{t_0+h}\}$ converges to $\chi\eta_{\pi,\tau}\tilde{\tau}$. Fix an arbitrary $\psi > 0$. Since $\{\tilde{v}_{t_0+h}\}_{h=0}^{\infty}$ converges to $\eta_{\pi,\tau}\tilde{\tau}$, there exists a h_ψ such that for $\forall h \geq h_\psi : |\tilde{v}_{t_0+h} - \eta_{\pi,\tau}\tilde{\tau}| < \frac{\psi}{2\chi}$. Next note that for $h > h_\psi$,

$$\begin{aligned} \tilde{n}_{t_0+h} &= \delta\chi \sum_{k=0}^{h-1} (1-\delta)^{h-1-k} \tilde{v}_{t_0+k} + (1-\delta)^h \tilde{n}_{t_0} \\ &= \delta\chi \sum_{k=h_\psi}^{h-1} (1-\delta)^{h-1-k} \tilde{v}_{t_0+k} + \delta\chi\,(1-\delta_\psi)^{h-h_\psi} \sum_{k=0}^{h_\psi-1} (1-\delta)^{h_\psi-1-k} \tilde{v}_{t_0+k} + (1-\delta)^h \tilde{n}_{t_0}. \end{aligned}$$

Then, for $h > h_\psi$,

$$\begin{aligned}
\tilde{n}_{t_0+h} - \chi\eta_{\pi,\tau}\tilde{\tau} &= \delta\chi \sum_{k=h_\psi}^{h-1} (1-\delta)^{h-1-k} \left(\tilde{v}_{t_0+k} - \eta_{\pi,\tau}\tilde{\tau} + \eta_{\pi,\tau}\tilde{\tau}\right) - \chi\eta_{\pi,\tau}\tilde{\tau} \\
&\quad + \delta\chi \left(1-\delta_\psi\right)^{h-h_\psi} \sum_{k=0}^{h_\psi-1} (1-\delta)^{h_\psi-1-k} \tilde{v}_{t_0+k} + (1-\delta)^h \tilde{n}_{t_0} \\
&= \delta\chi \sum_{k=h_\psi}^{h-1} (1-\delta)^{h-1-k} \left(\tilde{v}_{t_0+k} - \eta_{\pi,\tau}\tilde{\tau}\right) + \delta\chi\eta_{\pi,\tau}\tilde{\tau} \sum_{k=h_\psi}^{h-1} (1-\delta)^{h-1-k} - \chi\eta_{\pi,\tau}\tilde{\tau} \\
&\quad + \delta\chi \left(1-\delta_\psi\right)^{h-h_\psi} \sum_{k=0}^{h_\psi-1} (1-\delta)^{h_\psi-1-k} \tilde{v}_{t_0+k} + (1-\delta)^h \tilde{n}_{t_0} \\
&= \delta\chi \sum_{k=h_\psi}^{h-1} (1-\delta)^{h-1-k} \left(\tilde{v}_{t_0+k} - \eta_{\pi,\tau}\tilde{\tau}\right) \\
&\quad - \chi\eta_{\pi,\tau}\tilde{\tau} (1-\delta)^{h-h_\psi} + \delta\chi \left(1-\delta_\psi\right)^{h-h_\psi} \sum_{k=0}^{h_\psi-1} (1-\delta)^{h_\psi-1-k} \tilde{v}_{t_0+k} + (1-\delta)^h \tilde{n}_{t_0},
\end{aligned}$$

where we used that $\sum_{k=h_\psi}^{h-1} (1-\delta)^{h-1-k} = \frac{1-(1-\delta)^{h-h_\psi}}{\delta}$. Next note that

$$\begin{aligned}
\left| \delta\chi \sum_{k=h_\psi}^{h-1} (1-\delta)^{h-1-k} \left(\tilde{v}_{t_0+k} - \eta_{\pi,\tau}\tilde{\tau}\right) \right| &\leq \delta\chi \sum_{k=h_\psi}^{h-1} (1-\delta)^{h-1-k} \left|\tilde{v}_{t_0+k} - \eta_{\pi,\tau}\tilde{\tau}\right| \\
&< \delta\chi \sum_{k=h_\psi}^{h-1} (1-\delta)^{h-1-k} \frac{\psi}{2\chi} = \frac{\psi}{2}\left[1-(1-\delta)^{h-h_\psi}\right].
\end{aligned}$$

Hence,

$$\begin{aligned}
\left|\tilde{n}_{t_0+h} - \chi\eta_{\pi,\tau}\tilde{\tau}\right| &< \frac{\psi}{2}\left[1-(1-\delta)^{h-h_\psi}\right] + (1-\delta)^{h-h_\psi} \left|\chi\eta_{\pi,\tau}\tilde{\tau}\right| \\
&\quad + \left(1-\delta_\psi\right)^{h-h_\psi} \left| \delta\chi \sum_{k=0}^{h_\psi-1} (1-\delta)^{h_\psi-1-k} \tilde{v}_{t_0+k} \right| + (1-\delta)^h \left|\tilde{n}_{t_0}\right|.
\end{aligned}$$

Now choosing $h_\psi^* > h_\psi$ such that for all $h > h_\psi^*$ the last three terms are smaller than $\frac{\psi}{2}$, implies that $\tilde{n}_{t_0+h}$ converges to $\chi\eta_{\pi,\tau}\tilde{\tau}$.

Lastly note that $\tilde{X}_{t_0+h} = \varepsilon^0\tilde{\tau}_{t_0+h} + \tilde{n}_{t_0+h}$, and hence $\lim_{h\to\infty}\tilde{X}_{t_0+h} = \varepsilon^0\tilde{\tau} + \chi\eta_{\pi,\tau}\tilde{\tau} = \varepsilon\tilde{\tau}$. Since $\tilde{\tau} \neq 0$, $\lim_{h\to\infty}\varepsilon^h = \lim_{h\to\infty}\frac{\tilde{X}_{t_0+h}}{\tilde{\tau}_{t_0+h}} = \varepsilon$.

□

C.4.3 Details on Example 1

Plug $\Delta \ln \tau_{>t_0}$ into equation (6.10). This gives

$$\begin{aligned} \frac{d \ln n_{t_0+h}}{d \ln \tau_{t_0}} &= \chi \eta_{\pi,\tau} \Delta \ln \tau_{>t_0} \delta \sum_{k=0}^{h-1} (1-\delta)^{h-1-k} \\ &= \chi \eta_{\pi,\tau} \left[1-(1-\delta)^h\right] \Delta \ln \tau_{>t_0}. \end{aligned}$$

The claim now follows immediately.

C.4.4 Details on Example 2

Tariffs follow a first or autoregressive process with autoregressive root ρ. Then

$$\mathbb{E}_{t_0+k}\left[\frac{d \ln \tau_{t_0+k+\ell+1}}{d \ln \tau_{t_0}}\right] = \rho^{\ell+k+1}.$$

Plugging this expression into (6.10) gives

$$\begin{aligned} \frac{d \ln n_{t_0+h}}{d \ln \tau_{t_0}} &= \chi \eta_{\pi,\tau} \frac{\delta+r}{1+r} \delta \sum_{k=0}^{h-1} (1-\delta)^{h-1-k} \sum_{\ell=0}^{\infty} \left(\frac{1-\delta}{1+r}\right)^{\ell} \rho^{\ell+k+1} \\ &= \chi \eta_{\pi,\tau} \frac{\delta+r}{1+r} \delta \sum_{k=0}^{h-1} (1-\delta)^{h-1-k} (\rho)^{k+1} \left[\sum_{\ell=0}^{\infty} \left(\frac{1-\delta}{1+r}\rho\right)^{\ell}\right] \\ &= \chi \eta_{\pi,\tau} \frac{\delta+r}{1+r-(1-\delta)\rho} \delta \rho^h \sum_{k=0}^{h-1} \left(\frac{1-\delta}{\rho}\right)^{h-1-k} \\ &= \chi \eta_{\pi,\tau} \frac{\delta+r}{1+r-(1-\delta)\rho} \delta \rho^h \frac{1-\left(\frac{1-\delta}{\rho}\right)^h}{1-\frac{1-\delta}{\rho}}. \end{aligned}$$

Since

$$\frac{\frac{d \ln n_{t_0+h}}{d \ln \tau_{t_0}}}{\frac{d \ln \tau_{t_0+h}}{d \ln \tau_{t_0}}} = \chi \eta_{\pi,\tau} \frac{(\delta+r)\delta}{[1+r-(1-\delta)\rho]\left(1-\frac{1-\delta}{\rho}\right)} \left(1-\left(\frac{1-\delta}{\rho}\right)^h\right),$$

the claim follows immediately.

C.4.5 Proof of Proposition 3

Proposition 3. *The model delivers estimating equation* (2.2), *where*

$$\beta_X^h = \chi \eta_{\pi,\tau} \frac{\delta+r}{1+r} \delta \sum_{k=0}^{h-1} (1-\delta)^{h-1-k} \sum_{\ell=0}^{\infty} \left(\frac{1-\delta}{1+r}\right)^{\ell} \beta_\tau^{k+\ell+1} + \varepsilon^0 \beta_\tau^h.$$

β_τ^h is defined as the regression coefficient of $\Delta_h \ln \tau_{i,j,p,t}$ on $\Delta_0 \ln \tau_{i,j,p,t}$ in the population, and can be estimated from equation (2.3).

After augmenting the model with additional shocks, the fixed effects $\delta_{j,p,t}^{X,h}$ and $\delta_{i,p,t}^{X,h}$ capture a weighted sum of past, present, and expected future supply and demand shocks, respectively. The error term includes past, present, and expected future time-varying bilateral and product-specific demand shocks and non-tariff trade barriers, as well as the initial state.

Proof. Using equations (C.20) and (C.19), and the definition of $\varepsilon^0 = (1+\eta_{q,p})\,\eta_{p,\tau} + \eta_{q,\tau}$ (equation (6.6)), we have

$$\begin{aligned}
\tilde{q}_{t+h} - \tilde{q}_{t-1} + \tilde{p}_{t+h}^x - \tilde{p}_{t-1}^x &= \varepsilon^0 \left(\tilde{\tau}_{t+h} - \tilde{\tau}_{t-1}\right) \\
&\quad + \varepsilon^0 \left(\tilde{\kappa}_{t+h} - \tilde{\kappa}_{t-1}\right) + \left[(1+\eta_{q,p})\,\eta_{p,D} + \eta_{q,D}\right]\left(\tilde{\omega}_{t+h} - \tilde{\omega}_{t-1}\right) \\
&\quad + (1+\eta_{q,p})\,\eta_{p,c}\left(\tilde{c}_{t+h} - \tilde{c}_{t-1}\right) \\
&\quad + \left[(1+\eta_{q,p})\,\eta_{p,D} + \eta_{q,D}\right]\left(\tilde{D}_{t+h} - \tilde{D}_{t-1}\right).
\end{aligned}$$

Next, note that

$$\tilde{v}_{t+k} = \frac{\delta + r}{1+r}\mathbb{E}_{t+k}\left[\sum_{\ell=0}^{\infty}\left(\frac{1-\delta}{1+r}\right)^{\ell}\tilde{\pi}_{t+k+\ell+1}\right]$$

and

$$\tilde{n}_{t+h} = \delta\chi\sum_{k=0}^{h-1}(1-\delta)^{h-1-k}\,\tilde{v}_{t+k} + \tilde{n}_t\,(1-\delta)^h,$$

so that

$$\begin{aligned}
\tilde{n}_{t+h} - \tilde{n}_{t-1} &= \chi\frac{\delta+r}{1+r}\delta\sum_{k=0}^{h-1}(1-\delta)^{h-1-k}\,\mathbb{E}_{t+k}\left[\sum_{\ell=0}^{\infty}\left(\frac{1-\delta}{1+r}\right)^{\ell}\tilde{\pi}_{t+k+\ell+1}\right] + \tilde{n}_t\,(1-\delta)^h - \tilde{n}_{t-1} \\
&= \chi\frac{\delta+r}{1+r}\delta\sum_{k=0}^{h-1}(1-\delta)^{h-1-k}\,\mathbb{E}_{t+k}\left[\sum_{\ell=0}^{\infty}\left(\frac{1-\delta}{1+r}\right)^{\ell}\left(\tilde{\pi}_{t+k+\ell+1} - \tilde{\pi}_{t-1}\right)\right] \\
&\quad + \chi\tilde{\pi}_{t-1}\left[1 - (1-\delta)^h\right] + \tilde{n}_t\,(1-\delta)^h - \tilde{n}_{t-1}.
\end{aligned}$$

From (C.21) we obtain

$$\begin{aligned}
\tilde{\pi}_{t+k+\ell+1} - \tilde{\pi}_{t-1} &= \eta_{\pi,c}\left(\tilde{c}_{t+k+\ell+1} - \tilde{c}_{t-1}\right) + \eta_{\pi,\tau}\left(\tilde{\kappa}_{t+k+\ell+1} - \tilde{\kappa}_{t-1}\right) + \eta_{\pi,\tau}\left(\tilde{\tau}_{t+k+\ell+1} - \tilde{\tau}_{t-1}\right) \\
&\quad + \eta_{\pi,D}\left(\tilde{\omega}_{t+k+\ell+1} - \tilde{\omega}_{t-1}\right) + \eta_{\pi,D}\left(\tilde{D}_{t+k+\ell+1} - \tilde{D}_{t-1}\right).
\end{aligned}$$

Now putting the pieces together, and adding the subscripts back in, we have that

$$\begin{aligned}
\Delta_h \ln X_{i,j,p,t} &= \varepsilon^0\Delta_h \ln \tau_{i,j,p,t} + \eta_{\pi,\tau}\chi\frac{\delta+r}{1+r}\delta\sum_{k=0}^{h-1}(1-\delta)^{h-1-k}\,\mathbb{E}_{t+k}\left[\sum_{\ell=0}^{\infty}\left(\frac{1-\delta}{1+r}\right)^{\ell}\Delta_{k+\ell+1}\ln\tau_{i,j,p,t}\right] \\
&\quad + \delta_{j,p,t}^h + \delta_{i,p,t}^h + u_{i,j,p,t}^{X,h}
\end{aligned}$$

where we used that, for a generic variable x_t, $\Delta_h x_t = x_{t+h} - x_{t-1}$, and

$$\begin{aligned}
\delta_{i,p,t}^{X,h} &= [(1+\eta_{q,p})\,\eta_{p,D} + \eta_{q,D}]\,\Delta_h \ln D_{i,p,t} \\
&+\eta_{\pi,D}\chi\frac{\delta+r}{1+r}\delta\sum_{k=0}^{h-1}(1-\delta)^{h-1-k}\,\mathbb{E}_{t+k}\left[\sum_{\ell=0}^{\infty}\left(\frac{1-\delta}{1+r}\right)^{\ell}\Delta_{k+\ell+1}\ln D_{i,p,t}\right], \\
\delta_{j,p,t}^{X,h} &= (1+\eta_{q,p})\,\eta_{p,c}\Delta_h \ln c_{j,p,t} \\
&+\eta_{\pi,c}\chi\frac{\delta+r}{1+r}\delta\sum_{k=0}^{h-1}(1-\delta)^{h-1-k}\,\mathbb{E}_{t+k}\left[\sum_{\ell=0}^{\infty}\left(\frac{1-\delta}{1+r}\right)^{\ell}\Delta_{k+\ell+1}\ln c_{j,p,t}\right], \\
u_{i,j,p,t}^{X,h} &= \varepsilon^0\Delta_h \ln \kappa_{i,j,p,t} + \eta_{\pi,\tau}\chi\frac{\delta+r}{1+r}\delta\sum_{k=0}^{h-1}(1-\delta)^{h-1-k}\,\mathbb{E}_{t+k}\left[\sum_{\ell=0}^{\infty}\left(\frac{1-\delta}{1+r}\right)^{\ell}\Delta_{k+\ell+1}\ln \kappa_{i,j,p,t}\right] \\
&+ [(1+\eta_{q,p})\,\eta_{p,D} + \eta_{q,D}]\,\Delta_h \ln \omega_{i,j,p,t} \\
&+\eta_{\pi,D}\chi\frac{\delta+r}{1+r}\delta\sum_{k=0}^{h-1}(1-\delta)^{h-1-k}\,\mathbb{E}_{t+k}\left[\sum_{\ell=0}^{\infty}\left(\frac{1-\delta}{1+r}\right)^{\ell}\Delta_{k+\ell+1}\ln \omega_{i,j,p,t}\right] \\
&+\chi\tilde{\pi}_{i,j,p,t-1}\left[1-(1-\delta)^h\right] + \tilde{n}_{i,j,p,t}\,(1-\delta)^h - \tilde{n}_{i,j,p,t-1}.
\end{aligned}$$

Next define the regression coefficient of $\Delta_h \ln \tau_{i,j,p,t}$ on $\Delta_0 \ln \tau_{i,j,p,t}$ as β_τ^h in the population, where we assume that $\Delta_0 \ln \tau_{i,j,p,t}$ is an exogenous tariff shock. Clearly, β_τ^h can be estimated from equation (2.3). Then the estimating equation becomes

$$\Delta_h \ln X_{i,j,p,t} = \beta_X^h \Delta_0 \ln \tau_{i,j,p,t} + \delta_{j,p,t}^{X,h} + \delta_{i,p,t}^{X,h} + u_{i,j,p,t}^{X,h}.$$

where

$$\beta_X^h = \chi\eta_{\pi,\tau}\frac{\delta+r}{1+r}\delta\sum_{k=0}^{h-1}(1-\delta)^{h-1-k}\sum_{\ell=0}^{\infty}\left(\frac{1-\delta}{1+r}\right)^{\ell}\beta_\tau^{k+\ell+1} + \varepsilon^0\beta_\tau^h.$$

Note that β_τ^h, $h = 0, 1, ...$ are constants, so the expectation drops out.

□

C.5 Estimation in long differences

Proposition C.1. *(Part 1) Estimation as a horizon-h difference does generally not identify the horizon-h trade elasticity.*

(Part 2) If tariffs follow a random walk, a regression of $\Delta_h \ln X_t$ on $\Delta_h \ln \tau_t$ identifies the simple average of horizon-0 to horizon-h trade elasticities.

Proof. Since the first part of the proposition follows from the second part, we prove the second part.

Tariffs follow a random walk,

$$\tilde{\tau}_t = \tilde{\tau}_{t-1} + \sigma_u u_t^\tau,$$

where u_t^τ is white noise with unit variance, and σ_u denotes the standard deviation of the innovation to tariffs. Then

$$\tilde{\tau}_{t+k} - \tilde{\tau}_{t-1} = \sigma_u \sum_{j=0}^{k} u_{t+j}^\tau.$$

Consider the projection of $\tilde{\tau}_{t+k} - \tilde{\tau}_{t-1}$ on $\tilde{\tau}_{t+h} - \tilde{\tau}_{t-1}$ (i.e. the OLS estimator),

$$\frac{\mathbb{C}ov\left[\tilde{\tau}_{t+k} - \tilde{\tau}_{t-1}, \tilde{\tau}_{t+h} - \tilde{\tau}_{t-1}\right]}{\mathbb{V}\left[\tilde{\tau}_{t+h} - \tilde{\tau}_{t-1}\right]} = \frac{\mathbb{C}ov\left[\sum_{j=0}^{k} u_{t+k}^\tau, \sum_{j=0}^{h} u_{t+k}^\tau\right]}{\mathbb{V}\left[\sum_{j=0}^{h} u_{t+k}^\tau\right]} = \frac{k+1}{h+1}. \tag{C.25}$$

Next note that

$$\begin{aligned}\tilde{n}_{t+h} - \tilde{n}_{t-1} &= \chi \frac{\delta + r}{1+r} \delta \sum_{k=0}^{h-1} (1-\delta)^{h-1-k} \, \mathbb{E}_{t+k} \left[\sum_{\ell=0}^{\infty} \left(\frac{1-\delta}{1+r} \right)^{\ell} (\tilde{\pi}_{t+k+\ell+1} - \tilde{\pi}_{t-1}) \right] \\ &\quad + \chi \tilde{\pi}_{t-1} \left[1 - (1-\delta)^h \right] + \tilde{n}_t (1-\delta)^h - \tilde{n}_{t-1},\end{aligned}$$

which implies, together with

$$\tilde{\pi}_{t+k+\ell+1} - \tilde{\pi}_{t-1} = \eta_{q,\tau} \left(\tilde{\tau}_{t+k+\ell+1} - \tilde{\tau}_{t-1} \right)$$

that

$$\begin{aligned}\tilde{n}_{t+h} - \tilde{n}_{t-1} &= \chi \eta_{q,\tau} \frac{\delta + r}{1+r} \delta \sum_{k=0}^{h-1} (1-\delta)^{h-1-k} \, \mathbb{E}_{t+k} \left[\sum_{\ell=0}^{\infty} \left(\frac{1-\delta}{1+r} \right)^{\ell} (\tilde{\tau}_{t+k+\ell+1} - \tilde{\tau}_{t-1}) \right] \\ &\quad + \chi \tilde{\pi}_{t-1} \left[1 - (1-\delta)^h \right] + \tilde{n}_t (1-\delta)^h - \tilde{n}_{t-1}.\end{aligned}$$

Since $\mathbb{E}_{t+k}\left[\tilde{\tau}_{t+k+\ell+1}\right] = \tilde{\tau}_{t+k}$, this expression becomes

$$\begin{aligned}\tilde{n}_{t+h} - \tilde{n}_{t-1} &= \chi \eta_{q,\tau} \delta \sum_{k=0}^{h-1} (1-\delta)^{h-1-k} (\tilde{\tau}_{t+k} - \tilde{\tau}_{t-1}) \\ &\quad + \chi \tilde{\pi}_{t-1} \left[1 - (1-\delta)^h \right] + \tilde{n}_t (1-\delta)^h - \tilde{n}_{t-1}.\end{aligned}$$

Now

$$\begin{aligned}\tilde{X}_{t+h} - \tilde{X}_{t-1} &= \tilde{q}_{t+h} - \tilde{q}_{t-1} + \tilde{p}_{t+h}^x - \tilde{p}_{t-1}^x + \tilde{n}_{t+h} - \tilde{n}_{t-1} \\ &= \tilde{q}_{t+h} - \tilde{q}_{t-1} + \tilde{p}_{t+h}^x - \tilde{p}_{t-1}^x + \chi \eta_{q,\tau} \delta \sum_{k=0}^{h-1} (1-\delta)^{h-1-k} (\tilde{\tau}_{t+k} - \tilde{\tau}_{t-1}) \\ &\quad + \chi \tilde{\pi}_{t-1} \left[1 - (1-\delta)^h \right] + \tilde{n}_t (1-\delta)^h - \tilde{n}_{t-1}\end{aligned}$$

and regressing this on $(\tilde{\tau}_{t+h} - \tilde{\tau}_{t-1})$ gives

$$\frac{\mathbb{C}ov\left(\tilde{X}_{t+h} - \tilde{X}_{t-1}, \tilde{\tau}_{t+h} - \tilde{\tau}_{t-1}\right)}{\mathbb{V}\left(\tilde{\tau}_{t+h} - \tilde{\tau}_{t-1}, \tilde{\tau}_{t+h} - \tilde{\tau}_{t-1}\right)} \tag{C.26}$$

$$= \frac{\mathbb{C}ov\left(\tilde{q}_{t+h} - \tilde{q}_{t-1} + \tilde{p}^x_{t+h} - \tilde{p}^x_{t-1} + \chi\eta_{q,\tau}\delta\sum_{k=0}^{h-1}(1-\delta)^{h-1-k}\left(\tilde{\tau}_{t+k} - \tilde{\tau}_{t-1}\right), \tilde{\tau}_{t+h} - \tilde{\tau}_{t-1}\right)}{\mathbb{V}\left(\tilde{\tau}_{t+h} - \tilde{\tau}_{t-1}, \tilde{\tau}_{t+h} - \tilde{\tau}_{t-1}\right)}$$

$$= \frac{\mathbb{C}ov\left(\tilde{q}_{t+h} - \tilde{q}_{t-1} + \tilde{p}^x_{t+h} - \tilde{p}^x_{t-1}, \tilde{\tau}_{t+h} - \tilde{\tau}_{t-1}\right)}{\mathbb{V}\left(\tilde{\tau}_{t+h} - \tilde{\tau}_{t-1}, \tilde{\tau}_{t+h} - \tilde{\tau}_{t-1}\right)} + \frac{\mathbb{C}ov\left(\chi\eta_{q,\tau}\delta\sum_{k=0}^{h-1}(1-\delta)^{h-1-k}\left(\tilde{\tau}_{t+k} - \tilde{\tau}_{t-1}\right), \tilde{\tau}_{t+h} - \tilde{\tau}_{t-1}\right)}{\mathbb{V}\left(\tilde{\tau}_{t+h} - \tilde{\tau}_{t-1}, \tilde{\tau}_{t+h} - \tilde{\tau}_{t-1}\right)}$$

$$= \varepsilon^0 + \chi\eta_{q,\tau}\delta\sum_{k=0}^{h-1}(1-\delta)^{h-1-k}\frac{\mathbb{C}ov\left(\tilde{\tau}_{t+k} - \tilde{\tau}_{t-1}, \tilde{\tau}_{t+h} - \tilde{\tau}_{t-1}\right)}{\mathbb{V}\left(\tilde{\tau}_{t+h} - \tilde{\tau}_{t-1}, \tilde{\tau}_{t+h} - \tilde{\tau}_{t-1}\right)}$$

$$= \varepsilon^0 + \chi\eta_{q,\tau}\delta\sum_{k=0}^{h-1}(1-\delta)^{h-1-k}\frac{k+1}{h+1} \tag{C.27}$$

where the last equality uses equation (C.25) above.

Next note that

$$\begin{aligned}
\sum_{k=0}^{h-1}(1-\delta)^{h-1-k}\frac{k+1}{h+1} &= (1-\delta)^{h-1}\frac{1}{h+1} + (1-\delta)^{h-2}\frac{2}{h+1} + ... + (1-\delta)\frac{h-1}{h+1} + \frac{h}{h+1} \\
&= \frac{1}{h+1}[1] \\
&\quad + \frac{1}{h+1}[1 + (1-\delta)] \\
&\quad + ... \\
&\quad + \frac{1}{h+1}\left[1 + (1-\delta) + ... + (1-\delta)^{h-2}\right] \\
&\quad + \frac{1}{h+1}\left[1 + (1-\delta) + ... + (1-\delta)^{h-2} + (1-\delta)^{h-1}\right] \\
&= \frac{1}{h+1}\sum_{k=0}^{h-1}\sum_{j=0}^{k}(1-\delta)^j \\
&= \frac{1}{h+1}\sum_{k=0}^{h-1}\frac{1-(1-\delta)^{k+1}}{\delta}.
\end{aligned}$$

Plugging this expression into equation (C.27) gives

$$\begin{aligned}
\frac{\mathbb{C}ov\left(\tilde{X}_{t+h}-\tilde{X}_{t-1},\tilde{\tau}_{t+h}-\tilde{\tau}_{t-1}\right)}{\mathbb{V}\left(\tilde{\tau}_{t+h}-\tilde{\tau}_{t-1},\tilde{\tau}_{t+h}-\tilde{\tau}_{t-1}\right)} &= \varepsilon^0+\chi\eta_{q,\tau}\delta\frac{1}{h+1}\sum_{k=0}^{h-1}\frac{1-(1-\delta)^{k+1}}{\delta} \\
&= \varepsilon^0+\chi\eta_{q,\tau}\frac{1}{h+1}\sum_{k=0}^{h-1}\left[1-(1-\delta)^{k+1}\right] \\
&= \varepsilon^0+\chi\eta_{q,\tau}\frac{1}{h+1}\sum_{k=0}^{h}\left[1-(1-\delta)^{k}\right] \\
&= \frac{1}{h+1}\sum_{k=0}^{h}\varepsilon^k
\end{aligned}$$

where we used that $\varepsilon^h=\varepsilon^0+\chi\eta_{q,\tau}\left(1-(1-\delta)^h\right)$, see equation (6.11) of Example 1.

□