

The Long and Short (Run) of Trade Elasticities[†]

By CHRISTOPH E. BOEHM, ANDREI A. LEVCHENKO,
AND NITYA PANDALAI-NAYAR*

When countries change most favored nation (MFN) tariffs, partners that trade on MFN terms experience plausibly exogenous tariff changes. Using this variation, we estimate the trade elasticity at short and long horizons with local projections. We find that the elasticity of tariff-exclusive trade flows is -0.76 in the short run, and approximately -2 in the long run. Our long-run estimates are smaller than typical in the literature, and it takes 7 to 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 exports. (JEL C51, F13, F14)

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, trade balances, 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 between -1.75 and -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 in most static trade

*Boehm: University of Texas at Austin, and NBER (email: cboehm@utexas.edu); Levchenko: University of Michigan, and NBER (email: alev@umich.edu); Pandalai-Nayar: University of Texas at Austin, and NBER (email: npnayar@utexas.edu). Emi Nakamura was the coeditor for this article. We are grateful to the paper’s first-round coeditor, Pierre-Olivier Gourinchas, as well as four anonymous referees for extensive input that improved the paper, and to George Alessandria, Costas Arkolakis, Dominick Bartelme, Chad Bown, Emily Blanchard, Meredith Crowley, Alejandro Cuñat, Doireann Fitzgerald, 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 UC Berkeley, 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, SED-AM and the Minneapolis Fed for helpful comments, and to Barthélémy Bonadio, Jaedo Choi, George Cui, Jonathan Garita, Elliot Kang, Hiroshi Toma, and Florian Trouvain for excellent research assistance.

[†]Go to <https://doi.org/10.1257/aer.20210225> to visit the article page for additional materials and author disclosure statements.

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

models is around 1 in absolute value, and thus the gains from trade implied by these 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.

Our first contribution is to address the endogeneity of tariffs. Our estimating equations time-difference the data in order to account for omitted variables that vary by country-pair-product. However, 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 World Trade Organization (WTO) system: the most favored nation (MFN) principle. Under this principle, a country must apply the same tariffs to all its WTO 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 an importer's decision to change its MFN tariffs. Our estimation procedure then compares the changes in 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 ordinary least squares (OLS).

Our second contribution is to provide estimates over different time horizons, ranging from impact to ten 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.

In the process, we 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 shifters or other unobserved bilateral gravity variables. For instance, if buyers in a particular importing country have time-invariant 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.²

Our analysis uses data on global international trade flows from Base pour le'Analyse du Commerce International (BACI) (Gaulier and Zignago 2010; Centre d'Études Prospectives et d'Informations Internationales 2022a), and tariffs from the United Nations Trade Analysis Information System (UN TRAINS). The sample covers 183 economies, over 5,000 harmonized system (HS) six-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. It 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). We also present a multicountry multisector general equilibrium (GE) extension to show that the time path of our estimated elasticities is a key input for quantifying the GE responses of trade flows to tariff shocks.

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 five to six times larger than under the commonly used elasticity of -5 .

More broadly, our elasticity estimates can inform a range of long-standing questions in international macroeconomics. The parameters governing the response of

²While we are not the first to control for bilateral unobservables via time-differencing or fixed effects (see, e.g., Feenstra 1994; Head and Ries 2001; Lai and Trefler 2002, and the literature that followed), this approach is not common practice in trade elasticity estimation, and the systematic implications for the value of the estimated elasticities have not been emphasized by this body of work.

trade flows to relative price changes are crucial for understanding the impact of the exchange rate on the trade balance (e.g., Marshall 1923; Lerner 1944; Backus, Kehoe, and Kydland 1994; Obstfeld and Rogoff 2005; Imbs and Mejean 2015). For instance, Backus, Kehoe, and Kydland (1994) illustrate the role of a low Armington elasticity in matching the response of the trade balance to terms-of-trade changes. Our short-horizon estimates can be used to discipline business-cycle substitution elasticities between home and foreign goods, and so can also be used in quantifications of business cycle shock transmission across countries (e.g., Backus, Kehoe, and Kydland 1992; Heathcote and Perri 2002; Kose and Yi 2006; Johnson 2014; Huo, Levchenko, and Pandalai-Nayar 2020; Drozd, Kolbin, and Nosal 2021). Finally, the trade elasticity is also important for international risk sharing (Cole and Obstfeld 1991; Coeurdacier 2009; Heathcote and Perri 2013), among others. Indeed, the full time path of our elasticity estimates contains information useful for disciplining international macro models, as demonstrated recently by Auclert et al. (2021).

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; Hertel et al. 2007; Romalis 2007; Caliendo and Parro 2015; Imbs and Mejean 2015, 2017; Fontagné, Guimbard, and Orefice 2022). Other methods exploit differences in prices across locations (Eaton and Kortum 2002; Simonovska and Waugh 2014; Giri, Yi, and Yilmazkuday 2021) or rely on variation in transport costs (Hummels 2001; Shapiro 2016; Adão, Costinot, and Donaldson 2017). Existing estimates typically do not 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 elasticity of substitution (or demand elasticity) governs the trade elasticity, but in many others, such as the Melitz or Eaton-Kortum models, 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 exogenous 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

³A 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 2021). 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.

one) and years, making it challenging to control for multilateral resistance terms and/or exploit time-series variation in tariffs for identification.

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 are 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 the 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 I lays out the econometric framework and the identification strategy. Section II describes the data, and Section III the main results. Section IV explores the empirical estimates in a number of dimensions, while Section V connects the estimates to theory. Section VI concludes.

I. Estimation Framework

A. The Horizon- h Trade Elasticity

As the objective of this paper is to estimate elasticities of trade flows to tariffs 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 $\tau_{i,j,p,t}$ be the gross ad valorem tariff. 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 $\Delta_h \ln \tau_{i,j,p,t} \neq 0$ the horizon- h trade elasticity ε^h is defined as

$$(1) \quad \varepsilon^h = \frac{\Delta_h \ln X_{i,j,p,t}}{\Delta_h \ln \tau_{i,j,p,t}}.$$

Both conceptually and for the purposes of estimation, it is important to take into account the fact that tariffs follow a stochastic process, and the h -horizon change $\Delta_h \ln \tau_{i,j,p,t}$ is a cumulation 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 tariffs occurs at time t . The denominator $\Delta_h \ln \tau_{i,j,p,t}$ captures the effect of this shock on tariffs 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 $\ln \tau_{i,j,p,t}$ and of the subsequent changes in $\ln \tau_{i,j,p,t}$ 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 $\Delta_h \ln \tau_{i,j,p,t}$ can be thought of as sequences following the initial shock. They jointly inform the time path of the response of trade flows to tariff changes in dynamic models.

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

$$\varepsilon = \lim_{h \rightarrow \infty} \varepsilon^h$$

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

B. Estimating Equations

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):

$$(2) \quad \Delta_h \ln X_{i,j,p,t} = \beta_X^h \Delta_0 \ln \tau_{i,j,p,t} + \delta_{i,p,t}^{d,X,h} + \delta_{j,p,t}^{s,X,h} + \delta_{i,j,p}^{b,X,h} + u_{i,j,p,t}^{X,h}$$

where the δ s are fixed effects. As we will discuss throughout the paper, these fixed effects will help control for the effects of other determinants of trade flows that may be correlated with tariff changes. The coefficient β_X^h in equation (2) captures the change in trade flows h periods ahead that follows an initial one-period change in tariffs: $\Delta_h \ln X_{i,j,p,t} / \Delta_0 \ln \tau_{i,j,p,t}$. If tariff changes were always one-time permanent shocks, β_X^h would be an estimate of the horizon- h trade elasticity, a point we return to below.

This estimation approach affects the interpretation of the coefficients. Because the estimating equation (2) includes importer-product-time and exporter-product-time fixed effects $\delta_{i,p,t}^{d,X,h}$ and $\delta_{j,p,t}^{s,X,h}$, the version of (1) that we estimate in the data is a partial elasticity. The fixed effects control for determinants of trade flows that vary at the importer-product-time level (e.g., demand shifters such as the importer price index and aggregate consumption), and the exporter-product-time level (such as the marginal costs of production). In GE models a tariff change generally affects these demand and supply shifters, and hence the total response of trade flows to tariffs generally differs from the partial effect on trade flows estimated here. Section VB illustrates this point in a dynamic GE model. It also shows that the total response of trade flows in GE depends strongly on a set of structural parameters that can be disciplined with our estimates.

Estimation of the partial elasticity is preferable for several reasons. In part, it is driven by necessity. Isolating exogenous variation in tariff changes without these fixed effects is substantially more challenging. We claim instrument validity only conditional on including importer-product-time and exporter-product-time fixed effects. An additional advantage of the partial elasticity is that its mapping to model parameters is substantially cleaner, as we demonstrate in Section VA. Lastly, estimation of a partial elasticity is in line with virtually all of the modern literature on this topic. For instance, in a large class of static trade models, the long-run partial elasticity is a critical determinant of the welfare gains from trade (Arkolakis, Costinot, and Rodríguez-Clare 2012).

Autocorrelation in Tariffs.—In equation (2), if $\Delta_0 \ln \tau_{i,j,p,t}$ were 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 estimator of ε^h for all

h. We can, and do, estimate β_X^h , but it is often 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:

$$(3) \quad \Delta_h \ln \tau_{i,j,p,t} = \beta_\tau^h \Delta_0 \ln \tau_{i,j,p,t} + \delta_{i,p,t}^{d,\tau,h} + \delta_{j,p,t}^{s,\tau,h} + \delta_{i,j,p}^{b,\tau,h} + u_{i,j,p,t}^{\tau,h}$$

where the impact effect is $\beta_\tau^0 = 1$ by definition. The horizon-*h* trade elasticity can then be recovered as $\varepsilon^h = \beta_X^h / \beta_\tau^h$.

Further, we estimate the combined specification:

$$(4) \quad \Delta_h \ln X_{i,j,p,t} = \beta^h \Delta_h \ln \tau_{i,j,p,t} + \delta_{i,p,t}^{d,h} + \delta_{j,p,t}^{s,h} + \delta_{i,j,p}^{b,h} + u_{i,j,p,t}^h$$

When $\Delta_0 \ln \tau_{i,j,p,t}$ is used as an instrument for $\Delta_h \ln \tau_{i,j,p,t}$, (4) is equivalent to (2) and (3), and directly identifies the trade elasticity at horizon *h* using the impulse at time *t*: $\hat{\beta}^h$ is an estimator of ε^h . Estimating (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.⁴

In practice, the period length is a year and this estimation is carried out at different horizons $h = 0, \dots, 10$, to trace the full profile of ε^h over $h \leq 10$. If the estimates of β_X^h and β_τ^h stabilize within ten years of the shock, we interpret it as convergence of both the numerator and the denominator in (1), rendering our estimates informative about the long-run trade elasticity. Section VA provides a detailed discussion of the convergence to the long-run elasticity in the context of a conventional class of models. 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^h s.

The estimating equations (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 obtain 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

⁴Equation (3) includes the same fixed effects as equations (2) and (4). This is because when estimating the trade elasticity, these same fixed effects will also be partialled out from the tariff change on the right-hand side of equation (4). The estimates from (3) therefore illustrate the tariff dynamics after a tariff shock using the same variation that is used to estimate the trade elasticity.

(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 V. This model delivers estimating equations (2) and (3), and illustrates that the fixed effects capture dynamic analogues of multilateral resistance terms. In particular, the importer-product-time and exporter-product time effects absorb weighted averages of past, present and expected future demand and supply shifters.

Conventional Estimation.—A common approach to estimating the long-run elasticity ε starts from a static gravity equation: $\ln X_{i,j,p,t} = \beta \ln \tau_{i,j,p,t} + \delta_{i,p,t}^d + \delta_{j,p,t}^s + u_{i,j,p,t}$, 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, five-year differences in both tariffs and trade ignores the timing of tariff shocks. A five-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 five-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 V. Online 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 such as local projections 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 endpoint of the difference. We will additionally show below that estimation approaches based predominantly on cross-sectional variation—without differencing—likely suffer from omitted variable problems. Thus, we argue that our long-run estimates are preferable to the conventional alternatives even for researchers only interested in the long-run elasticity for calibrating a static trade model.

C. 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. 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; Shapiro 2016; 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-specific determinants of bilateral product-level trade flows. After presenting the main results, Section IVB 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, nontariff barriers, 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 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).

Instrument.—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 contemporaneous imports from an important partner country, or lobbying by an important partner country.⁵ 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 to whom the MFN tariff change does not apply because they do not trade on MFN terms. These are countries in preferential trade agreements (PTAs).

Our baseline instrument is

$$(5) \quad \Delta_0 \ln \tau_{i,j,p,t}^{instr} = \mathbf{1}\{\tau_{i,j,p,t} = \tau_{i,j,p,t}^{appliedMFN}\} \times \mathbf{1}\{\tau_{i,j,p,t-1} = \tau_{i,j,p,t-1}^{appliedMFN}\} \\ \times \left(\ln \tau_{i,j,p,t}^{appliedMFN} - \ln \tau_{i,j,p,t-1}^{appliedMFN} \right)$$

together with the sample restriction that observations are dropped if both

$$\mathbf{1}\{\tau_{i,j,p,t} = \tau_{i,j,p,t}^{appliedMFN}\} \times \mathbf{1}\{\tau_{i,j,p,t-1} = \tau_{i,j,p,t-1}^{appliedMFN}\} = 1$$

(6) and

$$\mathbf{1}\{j \text{ is a major trading partner of } i \text{ in } t - 1 \text{ in aggregate}\} \\ + \mathbf{1}\{j \text{ is a major trading partner of } i \text{ in } t - 1 \text{ in product } p\} \\ + \mathbf{1}\{j \text{ is a major trading partner of } i \text{ in } t \text{ in aggregate}\} \\ + \mathbf{1}\{j \text{ is a major trading partner of } i \text{ in } t \text{ in product } p\} > 0.$$

We estimate equations (2) and (3) with $\Delta_0 \ln \tau_{i,j,p,t}^{instr}$ as an instrument for $\Delta_0 \ln \tau_{i,j,p,t}$ and equation (4) using $\Delta_0 \ln \tau_{i,j,p,t}^{instr}$ as the instrument for the h -year endogenous tariff change $\Delta_h \ln \tau_{i,j,p,t}$. The two indicator functions on the first line of (5) simply say that the applied MFN tariff is binding for the countries and product in question both in the preperiod $t - 1$ and the impact period t . The term $\ln \tau_{i,j,p,t}^{appliedMFN} - \ln \tau_{i,j,p,t-1}^{appliedMFN}$ is simply the log change in the tariff from $t - 1$ to t .

In addition, we impose a sample restriction specified in (6). In words, we drop from the sample the MFN trade flows in which exporter j is a major trading partner of importer i , either in terms of j 's total exports to i , or in terms of exports to i in product p , where p is an HS4 product group. In our baseline a partner is major when it is a top ten exporter to market i . We presume that endogeneity concerns that persist after the fixed effects will mostly apply to the importer's major MFN trading partners. Thus, major MFN partners, either in terms of total trade flows or product-level trade flows, are dropped from the sample. We stress that the classification into major and minor trading partners is from the perspective of each individual

⁵Lagged imports are included as a pretrend control in most of our specifications, so they do not pose a threat to identification.

importer and product. Section IVA 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.⁶ 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 instrumental variable (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 difference-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 “nontreated” (non-MFN or control group) 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. We include such multilateral resistance terms in all baseline specifications.

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

D. Institutional Background

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 percent in the United States, 10.0 percent in China, and 48.6 percent in India. The number of products covered by the bounds is also negotiated and varies by country. In many countries, including the United States and China, 100 percent of products are covered by the bounds. By contrast, 74 percent of products are subject to MFN bounds in India, and 50 percent in Turkey. The bounds themselves vary substantially across products. In the United States in 2015, about 40 percent of products had a bound of 0, while about one-tenth

⁶Because our differenced specification includes importer-exporter-HS4 fixed effects, the statement above should be strictly speaking interpreted as referring to deviations from importer-exporter-HS4-specific trends. We show in Section IVB and online Appendix Figure B4 that this aspect of our empirical strategy is not crucial to the results.

of products had bounds above 10 percent.⁷ 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 postcrisis. 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. Such tariff changes are potentially endogenous, necessitating both the inclusion of a rich set of fixed effects (to remove business cycles and broad partner-specific variation), and an identification strategy to deal with residual endogeneity.

II. Data and Basic Patterns

Our trade dataset is the BACI version of UN-COMTRADE, covering years 1995–2018 (Gaulier and Zignago 2010). The data contain information on the trade partners, years, and product codes at the HS six-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 (UNCTAD 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 nonlinear/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.⁸ Detailed documentation on the construction of our tariff data from the raw TRAINS data is available in online Appendix A and in the replication package (Boehm, Levchenko, and Pandalai-Nayar 2023). For robustness exercises, we also use data containing information on standard gravity variables such as distance, common border, common language, etc., from Centre d'Études Prospectives et d'Informations Internationales (2022b); Mayer and Zignago (2011).

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.⁹ 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

⁷Further details can be found in Bown and Crowley (2016). We are grateful to Chad Bown for useful suggestions and examples.

⁸Among the 148 WTO members in 2013, the median fraction of HS6 products covered by non-ad valorem tariffs is 0.01 percent, and the mean fraction is 1.76 percent (World Trade Organization 2014).

⁹As far as we are aware, there is no double counting of trade flows reported under different HS revisions.

we do not observe transaction-level trade, any such split will introduce composition effects into our tariff measures. For instance, 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 ten years, which typically corresponds to only two HS revision transitions. Online Appendix Table A1 provides the fraction of codes that map uniquely across revisions. In a single revision transition, on average 89 percent of product codes have a unique mapping.¹⁰ In a small number of instances, the meanings of HS4 (and therefore HS6) codes change across revisions, which would imply that importer-exporter-product fixed effect categories combine substantively different products across time periods. We manually identified those instances and eliminated them.

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.¹¹

Patterns in Tariff Changes.—Online Appendix Figure A1 plots histograms of tariff changes. It shows that while tariff decreases are more frequent, still a substantial share of tariff changes in our data are increases. 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. Online Appendix 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.

Online Appendix A presents additional summary statistics about our sample: (i) the average share of imports by destination (online Appendix Figure A3) and by

¹⁰Naturally, 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.

¹¹Section VA contains the complete discussion. As an example, if the underlying model is Armington, our long-run estimates would correspond to the elasticity in the CES aggregator $-\sigma$, while the trade elasticity inclusive of tariffs would be $1 - \sigma$.

product (online Appendix Figure A4), and (ii) the incidence of MFN and non-MFN trade (online Appendix Figure A5) in the sample.

Examples of the Treatment/Control Assignments.—Online 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 four-digit HS code 6403, “footwear; with outer soles of rubber, plastics, leather or composition leather and uppers of leather.” For three large importers (the United States, 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 and 2 list the ten 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 ten countries in this category). Columns 3 and 4 list the ten major trade partners in terms of aggregate trade. These countries are disqualified from the treatment group. Columns 5 and 6 list the ten 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 HS4 code.

After these countries are dropped from the treatment group, columns 7–9 list the treatment, control, and excluded groups. As the table highlights, for the US North American Free Trade Agreement (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 United States. 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 the control group. For Germany, for instance, EU member countries do not appear in the treatment group, and are only part of the control group.

III. Main Results

We begin by estimating the effects of a one-period tariff change on h -periods ahead trade flows and tariffs, as in equations (2) and (3), using our instrument as described above. For the baseline estimation, the product disaggregation for the fixed effects is at the HS4 level.¹² We also exclude trading partners based on a classification into major and minor at the importer-HS4 level. The left panel of Figure 1 reports the time path of tariff changes h periods after the initial 1 percent 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 percent of the change remains after five years, and approximately 75 percent after ten years. At the same time, the pattern shown in the figure is contrary to the hypothesis that our low elasticity estimates

¹²With 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.

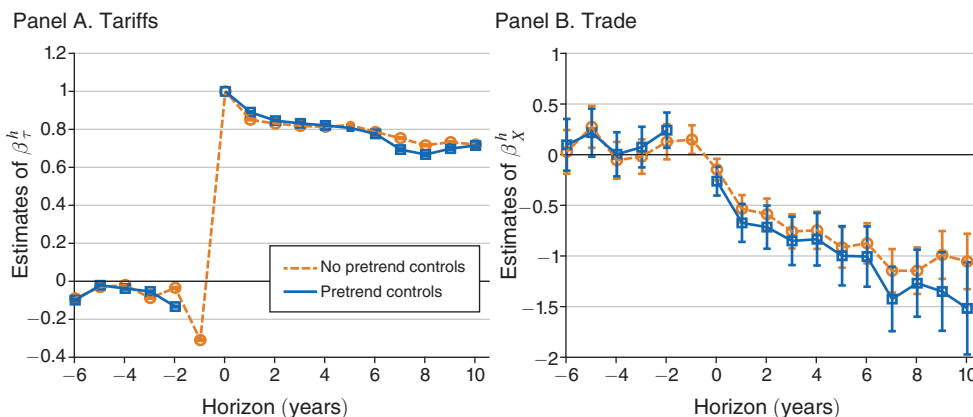


FIGURE 1. LOCAL PROJECTIONS: TARIFFS AND TRADE

Notes: This figure displays the results from estimating equations (2) and (3)—the local projection of h -period tariff growth (left panel) and h -period import growth (right panel) on the tariff change from $t - 1$ to t , instrumented with our baseline instrument (5). For negative time horizons the dependent variable is the one-year log change in tariffs (left panel) and trade (right panel). All specifications include exporter-HS4-year, importer-HS4-year, and exporter-importer-HS4 fixed effects. The specification with pretrend controls additionally includes log changes in tariffs from $t - 2$ to $t - 1$, instrumented with our baseline instrument, and log changes in trade from $t - 2$ to $t - 1$. The bars display 95 percent confidence intervals. Standard errors are clustered at the country-pair-product level.

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. These results illustrate the need for an estimation method that takes explicit account of the nontrivial time-series behavior of tariffs.

The figure suggests a presence of a pretrend. A tariff increase of 1 percent is preceded by a reduction of approximately 0.3 percent in the preperiod, reflecting again the negative first-order autocorrelation highlighted above. We thus include a lagged pretrend control for both tariffs and trade in our baseline estimates t throughout and instrument the lagged tariff change with a lag of our MFN tariff instrument. The blue lines in Figure 1 depict the estimates after including the pretrend 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 1 percent tariff change on trade flows. Trade falls gradually and stabilizes between 1 and 1.5 percent after seven to ten years. Unlike for tariffs, there is no evident pretrend in trade flows, regardless of whether we use pretrend 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 online Appendix Table B1 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




FIGURE 2. TRADE ELASTICITY: BASELINE AND ALL DATA/ALL TARIFFS 2SLS

Notes: This figure displays estimates of the trade elasticity based on specification (4), and including one lag of the changes in tariffs and trade as pretrend controls. Two-stage least squares (2SLS) estimates with the baseline instrument (5) are in blue, and all data/all tariffs 2SLS estimates are in red. All specifications include exporter-HS4-year, importer-HS4-year, and exporter-importer-HS4 fixed effects. The bars display 95 percent confidence intervals. Standard errors are clustered at the country-pair-product level.

the year following the tariff change is probably a better indicator of the short-run elasticity. At $h = 1$, the elasticity estimate is -0.76 . The 10-year estimate is -2.12 . Over the first seven years, the elasticity converges smoothly to the long-run value, and then is stable for years seven to ten. Online Appendix Table B2 reports the first stage F -statistics, which indicate that the instrument is very strong. The point estimates of the horizon-specific elasticities are displayed in the figure.

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 (4) on all the available data (i.e., without dropping major partners) and 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 I, 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 online Appendix Proposition C.1).¹³ At horizon 0, this approach amounts to a standard OLS estimation in differences. Note that because this strategy uses all tariff changes

¹³ Additionally, relying on higher frequency variation typically reduces confounding.

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 finding we revisit in Section IVB. 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 nontariff barriers increases coefficient estimates of their impact on 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 evolutions 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.

A. 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. Online 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 estimates 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) and 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 unrepresentative of world trade, if product groups with higher or lower elasticities predominate in the data. Online Appendix Figure B1 plots the baseline horizon-specific elasticities from Figure 2, together with the world-trade-weighted mean and median of the sector-specific elasticities reported in Figure 3. 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.

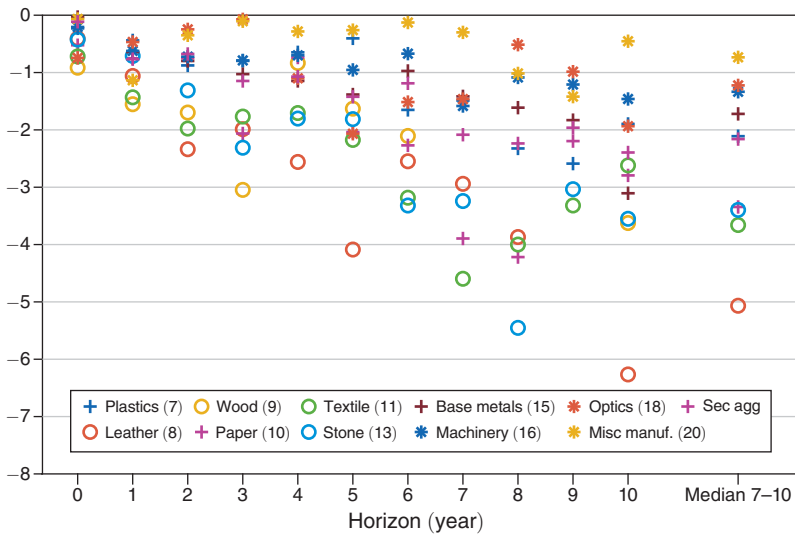


FIGURE 3. TRADE ELASTICITY: SECTORAL HETEROGENEITY

Notes: This figure displays the trade elasticity point estimates by HS section based on specification (4) and using the baseline instrument (5). All specifications include exporter-HS4-year, importer-HS4-year, and exporter-importer-HS4 fixed effects as well as one lag of the log change in tariffs and trade. Some HS sections are grouped into a single aggregate section “sec agg” as described in the text.

IV. 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 estimates are identified from broad variation representative of 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 artifact of nonrepresentative data or nonstandard estimation strategies.

A. Identifying Variation

One might be concerned that the coefficient estimates are identified from special and/or nonrepresentative 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, in 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

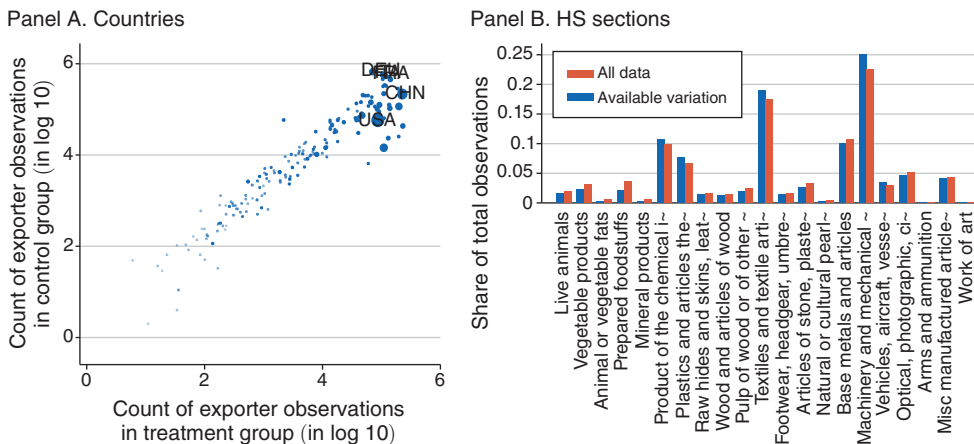


FIGURE 4. COUNTRY AND PRODUCT VARIATION

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-HS4-year, exporter-HS4-year, and importer-exporter-HS4 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 (blue bars), and the residualized sample after fixed effects have been taken out and the sample restrictions have been imposed (red bars).

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. 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. Online Appendix Figure B2 displays the frequency of country appearance in treatment or control group against 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 blue bars plot the shares of observations of all trade data. The red 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 unrepresentative trade flows. Online Appendix Figure B3 plots the frequency of different product groups in our residualized data at a finer level of sectoral disaggregation (HS2).

TABLE 1—ELASTICITY ESTIMATES: ALTERNATIVE APPROACHES

	Log levels			5-year log differences			10-year	
	No FE (1)	FE (2)	Multilateral + Bilateral FE (3)	OLS (4)	2SLS (5)	Baseline IV		
						(6)	(7)	(8)
$\ln\tau_{i,j,p,t}$	-3.70 (0.02)	-6.96 (0.05)	-1.04 (0.02)	-0.66 (0.02)	-0.47 (0.03)	-1.17 (0.11)	-1.24 (0.19)	-2.12 (0.33)
R^2	0.01	0.38	0.57	0.18				
Observations	107.1	106.2	104.9	38.1	38.1	21.4	16.7	8.3
First stage F					554,972	39,917	13,428	4,957
<i>Fixed effects</i>								
Imp \times HS4 \times year, Exp \times HS4 \times year		Yes	Yes	Yes	Yes	Yes	Yes	Yes
Imp \times exp \times HS4			Yes				Yes	Yes
Pretrend controls							Yes	Yes

Notes: This table compares alternative approaches of estimating trade elasticities. The dependent variables are log levels of trade values (columns 1–3) and log differences of trade flows (columns 4–8), and the independent variable of interest is the log of tariffs (columns 1–3), 5-year log differences of tariffs (columns 4–7), and the 10-year log-difference of tariffs (column 8). Column 1 reports the results with no fixed effects. Column 2 adds importer-HS4-year and exporter-HS4-year fixed effects. Column 3 further adds importer-exporter-HS4 fixed effects. Column 4 estimates the coefficient by OLS. Column 5 reports the all data/all tariffs 2SLS as explained in the text. Columns 6–8 present the results using our baseline IV. The specifications with pretrend controls additionally include log changes in tariffs from $t - 2$ to $t - 1$, instrumented with our lagged baseline instrument, and log changes in trade from $t - 2$ to $t - 1$. The reported R^2 s include the explanatory power of the fixed effects. Standard errors clustered by country-pair-product are in parentheses. Numbers of observations are reported in millions.

B. 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 of -5 to -10 (see for instance the review in Anderson and van Wincoop 2004). To get a sense of the range of existing values, online Appendix Table B3 summarizes the data, methods, and elasticity estimates of the set of papers closest to ours. These are the studies that use tariff variation in a gravity framework. Closest to our approach in terms of level of aggregation and spirit of the exercise are Hummels (2001); Hertel et al. (2007) and Fontagné, Guimbard and Orefice (2022) (disaggregate data, log-levels specification, some bilateral gravity controls); and Romalis (2007) (difference-in-differences in the cross section, HS6 data). Also closely related are the papers that use firm-level data and rely on variation in MFN versus non-MFN tariffs (Bas, Mayer, and Thoenig 2017; Fitzgerald and Haller 2018). All in all, the range is quite wide and encompasses our estimates. Our values are at the low end of the range in Romalis (2007), and broadly consistent with the horizon-specific aggregated elasticities in Fitzgerald and Haller (2018). There are a number of potential reasons we would not expect the estimates to line up exactly, including our substantially larger country sample, longer and nonoverlapping time period, use of local projections to estimate the entire time path while accounting for the timing of shocks, and use of an instrument.

Table 1 investigates the sources of these differences formally. Columns 1 and 2 of the table estimate the elasticity using a log-levels OLS specification, assuming all tariff variation is exogenous, similar to papers such as Head and Ries (2001). This specification, both without fixed effects and with the most commonly used multilateral resistance fixed effects (importer-product-time and exporter-product-time, as

in e.g., Hummels 2001), yields values between -3.7 and -7.0 , 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 and 4 include country-pair-product fixed effects and time-difference the data by five years, respectively. Such fixed-horizon differencing is similar to the approach used by Feenstra (1994), for instance. In both cases, the elasticity estimates fall sharply to around -1 . Thus, as recognized in several important contributions in the literature, controlling for unobserved bilateral determinants of trade either with fixed effects or by differencing is important for obtaining reliable estimates.

As discussed in Section IB, while long-differencing tackles the concern about bilateral unobservables, it ignores the timing of tariff shocks. Accounting for the timing of the shock might be important for estimating longer-run elasticities. Column 5 instruments the five-year change with the initial tariff change (“all-data/all-tariffs” estimation) to account for the timing of tariff changes. For closer comparison to column 4, we do not include pretrend controls. The estimates are similar to the five-year differencing. While quantitatively this adjustment does not change the estimates much, we caution that accounting for the timing of the shock may matter in other instances. As discussed in Section IB and online Appendix Proposition C.1, estimation in simple h -period differences does not generally identify the horizon- h trade elasticity.

Columns 6 and 7 report IV estimates using our baseline instrument, with and without the importer-exporter-HS4 effects and pretrend controls, respectively. The elasticity estimates have a tight range between -1.17 and -1.24 at the five-year horizon, demonstrating that neither the importer-exporter-HS4 fixed effects nor the pretrend controls play an important role in the differenced specifications (more on this below). Relative to the OLS estimates in column 4, and consistent with the main results above, instrumental variables push the estimates away from 0. Finally, column 8 presents the ten-year baseline estimates. The coefficient increases substantially in absolute value from -1.24 to -2.12 .

To summarize, there are four lessons from this exercise. First, controlling for bilateral unobservables drastically lowers the estimated elasticity, by a factor of about three to seven. Second, accounting for residual endogeneity by means of the instrument increases the estimated absolute value of the elasticity substantially, also evident in Figure 2. Using our baseline IV more than doubles the ten-year elasticity estimate. Third, the results in Figure 2 and Table 1 caution against using fixed-window differencing over a small number of years and interpreting the result as the long-run elasticity. In our estimates, the elasticity at ten years is 71.5 percent greater than at five years, indicating that the adjustment to the trade cost shock is not completed after five years. Estimation in levels does not solve this issue. As evident in column 3 of Table 1, a levels specification with bilateral fixed effects still produces an elasticity that is much too low when interpreted as a long-run elasticity. And fourth, simple differencing does not take into account the time path of tariffs, and therefore can produce incorrect elasticities depending on the timing of the shocks. While our local projections approach takes this explicitly into account, this adjustment does not make an economically significant difference in the case of the five-year horizon illustrated in Table 1.

TABLE 2—LOG-LEVEL ELASTICITY ESTIMATES VARYING THE FINENESS OF THE IMPORTER-EXPORTER EFFECTS

	No bilateral (1)	Gravity variables (2)	Country-pair (3)	Country-pair × HS2 (4)	Country-pair × HS3 (5)	Country-pair × HS4 (6)
$\ln \tau_{i,j,p,t}$	-6.96 (0.05)	-2.11 (0.02)	-1.39 (0.02)	-1.21 (0.03)	-1.17 (0.02)	-1.04 (0.02)
R^2	0.38	0.45	0.48	0.50	0.52	0.57
Observations	106.2	104.9	106.2	106.1	105.9	104.9
<i>Fixed effects and controls</i>						
Imp × HS4 × year, Exp × HS4 × year	Yes	Yes	Yes	Yes	Yes	Yes
<i>Bilateral gravity controls</i>						
Imp × exp		Yes	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 the log of trade values, and the independent variable of interest is the log of tariffs. Column 1 reports the results with importer-HS4-year and exporter-HS4-year fixed effects. Column 2 additionally includes common gravity variables as controls (distance, contiguity, common language, and colonial relationship). Columns 3–6 replace the observable bilateral gravity variables with progressively finer bilateral fixed effects, from importer-exporter to importer-exporter-HS4. The reported R^2 s include the explanatory power of the fixed effects. Standard errors clustered by country-pair-product are in parentheses. Numbers of observations are reported in millions.

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-level specifications in columns 1–3 of Table 1 have quite similar sample sizes, as 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 nonzero 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. Online Appendix Table B4 replicates Table 1 on a sample that is constant across columns. If anything, the difference between the cross-sectional estimates in columns 1 and 2 and the fixed effects/differenced specifications in columns 3 and 4 is even starker, as the cross-sectional variation implies even larger elasticity estimates in this subsample (as high as 8–11). Section IVC 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 in log levels without any bilateral fixed effects from Table 1. Column 2 adds the traditional gravity controls (distance, common border, common language, and colonial relationship). Columns 3–6 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 compared to the specification with no bilateral variables, the majority of the fall in the estimated elasticity comes from adding either the usual gravity controls, or the very coarsest bilateral fixed effect at the country-pair level (column 1 versus column 3). 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.^{14, 15}

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. Online Appendix Figure B4 plots the time path of trade elasticity estimates under different 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.

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 eight-digit or nine-digit 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

¹⁴While traditional gravity controls also lead to a large drop in the estimated elasticity, the bilateral fixed effects have the advantage of additionally absorbing unobserved gravity variables and bilateral preferences, which might be only imperfectly correlated with observed gravity variables.

¹⁵Online Appendix Table B5 replicates Table 2 with multilateral resistance terms at the HS6 rather than the HS4 level. The results are virtually identical.

to several countries, and are 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.

Multilateral Resistance.—Finally, we explore the impact of controlling for multilateral resistance terms on our estimates. This part of the empirical model is the least controversial, as it has been universally recognized since Anderson and van Wincoop (2003) that multilateral resistance terms are essential in gravity specifications. Nonetheless, it is still informative to know how estimates change if we depart from the conventional approach. Online Appendix Figure B5 displays the time paths of the elasticity estimates for the baseline, no multilateral resistance terms, and the multilateral resistance terms at intermediate levels of product disaggregation (importer/exporter-year, importer/exporter-year-HS2 and importer/exporter-year-HS3). What emerges is that it is important to control for some multilateral resistance terms. Including no multilateral resistance terms at all leads to larger elasticity estimates, as large as -4 at the eight- to nine-year horizon. Including any multilateral resistance fixed effects sharply lowers the elasticity estimates in absolute value. In fact, coarser fixed effects lead to even lower elasticities than the baseline.

C. Robustness

Pretrends 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 pretrend in tariffs. We account for differential pretrends 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 and 3 of Table 3 report results with no lags and five 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 online Appendix Table B1 reports the results of local projections of tariffs and trade flows directly on the initial tariff change, as in (2) and (3), while allowing for one, zero and five lags. Once again, the point estimates change little when adding lags.

A distinct concern is anticipation effects. Even if pretreatment 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 pretrends in trade flows. Figure 1 shows no evidence of pretrends in trade flows 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

TABLE 3: TRADE ELASTICITY, ROBUSTNESS: PRETRENDS, ALTERNATIVE CLUSTERING, ALTERNATIVE SAMPLES

	Baseline (1)	No lags (2)	Five Lags (3)	FE50 (4)	Two-way clustering (5)	Constant sample (6)	Extensive case 1 (7)	Extensive case 2 (8)
t	-0.26 (0.07)	-0.15 (0.05)	0.17 (0.14)	-0.23 (0.09)	-0.26 (0.10)	-0.59 (0.29)	0.02 (0.04)	-0.08 (0.06)
Observations	31.7	41.5	14.6	17.6	31.7	5.0	131.0	56.4
$t + 1$	-0.76 (0.11)	-0.63 (0.08)	-0.13 (0.21)	-0.60 (0.13)	-0.76 (0.14)	-0.10 (0.38)	-0.48 (0.06)	-0.81 (0.09)
Observations	26.2	32.8	12.5	15.2	26.2	5.0	108.1	49.1
$t + 3$	-1.02 (0.15)	-0.93 (0.10)	-0.63 (0.31)	-0.86 (0.17)	-1.02 (0.20)	-0.89 (0.47)	-0.59 (0.08)	-0.95 (0.11)
Observations	20.8	26.2	9.8	12.5	20.8	5.0	87.2	41.5
$t + 5$	-1.24 (0.19)	-1.11 (0.12)	-1.15 (0.43)	-1.01 (0.22)	-1.24 (0.25)	-0.92 (0.44)	-0.73 (0.10)	-1.18 (0.12)
Observations	16.7	21.1	7.3	10.2	16.7	5.0	69.7	34.6
$t + 7$	-2.06 (0.23)	-1.52 (0.15)	-2.33 (0.59)	-1.85 (0.27)	-2.06 (0.36)	-0.99 (0.49)	-0.90 (0.11)	-1.50 (0.14)
Observations	13.2	17.0	5.2	8.2	13.2	5.0	55.2	28.4
$t + 10$	-2.12 (0.32)	-1.46 (0.19)	-2.55 (1.02)	-1.76 (0.37)	-2.12 (0.33)	-1.82 (0.54)	-0.94 (0.15)	-1.64 (0.18)
Observations	8.3	11.3	3.2	5.2	8.3	5.0	35.1	19.2

Notes: This table presents robustness exercises for the results from estimating equation (4). All specifications include importer-HS4-year, exporter-HS4-year, and importer-exporter-HS4 fixed effects, and the baseline pretrend controls (one lag of each the log change in tariffs and trade) 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 per cluster. Column 6 restricts the sample to a constant sample across horizons. Column 7 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 every positive trade. Column 8 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. Standard errors are clustered at the importer-exporter-HS4 level, except in column 5 where they are additionally clustered by year. Observations are reported in millions.

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. Column 7 of online Appendix Table B6 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 percent of the initial shock surviving ten years.¹⁶ 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 General Agreement on Tariffs and Trade/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,

¹⁶Consistent with our estimates, Bown and Crowley (2014) document that most MFN tariff changes below bounds are permanent or very persistent.

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 (online Appendix Table B8). 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.

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 2021; 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.¹⁷

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 percent of the overall growth in North-American trade following NAFTA implementation. While instances of rapid economic growth and structural change—such as the Republic of Korea—can be associated with a contribution of the product-level extensive margin as high as 25 percent, the extensive margin plays a negligible role in trade growth under more conventional circumstances (such as US-UK trade).

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.¹⁸

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.

¹⁷ 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 changes in trade. Fitzgerald and Haller (2018) estimate that in Ireland, the contribution of the firm-level extensive margin to the long-run elasticity of trade to tariffs is less than 10 percent.

¹⁸ 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.

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.¹⁹ 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 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 be 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).²⁰

The resulting estimates in columns 7 and 8 of Table 3 can be interpreted as the elasticity inclusive of both the intensive and product-level extensive margins. When including more zeros (column 7), the point estimates are similar to the baseline initially, and 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 8 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 ten-year elasticity including the extensive margin is slightly higher (-1.64) than the corresponding intensive margin instrumented specification without pretrend controls (-1.46).

Alternative Instruments, Outcome Variables, and Samples.—The baseline instrument excludes the top ten largest trading partners from the treatment group. 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”). As in the baseline, the instrument is 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. Column 3 implements a different cutoff for major partner, by only classifying the top five importers as major. The resulting long-run elasticity estimate of -1.5 is between the baseline and the version in which none of the major partners are

¹⁹That is, if country A ever exports any product in HS one-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.

²⁰Including pretrend 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 pretrend 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 2,000). 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 pretrend controls thus gives the extensive margin maximum chance to produce larger elasticities relative to the baseline.

TABLE 4—TRADE ELASTICITY, ROBUSTNESS: ALTERNATIVE INSTRUMENTS, OUTCOMES, AND SAMPLES

	All data/MFN		Top 5 maj.		Unit Values	Weighted	SD1	PTA
	Baseline	tariffs	Partners	Quantities				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
t	-0.26 (0.07)	-0.28 (0.03)	-0.24 (0.05)	-0.18 (0.09)	-0.05 (0.06)	-0.24 (0.08)	-0.31 (0.10)	-0.24 (0.07)
Observations	31.7	57.1	39.0	31.7	31.7	31.6	28.7	31.8
$t + 1$	-0.76 (0.11)	-0.62 (0.05)	-0.67 (0.08)	-0.66 (0.14)	-0.03 (0.08)	-0.82 (0.12)	-0.89 (0.14)	-0.73 (0.11)
Observations	26.2	47.2	32.1	26.2	26.2	26.2	23.8	26.3
$t + 3$	-1.02 (0.15)	-0.65 (0.06)	-0.86 (0.10)	-0.81 (0.18)	-0.13 (0.10)	-1.05 (0.16)	-1.23 (0.20)	-0.97 (0.15)
Observations	20.8	38.2	25.6	20.8	20.8	20.8	18.9	20.9
$t + 5$	-1.24 (0.19)	-0.72 (0.07)	-1.08 (0.12)	-1.42 (0.23)	0.29 (0.13)	-1.21 (0.21)	-1.18 (0.25)	-1.12 (0.19)
Observations	16.7	30.9	20.6	16.7	16.7	16.7	15.1	16.8
$t + 7$	-2.06 (0.23)	-0.94 (0.09)	-1.53 (0.16)	-2.17 (0.29)	0.16 (0.16)	-2.14 (0.26)	-1.99 (0.32)	-1.97 (0.23)
Observations	13.2	24.6	16.3	13.2	13.2	13.2	12.0	13.3
$t + 10$	-2.12 (0.32)	-0.87 (0.12)	-1.48 (0.22)	-1.76 (0.41)	-0.08 (0.22)	-2.37 (0.37)	-2.36 (0.44)	-2.08 (0.33)
Observations	8.3	15.9	10.3	8.3	8.3	8.3	7.5	8.4

Notes: This table presents alternative estimates for the results from estimating equation (4), varying the instrument, outcome variable, or sample. All specifications include importer-HS4-year, exporter-HS4-year, and importer-exporter-HS4 fixed effects, and the baseline pretrend controls (one lag of each the log change in tariffs and trade). Column 2 uses an alternative sample where all trade partners subject to the MFN regime are included. Column 3 presents results where the sample excludes only the top five major MFN trade partners. Column 4 reports results for quantities, and column 5 the results for unit values. Column 6 presents results for a weighted specification where $t - 1$ log trade values are used as weights. Column 7 reports the results based on a sample where tariffs do not vary within an importer-exporter-HS6-year observation. Column 8 presents results where we assign observations covered by a PTA listed in the WTO PTA database to the control group. Standard errors are clustered at the importer-exporter-HS4 level. Observations are reported in millions.

dropped (column 2), which is intuitive. Columns 4 and 5 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. Column 6 implements a weighted regression, with weights given by the initial log imports. The estimates are very similar to the baseline.

Column 7 of Table 4 estimates the elasticity on a sample where tariffs do not vary within an importer-HS6. This specification drops importer-product instances where tariffs are set at finer levels of disaggregation, such as HS8 or HS10. Again, the results are very similar to the baseline at all horizons. In the baseline analysis, we place a country in the control group if its applied tariff is below the MFN tariff in either period $t - 1$ or period t . If an applied tariff is equal to the MFN tariff in period $t - 1$ and period t , we assume it trades on MFN terms. This observation is then either in the treated group, or in the excluded (dropped) group if it is one of the top ten trading partners. Thus, we do not use outside information on PTA membership to place countries in the treated/excluded or control groups. There is

a possibility, then, that a country-pair-product is technically in a PTA, but the PTA tariff coincides with the MFN tariff in both $t - 1$ and t . This is a gray area, in the sense that these observations trade on de facto MFN terms. Without knowledge of the political process that led these PTA tariffs to coincide with MFN tariffs, we cannot be sure whether to assign them to the control group or if these more closely resemble the treated/excluded groups. The baseline analysis assigns these observations to the treated (if minor) or excluded (if major partner) groups. A reasonable alternative is to assign them to the control group on account of the fact that they are legally PTA observations. The best available information on product-specific PTA tariffs comes from the WTO Tariff Download Facility (World Trade Organization Integrated Database and Consolidated Tariff Schedules 1995–2018). We merged these data with ours, and used it to reclassify those instances into the control group. Column 8 presents the results. The elasticity estimates are quite similar to the baseline.²¹ Finally, column 9 of online Appendix Table B7 drops the instances in which a trade flow is subject to a temporary trade barrier (TTB), such as antidumping, countervailing, or safeguard duties. The data on TTBs come from Bown (2011), updated to 2019 by the World Bank (2021). Dropping observations covered by the TTBs leaves the results virtually unchanged.²²

Additional Results, Diagnostics, and Robustness.—Online Appendix B presents further robustness and diagnostics. Online Appendix Tables B1, B6, and B7 report the results for all the specifications at every horizon. Online Appendix Table B2 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 ten. Columns 3 and 4 of online Appendix Table B8 report results for the elasticity estimated with the multi-lateral resistance terms at the HS6 level. The estimates are somewhat smaller than the baseline, though the sample shrinks and the standard errors widen.²³ Column 5 of online Appendix Table B8 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 nonmissing 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.²⁴

²¹ We do not adopt this approach as the baseline in part because the WTO product-specific PTA tariff data are self-reported and turned out to be highly incomplete.

²² We 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.

²³ Note that in our baseline estimation, time differencing already eliminates importer-exporter-HS6 fixed effect in levels.

²⁴ Formally, we estimate the equation $\Delta_0 \ln X_{i,j,p,t} = \sum_{k=0}^{10} \gamma^k \Delta_0 \ln \tau_{i,j,p,t-k} + \delta_{i,p,t}^d + \delta_{j,p,t}^s + \delta_{i,j,p}^b + u_{i,j,p,t}$ instrumenting $\Delta_0 \ln \tau_{i,j,p,t-k}$ with $\Delta_0 \ln \tau_{i,j,p,t-k}^{instr}$, for all k . The trade elasticity at horizon h reported in online Appendix Table B8 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 online Appendix Table B2. 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

V. Theory and Applications

We stress that equations (2), (3), and (4) are not tied to a particular theory, 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 partial equilibrium (PE) dynamic model of sluggish adjustment to trade cost shocks. PE is a natural starting point. Since our econometric estimates identify a partial elasticity this framework fits tightly with the empirics. The recent literature on trade dynamics is rich in both substantive mechanisms and quantification (see, among many others, Costantini and Melitz 2007; Ruhl 2008; Drozd and Nosal 2012; Burstein and Melitz 2013; Alessandria and Choi 2014; Alessandria, Choi, and Ruhl 2021; Ruhl and Willis 2017; Blaum 2019; Fitzgerald, Haller, and Yedid-Levi 2016; Leibovici and Waugh 2019; Alessandria, Arkolakis, and Ruhl 2021; Steinberg 2022). 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. Online 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 develop a dynamic multicountry multisector general-equilibrium (GE) model, that embeds a special case of this simple PE framework. We use this GE model to illustrate that our estimated partial elasticities are key for disciplining the total (GE) responses of trade flows to shocks. 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.

A. Dynamics of Trade Elasticities

Setup.—The minimalist model that can capture differing trade elasticities in the short versus 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 PE decisions of firms from one market selling to another, and thus suppress importer, exporter, and product subscripts. Consistent with the gravity tradition, GE objects such as domestic unit costs or foreign demand shifts are absorbed by country-product-time fixed effects, and thus we ignore GE forces in most of this

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 .

subsection. 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 versus long-run distinction, we assume that p_t^x and q_t adjust instantaneously to tariff changes, whereas n_t is predetermined by one period, and can only change from the next period onwards. Quantity and price are functions 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 $\pi(\tau_t)$ per unit mass n_t . Define the following elasticities:

$$(7) \quad \eta_{q,p} := \frac{\partial \ln q}{\partial \ln p^x}, \quad \eta_{q,\tau} := \frac{\partial \ln q}{\partial \ln \tau}, \quad \eta_{p,\tau} := \frac{\partial \ln p^x}{\partial \ln \tau}, \quad \eta_{\pi,\tau} := \frac{\partial \ln \pi}{\partial \ln \tau},$$

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:

$$(8) \quad v_t = \frac{1}{1+r} E_t [\pi_{t+1} + (1-\delta)v_{t+1}],$$

$$(9) \quad n_t = n_{t-1}(1-\delta) + G(v_{t-1}),$$

subject to the transversality condition $\lim_{t \rightarrow \infty} [(1-\delta)/(1+r)]^t v_t = 0$. The forward-looking equation (8) states that the value of exporting v_t is the expected present value of future flow profits from exporting. The backward-looking equation (9) 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 (8) and (9) can be solved sequentially. For any stochastic process for tariffs $\{\tau_t\}_{t=0}^\infty$, equation (8) can be solved forward to obtain

$$(10) \quad v_t = \frac{1}{1+r} E_t \left[\sum_{k=0}^\infty \left(\frac{1-\delta}{1+r} \right)^k \pi(\tau_{t+k+1}) \right].$$

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 (9) backward,

$$(11) \quad n_t = \sum_{\ell=0}^{t-1} (1-\delta)^\ell G(v_{t-1-\ell}) + (1-\delta)^t n_0.$$

For a given initial value of n_0 and a stochastic process for tariffs $\{\tau_t\}_{t=0}^\infty$, equations (10) and (11) and elasticities (7) 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 (7) 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 , $\eta_{p,\tau} = -\partial \ln \bar{\varphi} / \partial \ln \tau$, where $\bar{\varphi}$ 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 exporting 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. Online Appendix C.2 provides a detailed discussion of the 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

$$(12) \quad \varepsilon^0 := \frac{d \ln X_{t_0}}{d \ln \tau_{t_0}} = (1 + \eta_{q,p}) \eta_{p,\tau} + \eta_{q,\tau}$$

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 (10) becomes $v = \pi(\tau) / (\delta + r)$. Equation (11) then implies that n_t monotonically converges to $n = G(v) / \delta$.²⁵ It follows that $d \ln n / d \ln \tau = \chi \eta_{\pi,\tau}$ where $\chi := g(v) v / G(v)$. These two expressions characterize the nonstochastic steady state of the model. Hence, the long-run trade elasticity is

$$(13) \quad \varepsilon := \frac{d \ln X}{d \ln \tau} = \varepsilon^0 + \frac{d \ln n}{d \ln \tau} = \varepsilon^0 + \chi \eta_{\pi,\tau}$$

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

²⁵The convergence of n_t to its steady state value is geometric and monotone. The rate of convergence is δ . We provide details in online Appendix C.3.

sensitive are profits to tariffs, the greater the absolute value of the long-run trade elasticity.

The long-run trade elasticity increases (in absolute value) in the elasticity χ of mass n with respect to value v . The 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, $\chi = [f''(a)a/f'(a)]^{-1}$. 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 $\{d\ln\tau_{t_0+h}/d\ln\tau_{t_0}\}_{h=0}^\infty$. 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 (10).

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

$$(14) \quad \frac{d\ln X_{t_0+h}}{d\ln\tau_{t_0}} = \varepsilon^0 \frac{d\ln\tau_{t_0+h}}{d\ln\tau_{t_0}} + \frac{d\ln n_{t_0+h}}{d\ln\tau_{t_0}}.$$

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

$$(15) \quad \varepsilon^h := \frac{\frac{d\ln X_{t_0+h}}{d\ln\tau_{t_0}}}{\frac{d\ln\tau_{t_0+h}}{d\ln\tau_{t_0}}} = \varepsilon^0 + \frac{\frac{d\ln n_{t_0+h}}{d\ln\tau_{t_0}}}{\frac{d\ln\tau_{t_0+h}}{d\ln\tau_{t_0}}}$$

as long as this object is finite (i.e., $d\ln\tau_{t_0+h}/d\ln\tau_{t_0} \neq 0$). Note that this definition of the horizon- h trade elasticity coincides with equation (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 (15), the adjustment of n_t to the tariff shock.

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

$$(16) \quad \frac{d\ln n_{t_0+h}}{d\ln\tau_{t_0}} = \chi\eta_{\pi,\tau} \frac{\delta + r}{1 + r} \delta \sum_{k=0}^{h-1} (1 - \delta)^{h-1-k} E_{t_0+k} \left[\sum_{\ell=0}^\infty \left(\frac{1 - \delta}{1 + r} \right)^\ell \frac{d\ln\tau_{t_0+k+\ell+1}}{d\ln\tau_{t_0}} \right].$$

PROOF:

See online Appendix C.4.

Plugging (16) into (15) delivers the horizon- h trade elasticity. As is clear from equations (14) and (15), 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 16), 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 $\ell = 1, 2, \dots$. 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 $\lim_{h \rightarrow \infty} (d \ln \tau_{t_0+h} / d \ln \tau_{t_0}) \neq 0$ and is finite, then $\lim_{h \rightarrow \infty} \varepsilon^h = \varepsilon$.*

PROOF:

See online 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 One Period): Let there be a surprise change in the tariff sequence of the form $\{d \ln \tau_{t_0+h} / d \ln \tau_{t_0}\}_{h=0}^{\infty} = \{1, \Delta \ln \tau_{>t_0}, \Delta \ln \tau_{>t_0}, \Delta \ln \tau_{>t_0}, \dots\}$. 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

$$(17) \quad \varepsilon^h = \varepsilon^0 + \chi \eta_{\pi, \tau} [1 - (1 - \delta)^h],$$

with ε^0 given by (12). 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

$$(18) \quad \varepsilon^h = \varepsilon^0 + \chi\eta_{\pi,\tau} \frac{(\delta + r)\delta}{\left[1 + r - (1 - \delta)\rho\right]\left(1 - \frac{1 - \delta}{\rho}\right)} \left[1 - \left(\frac{1 - \delta}{\rho}\right)^h\right].$$

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 one, the horizon- h trade elasticity in the AR(1) case (18) converges pointwise to the horizon- h trade elasticity under a permanent tariff change (17). This property is important for our empirical application. Although tariff changes in our sample retain 75 percent of their initial impulse ten 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), where*

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

and 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 (3).

The fixed effects $\delta_{j,p,t}^{s,X,h}$ and $\delta_{i,p,t}^{d,X,h}$ capture a weighted sum of past, present, and expected future changes in interest rates, demand, the cost of production, the cost of entry, and nontariff trade barriers that vary at the exporter-product-time (j, p, t) and importer-product-time (i, p, t) level, respectively, in model extensions in which these vary over time. The error term includes past, present, and expected future time-varying importer-exporter-product-specific demand shocks and nontariff trade barriers, as well as the initial state.

PROOF:

See online Appendix C.4.

We stress that this proposition extends to GE settings. The changes in supply and demand absorbed by the multilateral resistance terms include both exogenous (shocks), and endogenous GE changes in prices and aggregate consumption. As a result, the econometric estimates identify a partial elasticity of trade with respect to trade costs holding these terms constant. The intuition for this result is similar to conventional static microfounded gravity equations. Online Appendix D.4 presents a fully articulated special case of this proposition in the context of our GE model and clarifies the economic interpretation of the objects absorbed by the exporter-product-time and importer-product-time fixed effects.

Trade Elasticities to Tariffs and Nontariff Barriers.—Before moving to the quantification, we note a distinction that will matter for connecting the model to the data. Our estimates are of trade elasticities with respect to tariffs, and in our data trade flows do not include tariff payments. Theoretical models usually also include nontariff iceberg trade barriers such as transport costs. The elasticity of tariff-exclusive trade flows with respect to nontariff trade barriers is in general related to but distinct from the tariff elasticity. This is because nontariff iceberg trade costs shift prices received by the exporter, whereas tariffs do not. Letting κ_t denote the nontariff iceberg trade cost and c_t denote the domestic marginal cost, in the CES model the price received by the exporter is $p_t^x = [\sigma/(\sigma - 1)]\kappa_t c_t$, and the quantity produced is $q_t = (\tau_t p_t^x)^{-\sigma} D_t$, where D_t is the demand shifter. As a result, the elasticity of tariff-exclusive trade flow per unit mass of firms $p_t^x q_t$ is $-\sigma$ with respect to tariffs, and $1 - \sigma$ with respect to nontariff trade costs. By the same token, those two elasticities would coincide for trade flows inclusive of tariff payments.

Thus, in many static models the mapping between our estimates and the elasticity of trade to nontariff trade barriers is particularly simple: we should simply add one to our estimates. This will be relevant when we apply our estimates to the Arkolakis, Costinot, and Rodríguez-Clare (2012) gains from trade exercise in Section VC. The mapping between the two elasticities is more complex in dynamic models, and online Appendix C.6 provides details for the class of models considered in this section. With CES demand, the short-run (respectively long-run) elasticity to tariffs is $-\sigma$ (respectively $-\sigma(1 + \chi)$) while the elasticity to nontariff trade barriers is $1 - \sigma$ (respectively $(1 - \sigma)(1 + \chi)$). Importantly, our estimates of the elasticity of trade flows to tariffs can be used to infer the nontariff trade elasticity in many static and dynamic models.²⁶

Quantification.—Next, we explore the time path of tariff 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 (13), 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 $\varepsilon = -2$. We further set the depreciation rate to $\delta = 0.25$ to roughly match the

²⁶Our estimates are informative about trade elasticities to tariffs and nontariff trade barriers when trade costs take the iceberg form. Mappings to models with non-iceberg trade costs would have to be considered on a case-by-case basis.

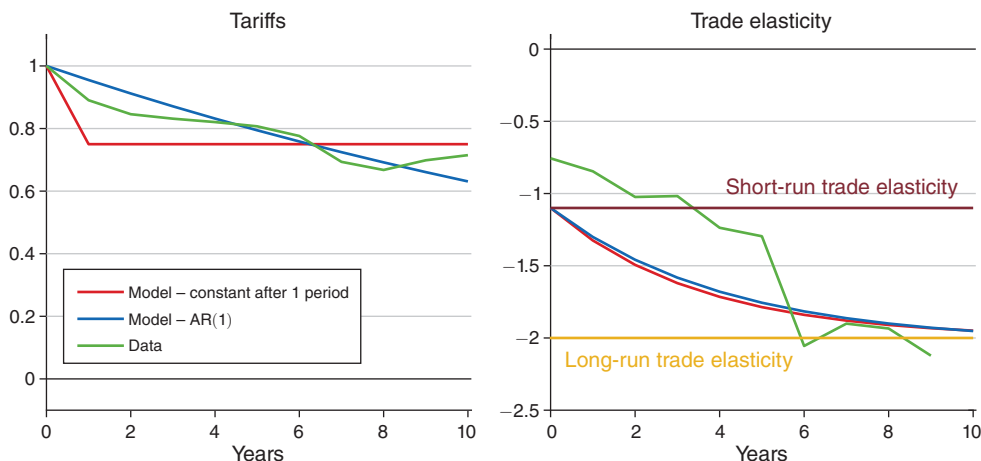


FIGURE 5. TIME PATH OF ELASTICITIES IN THE DYNAMIC MODEL

Notes: This figure illustrates the trade elasticities as implied by the model in Examples 1 and 2, and compares them to the baseline estimates. The parameters are set as follows: $\sigma = 1.1$, $\chi = 0.82$, $\delta = 0.25$, $r = 0.03$, and $\rho = 0.955$.

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 onward. 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 $\sigma > 1$.²⁷

²⁷ 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_i^x falls when a tariff increases, helping push down the trade elasticity all else equal. However, the flip side of a fall in p_i^x is a ceteris paribus increase in the quantity imported. It turns out there is no combination of $\sigma > 1$ and local distribution cost share between 0 and 1 that delivers

B. General Equilibrium

The setup above is in PE, which allows for a precise mapping to the empirical estimation and unified analytical results in a number of environments. It is well understood that the GE response of trade will not coincide with the PE response, and thus the partial elasticities that we estimate do not capture the total change in trade following a trade cost shock. However, these estimates are a crucial input for disciplining the parameters of GE models. While this is well understood in static trade models, it is an open question to what extent the partial elasticity estimated in the data matters for the GE trade response in dynamic models.

To show that our estimates are important for disciplining GE models, we take one of the PE models laid out above—the dynamic Krugman model—and embed it in GE. Online Appendix D lays out the details of the model and the calibration. In the quantitative implementation we limit the size of the economy to six countries with five sectors for computational reasons. The model is calibrated to standard data on import and expenditure shares from the Integrated Industry-Level Product Account (KLEMS) and the World Input-Output database. The calibration is summarized in online Appendix Table D1.

Figure 6 displays the impulse responses of US imports for sets of unexpected and permanent tariff shocks. The panels differ in how broad-based the tariff shock is, starting from the most localized in panel A to the most pervasive in panel D. Panel A reports the responses of trade flows to a 1 percent tariff shock on one product from one importer, while panel D displays the results for an across-the-board 1 percent tariff shock on all US imports from all source countries.

The solid blue lines and the dashed red lines display the PE responses under different trade elasticities. The blue line is calibrated to our estimates, setting $\sigma = 1.1$ and targeting a long-run elasticity of 2. Thus, it matches the impulse response reported in Figure 5. The dashed red line is instead calibrated to match the long-run partial elasticity of 6, more common in the trade literature, and assumes a short-run elasticity of $\sigma = 3$.²⁸ By construction, since these are partial elasticities, the solid blue and dashed red lines are the same in every panel.

The blue (respectively red) shaded areas are the ranges, across source countries and sectors, of the GE trade responses to the same tariff shocks under the two alternative calibrations. The main finding is that the GE trade responses are very different across elasticities. That is, the parameter values required to match a particular partial elasticity matter a great deal for the GE responses of trade. Not surprisingly, higher partial elasticities translate to higher GE elasticities. This main conclusion is not sensitive to whether we consider isolated or pervasive tariff shocks. Thus, our estimates are informative and quantitatively important in GE settings. Online

a less than unitary trade elasticity as we measure it (of $p_i^x q_i$ with respect to τ_i). In addition, Table 4 shows a virtually nil response of p_i^x to tariffs, a finding consistent with recent estimates using the US-China trade war (Fajgelbaum et al., 2020; Cavallo et al., 2021). 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. (2021) document a fall in retail margins for US imports affected by the trade war.

²⁸A long-run tariff-exclusive elasticity of -6 , as in Figure 6, implies a long-run nontariff iceberg trade cost elasticity of -5 , a common value in the trade literature (e.g., Costinot and Rodríguez-Clare 2014). See also the discussion in Section VA.

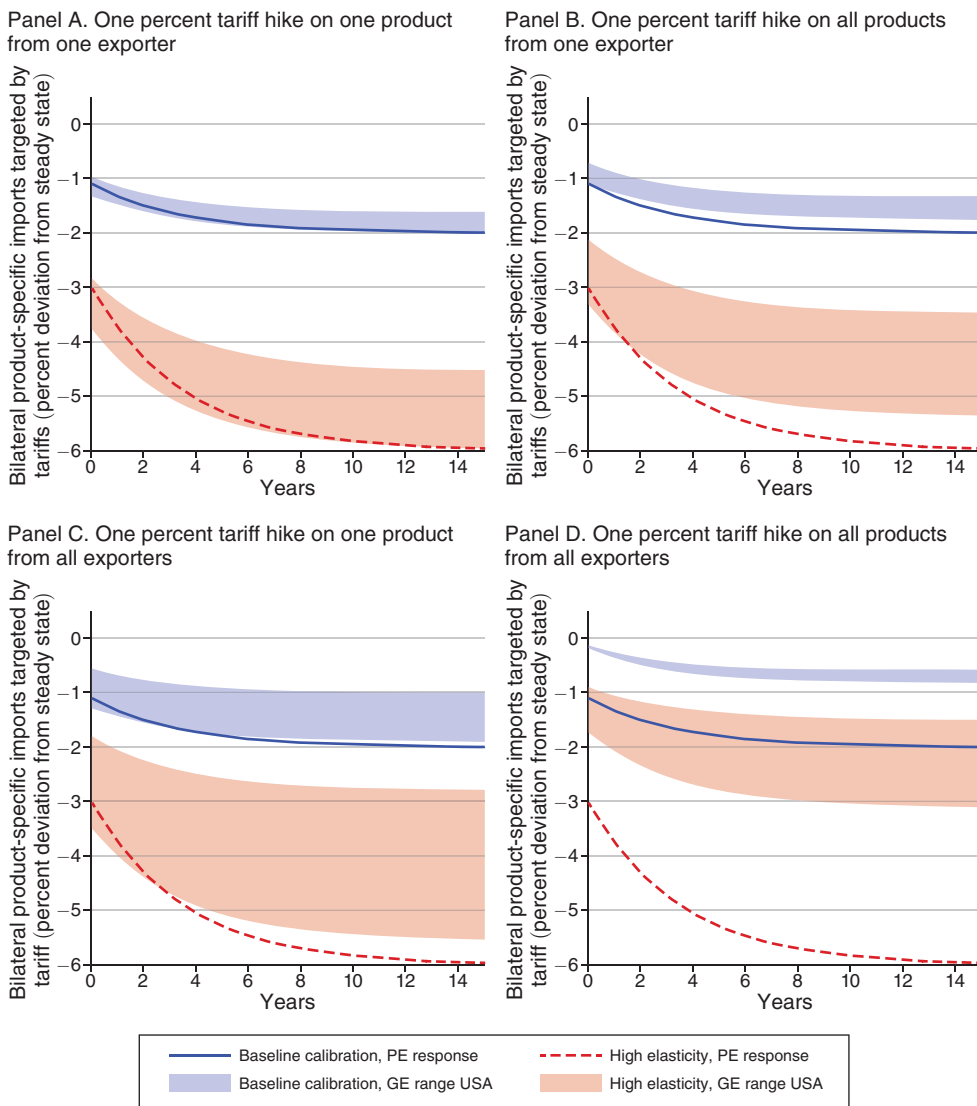


FIGURE 6. GENERAL EQUILIBRIUM TRADE RESPONSES: US IMPORTS

Notes: This figure reports the impulse responses of US imports to unexpected and permanent 1 percent tariff hikes in partial equilibrium (solid blue lines and dashed red lines) and in general equilibrium (shaded areas). The shaded areas represent the ranges of impulse response functions in GE taken over exporters and sectors. In the baseline calibration $\sigma = 1.1$ and $\chi = 0.82$, so that the long-run elasticity $\varepsilon = 2$. In the high elasticity calibration, $\sigma = 3$ and $\chi = 1$, so that $\varepsilon = 6$. See online Appendix Table D1 for details on the calibration.

Appendix Figure B6 reports analogous results for Canada, a smaller economy than the United States. The results are very similar.

C. 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 static 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 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. ACR formulas normally apply the tariff-inclusive elasticity, or alternatively, the elasticity of trade with respect to nontariff iceberg costs. To translate our tariff-exclusive estimates to the welfare-relevant elasticity θ we must add 1, as discussed in Section VA. Thus, our estimates imply that the welfare-relevant elasticity θ is around -1 .

Figure 7 displays the gains from trade as a function of λ_{jj} , under our value of θ and under a tariff-inclusive elasticity of -5 considered by ACR.²⁹ As expected, the gains from trade are substantially larger with our elasticity. For the United States, gains from trade are 5.27 percent for $\theta = -1$, compared to 1.0 percent for $\theta = -5$. The median welfare gain is 22.9 percent in a sample of 64 countries, compared to 4.2 percent implied by $\theta = -5$. Online Appendix Table B9 reports the gains from trade under $\theta = -1, -5$, and -10 for selected countries in the sample.

The blue bars in online Appendix Figure B7 report the gains from trade using the multisector ACR formula and our sector-specific elasticity values (Section IIIA). We benchmark these to the sector-specific trade elasticity estimates from Ossa (2015), who explores the properties of multisector 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 multisector formula imply average gains from trade of 26.7 percent, compared to 12.8 percent 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) and Alessandria, Choi, and Ruhl (2021). 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. The notion that the value of the partial trade elasticity matters for the size of the gains from trade of course applies to dynamic settings, even if there are no known analytical formulas. To illustrate this, online Appendix Figure B8 displays the gains from trade in the dynamic GE model used in Section VB and detailed in online Appendix D. Because of the transition path, there is no unique way of calculating the gains from trade, as the total welfare change depends on the time path of trade costs. We start from the autarky equilibrium and implement a one-time unexpected permanent change in trade costs large enough to deliver a new steady state that matches the current level of trade. We then record the change in welfare between autarky and trade, taking into account the transition path to the new steady state. As in the ACR formula application in Figure 7, in the

²⁹We use data from the 2006 World Input-Output Database.




FIGURE 7. GAINS FROM TRADE

Notes: This figure displays the gains from trade as a function of the domestic absorption share 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 share from the 2006 World Input-Output Database over 43 countries.

dynamic model the trade elasticity matters a great deal for the magnitude of the gains from trade, with a lower elasticity producing larger gains.³⁰

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 multisector variant of the formula, that aligns more closely 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.

VI. Conclusion

We develop a novel method to estimate the trade elasticity, a key parameter in virtually all models in international economics. Our main contributions are to (i) tackle the endogeneity problem that tariffs and trade flows are jointly determined, and (ii) estimate the full time path of trade elasticities at different horizons. Our main

³⁰The dynamic gains in Appendix Figure B8 are larger than the ACR gains in Figure 7. This is in part due to the different interpretation of a long-run elasticity of 2 in the two settings. In static models, since there is no time dimension, a long-run elasticity of 2 is rationalized by setting $\sigma = 2$ in an Armington/Krugman setting. In the dynamic model, σ governs the short-run response of trade flows to tariffs, and in the long run the trade elasticity keeps increasing because of the adjustments to the mass of firms over the transition. Setting $\sigma = 1.1$ in the dynamic model appears to generate even larger gains from trade.

findings are that the trade elasticity falls from about -0.76 in the short run to about -2 in the long run. It takes seven to ten years for the point estimate to stabilize at the long-run value. While our estimation approach is not specific to a particular theoretical framework, we relate our empirical strategy and results to several theoretical applications. Our finding that the trade elasticity differs by horizon and converges to the “long-run” after about seven to ten years implies substantial adjustment costs to changing trade flows. The long-run estimates imply that the welfare-relevant trade elasticity is significantly smaller than conventional wisdom in the literature, suggesting the welfare gains from trade are larger than previously thought.

REFERENCES

- Adão, Rodrigo, Arnaud Costinot, and Dave Donaldson.** 2017. “Nonparametric Counterfactual Predictions in Neoclassical Models of International Trade.” *American Economic Review* 107 (3): 633–89.
- Alessandria, George, Costas Arkolakis, and Kim J. Ruhl.** 2021. “Firm Dynamics and Trade.” *Annual Review of Economics* 13 (1): 253–80.
- Alessandria, George, and Horag Choi.** 2014. “Establishment Heterogeneity, Exporter Dynamics, and the Effects of Trade Liberalization.” *Journal of International Economics* 94 (2): 207–23.
- Alessandria, George, and Horag Choi.** 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.** 2021. “Trade Adjustment Dynamics and the Welfare Gains from Trade.” *Journal of International Economics* 131: 103458.
- Anderson, James E.** 1979. “A Theoretical Foundation for the Gravity Equation.” *American Economic Review* 69 (1): 106–16.
- Anderson, James E., and Eric van Wincoop.** 2003. “Gravity with Gravitas: A Solution to the Border Puzzle.” *American Economic Review* 93 (1): 170–92.
- Anderson, James, and Eric van Wincoop.** 2004. “Trade Costs.” *Journal of Economic Literature* 42 (3): 691–751.
- Antràs, Pol, and Gerard Padró i Miquel.** 2011. “Foreign Influence and Welfare.” *Journal of International Economics* 84 (2): 135–48.
- Arkolakis, Costas.** 2010. “Market Penetration Costs and the New Consumers Margin in International Trade.” *Journal of Political Economy* 118 (6): 1151–99.
- 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.” Unpublished.
- Atkeson, Andrew, and Ariel Burstein.** 2008. “Pricing-to-Market, Trade Costs, and International Relative Prices.” *American Economic Review* 98 (5): 1998–2031.
- Auclert, Adrien, Matthew Rognlie, Martin Souchier, and Ludwig Straub.** 2021. “Exchange Rates and Monetary Policy with Heterogeneous Agents: Sizing up the Real Income Channel.” NBER Working Paper 28872.
- Backus, David K., Patrick J. Kehoe, and Finn E. Kydland.** 1992. “International Real Business Cycles.” *Journal of Political Economy* 100 (4): 745–75.
- Backus, David K., Patrick J. Kehoe, and Finn E. Kydland.** 1994. “Dynamics of the Trade Balance and the Terms of Trade: The J-Curve?” *American Economic Review* 84 (1): 84–103.
- Bagwell, Kyle, and Robert W. Staiger.** 2016. *Handbook of Commercial Policy*, Vol. 1. Amsterdam: 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.” Unpublished.

- Boehm, Christoph, Andrei Levchenko, and Nitya Pandalai-Nayar.** 2023. "Replication Data for: The Long and Short (Run) of Trade Elasticities." American Economic Association [publisher], Inter-university Consortium for Political and Social Research [distributor]. <https://doi.org/10.3886/E182781V1>.
- Bown, Chad P.** 2011. "Taking Stock of Antidumping, Safeguards and Countervailing Duties, 1990–2009." *World Economy* 34 (12): 1955–98.
- 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: 261–73.
- Bown, Chad P., and Meredith A. Crowley.** 2016. "The Empirical Landscape of Trade Policy." In *Handbook of Commercial Policy*. Vol. 1, edited by Kyle Bagwell and Robert W. Staiger, 3–108. Amsterdam: Elsevier-North Holland.
- 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): 269–83.
- 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): 123–27.
- Burstein, Ariel, and Marc Melitz.** 2013. "Trade Liberalization and Firm Dynamics." In *Advances in Economics and Econometrics Tenth World Congress. Applied Economics*, Vol. 2, edited by Daron Acemoglu, Manuel Arellano, and Eddie Dekel, 283–328. 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.
- 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): 431–49.
- Cavallo, Alberto, Gita Gopinath, Brent Neiman, and Jenny Tang.** 2021. "Tariff Pass-through at the Border and at the Store: Evidence from US Trade Policy." *American Economic Review: Insights* 3 (1): 19–34.
- Centre d'Études Prospectives et d'Informations Internationales.** 2022a. "BACI Database." Centre d'Études Prospectives et d'Informations Internationales. http://www.cepii.fr/CEPII/en/bdd_modele/bdd_modele_item.asp?id=37 (accessed April 15, 2020).
- Centre d'Études Prospectives et d'Informations Internationales.** 2022b. "GeoDist database." Centre d'Études Prospectives et d'Informations Internationales. http://www.cepii.fr/CEPII/en/bdd_modele/bdd_modele.asp (accessed November 8, 2022).
- Coeurdacier, Nicolas.** 2009. "Do Trade Costs in Goods Market Lead to Home Bias in Equities?" *Journal of International Economics* 77 (1): 86–100.
- Cole, Harold L., and Maurice Obstfeld.** 1991. "Commodity Trade and International Risk Sharing: How Much Do Financial Markets Matter?" *Journal of Monetary Economics* 28 (1): 3–24.
- Correia, Sergio.** 2015. "Singletons, Cluster-Robust Standard Errors and Fixed Effects: A Bad Mix." Unpublished.
- 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 Elhanan Helpman, Dalia Marin and Thierry Verdier, 107–42. 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*, Vol. 4, edited by Gita Gopinath, Elhanan Helpman, and Kenneth Rogoff, 197–261. Amsterdam: Elsevier.
- Donaldson, Dave.** 2018. "Railroads of the Raj: Estimating the Impact of Transportation Infrastructure." *American Economic Review* 108 (4-5): 899–934.
- Drozd, Lukasz A., Sergey Kolbin, and Jaromir B. Nosal.** 2021. "The Trade-Comovement Puzzle." *American Economic Journal: Macroeconomics* 13 (2): 78–120.
- 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): 1741–79.

- 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): 135–50.
- Fitzgerald, Doireann, and Stefanie Haller.** 2018. "Exporters and Shocks." *Journal of International Economics* 113 (C): 154–71.
- Fitzgerald, Doireann, Stefanie Haller, and Yaniv Yedid-Levi.** 2016. "How Exporters Grow." NBER Working Paper 21935.
- Fontagné, Lionel, Houssein Guimbard, and Gianluca Orefice.** 2022. "Tariff-Based Product-Level Trade Elasticities." *Journal of International Economics* 137: 103593.
- Fontagné, Lionel, Philippe Martin, and Gianluca Orefice.** 2018. "The International Elasticity Puzzle Is Worse than You Think." *Journal of International Economics* 115: 115–29.
- Gallaway, Michael P., Christine A. McDaniel, and Sandra A. Rivera.** 2003. "Short-Run and Long-Run Industry-Level Estimates of U.S. Armington Elasticities." *North American Journal of Economics and Finance* 14 (1): 49–68.
- Gaulier, Guillaume, and Soledad Zignago.** 2010. "BACI: International Trade Database at the Product-Level. The 1994–2007 Version." Centre d'Études Prospectives et d'Informations Internationales Working Paper 2010–23.
- Gawande, Kishore, Pravin Krishna, and Michael J. Robbins.** 2006. "Foreign Lobbies and U.S. Trade Policy." *Review of Economics and Statistics* 88 (3): 563–71.
- Giri, Rahul, Kei-Mu Yi, and Hakan Yilmazkuday.** 2021. "Gains from Trade: Does Sectoral Heterogeneity Matter?" *Journal of International Economics* 129: 103429.
- Hamilton, James D.** 1994. *Time Series Analysis*. Princeton: 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, 131–95. Amsterdam: Elsevier.
- 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): 858–76.
- Heathcote, Jonathan, and Fabrizio Perri.** 2002. "Financial Autarky and International Business Cycles." *Journal of Monetary Economics* 49 (3): 601–27.
- Heathcote, Jonathan, and Fabrizio Perri.** 2013. "The International Diversification Puzzle Is Not as Bad as You Think." *Journal of Political Economy* 121 (6): 1108–59.
- Hertel, Thomas, David Hummels, Maros Ivanic, and Roman Keeney.** 2007. "How Confident Can We Be of CGE-Based Assessments of Free Trade Agreements?" *Economic Modelling* 24 (4): 611–35.
- Hooper, Peter, Karen Johnson, and Jaime Marquez.** 2000. "Trade Elasticities for the G-7 Countries." *Princeton Studies in International Economics* 87.
- Hummels, David.** 2001. "Toward a Geography of Trade Costs." Unpublished.
- Huo, Zhen, Andrei A. Levchenko, and Nitya Pandalai-Nayar.** 2020. "International Comovement in the Global Production Network." NBER Working Paper 25978.
- Imbs, Jean, and Isabelle Mejean.** 2015. "Elasticity Optimism." *American Economic Journal: Macroeconomics* 7 (3): 43–83.
- Imbs, Jean, and Isabelle Mejean.** 2017. "Trade Elasticities." *Review of International Economics* 25 (2): 383–402.
- Johnson, Robert C.** 2014. "Trade in Intermediate Inputs and Business Cycle Comovement." *American Economic Journal: Macroeconomics* 6 (4): 39–83.
- Jordà, Òscar.** 2005. "Estimation and Inference of Impulse Responses by Local Projections." *American Economic Review* 95 (1): 161–82.
- Kehoe, Timothy J., and Kim J. Ruhl.** 2013. "How Important Is the New Goods Margin in International Trade?" *Journal of Political Economy* 121 (2): 358–92.
- Khan, Shafaat Yar, and Armen Khederlarian.** 2021. "How Does Trade Respond to Anticipated Tariff Changes? Evidence from NAFTA." *Journal of International Economics* 133: 103538.
- Kose, M. Ayhan, and Kei-Mu Yi.** 2006. "Can the Standard International Business Cycle Model Explain the Relation between Trade and Comovement?" *Journal of International Economics* 68 (2): 267–95.
- 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: 124–46.

- Leibovici, Fernando, and Michael E. Waugh.** 2019. "International Trade and Intertemporal Substitution." *Journal of International Economics* 117: 158–74.
- Lerner, Abba.** 1944. *The Economics of Control: Principles of Welfare Economics*. New York: Macmillan Co.
- Marshall, Alfred.** 1923. *Money, Credit and Commerce*. London: Macmillan and Co.
- Mayer, Thierry, and Soledad Zignago.** 2011. "Notes on CEPII's Distances Measures: The GeoDist Database." Unpublished.
- Melitz, Marc J.** 2003. "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity." *Econometrica* 71 (6): 1695–1725.
- Obstfeld, Maurice, and Kenneth Rogoff.** 2005. "Global Current Account Imbalances and Exchange Rate Adjustments." *Brookings Papers on Economic Activity* 36 (1): 67–123.
- Ossa, Ralph.** 2015. "Why Trade Matters after All." *Journal of International Economics* 97 (2): 266–77.
- 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): 416–35.
- Ruhl, Kim J.** 2008. "The International Elasticity Puzzle." Unpublished.
- Ruhl, Kim J., and Jonathan L. Willis.** 2017. "New Exporter Dynamics." *International Economic Review* 58 (3): 703–26.
- Shapiro, Joseph S.** 2016. "Trade Costs, CO, and the Environment." *American Economic Journal: Economic Policy* 8 (4): 220–54.
- 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.
- Soderbery, Anson.** 2018. "Trade Elasticities, Heterogeneity, and Optimal Tariffs." *Journal of International Economics* 114: 44–62.
- Steinberg, Joseph B.** 2022. "Export Market Penetration Dynamics." Unpublished.
- 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): 138–60.
- United Nations Conference on Trade and Development (UNCTAD).** 1995–2018. "UNCTAD TRAINS Database." UNCTAD. <http://wits.worldbank.org/WITS/WITS/Default-A.aspx?Page=Default> (accessed May 5, 2020).
- World Bank.** 2021. "Temporary Trade Barriers Database." World Bank. <https://www.worldbank.org/en/data/interactive/2021/03/02/temporary-trade-barriers-database> (accessed February 18, 2022).
- World Trade Organization.** 2014. *World Tariff Profiles 2014*. Geneva: World Trade Organization, United Nations Conference on Trade and Development, and International Trade Centre.
- World Trade Organization Integrated Database and Consolidated Tariff Schedules.** 1995–2018. "WTO Tariff Database." World Trade Organization. <http://stats.wto.org> (accessed February 8, 2022).
- Yilmazkuday, Hakan.** 2019. "Estimating the Trade Elasticity over Time." *Economics Letters* 183: 108579.